When is Information Sufficient for Action? Search with Unreliable Yet Informative Intelligence

Michael P. Atkinson,* Moshe Kress,† Rutger-Jan Lange‡

January 13, 2016

Abstract

We analyze a variant of the whereabouts search problem, in which a searcher looks for a target hiding in one of n possible locations. Unlike in the classic version, our searcher does not pursue the target by actively moving from one location to the next. Instead, the searcher receives a stream of intelligence about the location of the target. At any time, the searcher can engage the location he thinks contains the target or wait for more intelligence. The searcher incurs costs when he engages the wrong location, based on insufficient intelligence, or waits too long in the hopes of gaining better situational awareness, which allows the target to either execute his plot or disappear. We formulate the searcher's decision as an optimal stopping problem and establish conditions for optimally executing this search-and-interdict mission.

Keywords: whereabouts search, optimal stopping, multinomial selection

---

*mpatkins@nps.edu Operations Research Department, Naval Postgraduate School

†mkress@nps.edu Operations Research Department, Naval Postgraduate School

‡rutger-jan.lange@cantab.net Finance Department, VU University Amsterdam & School of Economics, Erasmus University Rotterdam

1 Introduction

Operation Neptune Spear led to the capture and elimination of Osama bin Laden by the US in 2011. While US intelligence agencies had continuously collected information regarding his whereabouts, the dilemma was when to act. Raiding a wrong location, based on insufficient or false information, would cause collateral damage, diplomatic blowback and loss of intelligence assets. On the other hand, waiting too long for more information could result in bin Laden escaping. The dilemma between "act now" or "wait and see" was acute but fortunately was resolved successfully in this case. Another example of such a dilemma concerns a "ticking bomb" scenario (Kaplan, 2012). In this scenario a hiding terrorist plots to attack a target (e.g., a suicide bomber), and the authorities must race to stop the attack. A final example involves an operation to rescue hostages held by an insurgency group. The insurgents may kill the hostages (e.g., in an escape attempt) if the authorities delay the operation for too long. However, a failed rescue attempt may alert the insurgents resulting in the deaths of the hostages. Many military, law enforcement, and intelligence investigations face a similar tradeoff decision concerning timing and cost of premature action.

Motivated by the aforementioned examples, we consider a search situation called the *whereabouts search* problem (Kadane, 1971; Stone, 1975). In its simplest form, a target lies hidden in one of n cells, where p_i is the probability that the target resides in cell i, $\sum_{i=1}^{n} p_i = 1$, and c_i is the cost of searching cell i. The searcher examines one cell at a time and the search is error-free; if a cell contains the target, the searcher will detect it. The objective is to find a search strategy – an order in which to search the cells – to minimize the expected total search cost. Several variations of this problem include, among others, situations where a search is subject to error (Kress et al., 2008; Wilson et al., 2011), the target moves (Komiya et al., 2006) or acts strategically (An et al., 2013) , and multiple targets arrive and disappear in a random fashion (Szechtman et al., 2008). However, all of the aforementioned cases share the same definition of a *strategy*, namely, a search sequence for an active searcher.

In this paper, we consider the same physical description of the whereabouts problem: a

single static target hidden in one of n cells. However, the operational setting is different in two major aspects: (a) the searcher does not actively search the cells but instead relies on occasional pieces of intelligence of the form "cell i contains the target", and (b) the search mission is time-critical. The searcher does not control the arrival rate of intelligence, and an intelligence item may be wrong. At a certain point the searcher may choose a cell to engage in the hope of interdicting the target. If the searcher chooses the wrong cell, he incurs a cost comprising collateral damage, loss of intelligence assets, political ramifications, etc.

We describe the problem in Section 2 and formulate the mathematical model in Section 3. The cases of $n = 2$ and $n = \infty$ appear in Sections 4 and 5, respectively. Section 6 examines the optimal strategy when $2 < n < \infty$. We present numerical illustrations in Section 7. Section 8 discusses extensions. All proofs appear in the Appendix.

2   The Problem

A searcher wants to interdict a target, residing in one of n possible cells, before some event occurs. Such an event would be, for example, the disappearance of bin Laden from a certain region or an execution of a terror plot, which we use as our reference scenario. An attack occurs when the plot fully *matures*, and the plotting time is exponentially distributed with mean $1/\mu$ (a similar assumption appears in Kaplan (2010)). While the searcher may have some initial notion regarding the target's location based on exogenous intelligence, we will often focus on the case where there is none: the uniform prior location distribution.

Independent intelligence items from human informants, intercepted communications and interrogations of the form "cell i contains the target" arrive according to a Poisson process with rate λ. The searcher has no control over the timing or content of the items. Thus, scheduled sensor cues (e.g., RADAR, SONAR, images, videos) from cells do not apply here. Although our model applies to a variety of intelligence sources, we use, as a reference setting, human informants who provide tips. For most of our analysis the parameters μ and λ only

appear via the *intensity ratio* $\rho = \mu/\lambda$.

If, following a certain number of tips, the searcher decides to engage a specific cell, the search ends, even if the searcher chooses incorrectly. If the searcher engages the correct cell, the target is interdicted. However, if the searcher engages the wrong cell, then the target realizes that he is being hunted, and therefore immediately executes his (not fully mature) plot before the searcher finds him. In Section 8.1 we consider a variant where the target only executes mature plots and the searcher continues obtaining intelligence and engaging cells until he either finds the target or the target attacks.

The searcher desires to minimize the expected cost of two possible negative outcomes: (a) engaging a wrong cell, or (b) execution of a mature attack by the target. The costs of (a) and (b) are c and d, respectively. The false positive cost c comprises collateral damage resulting from engaging an innocent cell and the (possible) cost of a premature attack. We neither need nor make any assumption regarding the relative values of c and d. Because the results to follow only depend on the the *cost-ratio* $\alpha = d/c$, we assume, without loss of generality, that $c = 1$ and $d = \alpha$.

A tip specifies the correct cell with probability q. We often refer to q as the informant's *reliability*. Informants are neither clueless nor malevolent, that is, $q > \frac{1}{n}$. If the informant provides an incorrect tip (with probability $1 - q$), then the error is uniform; the informant specifies each one of the $n - 1$ incorrect cells with equal probability.

The question is: *when should the searcher engage a cell?* We have here a "race" between the flow of tips and the time of attack. On the one hand, the searcher wants to receive as many tips as possible to reduce his uncertainty about the target's location. On the other hand, this "wait and see" approach may lead to the target attacking before the searcher has the chance to do so. If the searcher instead rushes to engage a cell, the likelihood of a false positive error increases. The searcher knows the values of all the parameters involved in this process: n, q, α, and ρ.

This search problem is an example of an optimal stopping problem (Chow et al., 1971;

4

Shiryaev, 2007; Ferguson, 2004). Wald and Wolfowitz (1948) examine a similar problem in their work on the sequential probability ratio test. They show that the decision between selecting a hypothesis and receiving another observation is optimally determined by a threshold policy. In our model, when $n = 2$ cells, we find a similar threshold result (see Section 4), which does not hold for $n > 2$. For $n > 2$, our problem can be framed as a higher dimensional stopping problem. Lange (2012) examines optimal stopping of an n-dimensional Brownian motion, and shows that the continuation region is generally also n-dimensional. While standard one-dimensional techniques do not apply, he shows that the continuation region can be found by reformulating the problem as a free-boundary problem in n dimensions.

When $n > 2$ cells, our problem relates to the family of multinomial selection problems (Kim and Nelson, 2006) in which an observation specifies the "winner" among n competing alternatives. A decision maker may either observe a fixed number of samples before choosing the best option (Bechhofer et al., 1959) or may dynamically decide, after each observation, whether to pick an alternative or receive another observation (Ramey and Alam, 1979). Most formulations desire to achieve a lower bound on the probability of choosing the correct alternative, provided certain conditions about the system hold. These conditions usually relate to the relationship between the true probabilities of the two best alternatives (Chen, 1988). A good survey of the techniques used in multinomial selection problems appears in Vieira et al. (2014). Most selection problems assume a deterministic number of observations. In our problem the number of tips is random because the time until the plot matures is random. We found only two multinomial selection papers that examine a random maximum number of observations (Frazier and Yu, 2007; Dayanik and Yu, 2013). The model in Frazier and Yu (2007) considers only the $n = 2$ case and allows for a general stochastic deadline, which is analogous to the time until the attack occurs in our model. The approach in Dayanik and Yu (2013) does allow for $n > 2$ alternatives. It focuses on neuroscience applications and considers a cost-rate, as opposed to total cost in our model.

Finally, note that our model has one decision maker, the searcher. One could view the

problem as having three strategic players: the searcher, the target, and the informant. We consider here a simpler, yet, we believe, a realistic situation where the target does not really know the searcher's operational options and the informant is incentivized by the searcher to do the best he can. One could develop a two-player Markov game between the searcher and target similar to the Inspection Game (see Chapter 4 of Washburn (2014)). However, the formulation would quickly become unwieldy because one would need to specify not only the intelligence picture of each player, but also each player's *perceived* intelligence picture.

3   Mathematical Preliminaries

The decision to *engage* a cell or *wait* for more tips depends on the expected cost of each option. In this section we develop the mathematical building blocks to compute these expected costs. Two factors determining the expected costs are *Location* probability, which specifies the likelihood that cell i contains the target, and *Pointing* probability that specifies the likelihood that the next tip points at cell i. In Section 3.1 we compute these probabilities, and in Section 3.2 we use these probabilities to derive the expected costs.

3.1   Location and Pointing Probabilities

Let $p = (p_1, ..., p_n)$ denote the current location probabilities and let $\tilde{p}$ denote the initial location probabilities, before the first tip. Let s_i be the number of tips thus far specifying cell i as the target's location, and $s = (s_1, ..., s_n)$. In this subsection we assume that $s_1 \geq ... \geq s_n$. The location probability of cell i given s is:

$$p_i(s) = \mathbf{P}[\text{target in } i \mid s] = \frac{\mathbf{P}[s \mid \text{target in } i]\tilde{p}_i}{\sum_{j=1}^{n} \mathbf{P}[s \mid \text{target in } j]\tilde{p}_j}. \tag{1}$$

An informant points to the correct cell with probability q and a specific incorrect cell with probability $\frac{1-q}{n-1}$. Thus, utilizing the multinomial nature of s, we have

$$
\begin{aligned}
\mathbf{P}[s \mid \text{target in } i] &= \frac{(\sum_k s_k)!}{\prod_k s_k!} q^{s_i} \left(\frac{1-q}{n-1}\right)^{\sum_{k \neq i} s_k} \\
&= \frac{(\sum_k s_k)!}{\prod_k s_k!} \left(\frac{1-q}{n-1}\right)^{\sum_{k=1}^n s_k} \left(\frac{q}{\frac{1-q}{n-1}}\right)^{s_i} \\
&= \frac{(\sum_k s_k)!}{\prod_k s_k!} \left(\frac{1-q}{n-1}\right)^{\sum_{k=1}^n s_k} \gamma^{s_i},
\end{aligned}
\tag{2}
$$

where

$$
\gamma = \frac{q}{\frac{1-q}{n-1}}.
\tag{3}
$$

Note that only the γ^{s_i} portion of (2) depends on i. This is a direct consequence of our assumption that each wrong cell is equal likely to be pointed at. When we substitute (2) back into (1), most terms cancel, and the location probability simplifies to:

$$
p_i(s) = \frac{\gamma^{s_i} \tilde{p}_i}{\sum_{j=1}^n \gamma^{s_j} \tilde{p}_j}.
\tag{4}
$$

Note from equation (4) that $p_i(s)$ is invariant to additive shifts in s. If $\hat{s}$ is such that $\hat{s}_i = s_i + L$ for some integer L, then $p_i(\hat{s}) = p_i(s)$. Specifically, if we set $L = -s_n = -\min(s)$ and use $\hat{s}_i = s_i - s_n$, then we can write $\hat{s}_i = \sum_{j=i}^{n-1} \Delta_j$, where $\Delta_j = s_j - s_{j+1} \geq 0$. Therefore, $p_i(s)$ is uniquely determined by the tip-differentials $\Delta_j, j = 1, ..., n-1$.

While s or Δ are natural state vectors, it is simpler to use the location probabilities $p = (p_1, ..., p_n)$ as the state vector for most of the mathematical analysis in Sections 4–6.

Specifically, if the next tip points to cell i, then the updated probability $p_j^{(+i)}$ for cell j is:

$$p_j^{(+i)} = \begin{cases} \frac{\gamma p_i}{\gamma p_i + (1 - p_i)} & \text{if } j = i \\[3mm] \frac{p_j}{\gamma p_i + (1 - p_i)} & \text{if } j \neq i. \end{cases} \tag{5}$$

Recall that according to our assumption $q > \frac{1}{n}$ and therefore $\gamma > 1$. Consequently, a tip pointing to cell i increases the posterior location probability of cell i ($p_i^{(+i)} \geq p_i$) and decreases the posterior probability of other cells ($p_j^{(+i)} \leq p_j$ for $j \neq i$).

We next define $B(p)$ as the set of cells with the highest location probability:

$$B(p) = \{i \ : \ p_i = \max_j p_j, \ 1 \leq i \leq n\}. \tag{6}$$

The following proposition defines a lower bound on $\max_j p_j$. The proof appears in Appendix A.

Proposition 1. *If $|B(p)| = 1$ and the prior distribution for the target's location is uniform, then $\max_j p_j \geq q$.*

Next we consider the pointing probability $r_i(p)$ that the next tip points to cell i, given the current location probabilities p:

$$r_i(p) \equiv \mathbf{P}[\text{informant says } i \mid p] = \sum_{k=1}^n \mathbf{P}[\text{informant says } i \mid p, \text{target in } k]\mathbf{P}[\text{target in } k \mid p]$$

$$= qp_i + \frac{1-q}{n-1}\sum_{k \neq i} p_k = qp_i + \frac{1-q}{n-1}(1 - p_i). \tag{7}$$

Inspection of (7) reveals that $r_i(p) \in [\frac{1-q}{n-1}, q]$. Thus, a tip may point at a cell other than i, even if p_i is close to 1, if $q << 1$. Note also that $r_i(p)$ only depends on p_i, it does not depend upon how the remaining $(1 - p_i)$ probability mass is spread among the other $n - 1$ cells.

3.2 Expected Cost

Define $C(p)$ as the expected cost if the searcher acts optimally in state p. Since an optimal stopping problem is a dynamic programming problem (Chow et al., 1971), we compute $C(p)$ by comparing the expected costs of two decisions: *engage* or *wait*. That is,

$$C(p) = \min \Bigg(\text{expected cost if the searcher engages a cell,}$$
$$\frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\text{expected cost after receiving the next tip} \Bigg). \qquad (8)$$

If the searcher decides to wait, the target may attack before the searcher receives the next tip. In that case, which happens with probability $\frac{\rho}{1+\rho}$, the mature attack produces a cost of α. If the next tip arrives before the target's attack, the system transitions, and we assume the searcher behaves optimally in the future. Next we compute the expected costs of the two possible options: *engage* or *wait*.

If the searcher decides to engage cell j while in state p, the expected cost is $1 - p_j$. Obviously, the searcher should engage a cell in $B(p)$; the searcher can use any tie-breaking mechanism if $B(p)$ contains multiple cells. To simplify notation, we henceforth assume, without loss of generality, that $p_1 \geq ... \geq p_n$. Therefore $B(p)$ contains cell 1 and

$$\mathbf{E}[\text{Cost if searcher decides to engage } | p] = 1 - \max_j p_j = 1 - p_1. \qquad (9)$$

If the searcher decides to wait, and an informant next points to cell i, then p transitions to $p^{(+i)}$ according to equation (5). The informant points to cell i with probability $r_i(p)$, and the searcher will incur an expected cost of $C(p^{(+i)})$ if this occurs. Putting these pieces together,

we have

$$\mathbf{E}[\text{Cost if waiting for and receiving the next tip} \mid p] = \sum_{i=1}^{n} \mathbf{P}[\text{informant says } i \mid p] C(p^{(+i)})$$

$$= \sum_{i=1}^{n} r_i(p) C(p^{(+i)}). \tag{10}$$

Moving to the general case, we combine equations (8), (9), and (10) to produce the complete cost function:

$$C(p) = \min\left(1 - p_1, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p) C(p^{(+i)})\right). \tag{11}$$

If the searcher is indifferent between engaging and waiting, he engages. In Appendix B we present characteristics of $C(p)$, such as its concavity. As most of these results are fairly intuitive (e.g., $C(p)$ decreases if the informant next points to cell 1), we defer this discussion to the Appendix.

4 The Case of Two Cells

Arguably, the simpler the form of the optimal policy the more attractive it is operationally. One such simple form is a threshold policy: the searcher engages if and only if $p_1 \geq \tau$ for some threshold τ (recall we assume that $p_1 \geq p_2$). The next corollary follows from the convexity of the engage region (see Proposition EC.2 in Appendix B).

Corollary 1. *For $n = 2$, the searcher should engage if and only if $p_1 \geq \tau$ for some threshold $\tau \in [0.5, 1)$.*

We prove this corollary in Appendix C. While there is an explicit expression for the threshold τ, its derivation is cumbersome and therefore we defer most of its details to Appendix D.

A necessary and sufficient condition to engage in all states (i.e., $\tau = 0.5$) is

$$\frac{1}{2} \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q. \tag{12}$$

If condition (12) does not hold then $\tau > 0.5$. See Appendix D for the general expression for τ when $\tau > 0.5$. The implication is straightforward; if damage from a mature attack exceeds the false positive cost ($\alpha \geq 1$) and the informant has low reliability ($q \approx 0.5$), the searcher should always engage. The benefits from future tips are small, and the risk of waiting is high.

To derive τ we leverage off the rich results related to the gambler's ruin problem. Denote $\tilde{p}$ as the the prior state before the arrival of the $s_1 + s_2$ tips. Using equation (4) we transform $\tilde{p}$ to p :

$$p_1 = \frac{\gamma^{s_1}\tilde{p}_1}{\gamma^{s_1}\tilde{p}_1 + \gamma^{s_2}(1-\tilde{p}_1)} = \frac{\gamma^{s_1-s_2}\tilde{p}_1}{\gamma^{s_1-s_2}\tilde{p}_1 + (1-p_1)} \tag{13}$$

$$p_2 = \frac{\gamma^{s_2}(1-\tilde{p}_1)}{\gamma^{s_1}\tilde{p}_1 + \gamma^{s_2}(1-\tilde{p}_1)} = \frac{1-\tilde{p}_1}{\gamma^{s_1-s_2}\tilde{p}_1 + (1-\tilde{p}_1)}. \tag{14}$$

To update the probabilities we only need to know the tip-differential $s_1 - s_2$. We model $\Delta \equiv s_1 - s_2$ as a random walk. For a given prior $\tilde{p}$, we can transform the threshold policy from the real number τ to two non-negative integers $A(\tilde{p}, \tau)$ and $B(\tilde{p}, \tau)$ such that the searcher waits as long as $-B(\tilde{p}, \tau) < \Delta < A(\tilde{p}, \tau)$. If Δ first hits $A(\tilde{p}, \tau)$ $(-B(\tilde{p}, \tau))$, the searcher engages cell 1 (cell 2). This approach facilitates the use of gambler's ruin machinery to compute relevant parameters (See Appendix D for details).

It is difficult to gain much insight about the optimal threshold τ using purely analytic approaches. Thus, we illustrate its behavior using several figures. Figure 1 presents how the threshold τ varies with informant reliability q for fixed cost-ratio α and intensity-ratio ρ. As we move from Figure 1a to Figure 1c, we increase α from 0.5 to 2. Each curve on a figure corresponds to a fixed value of $\rho \in \{0.01, 0.1, 1\}$. The threshold τ is a non-decreasing function of q. A more reliable informant reduces the engage region and makes the searcher more likely to wait because future tips are more valuable. The threshold decreases as we increase either

11

α (mature attack becomes more costly) or ρ (mature attack becomes more imminent) and hence the engage region expands. In particular, in some situations with large α and/or large ρ, the searcher immediately engages regardless of the current state p or informant reliability q.

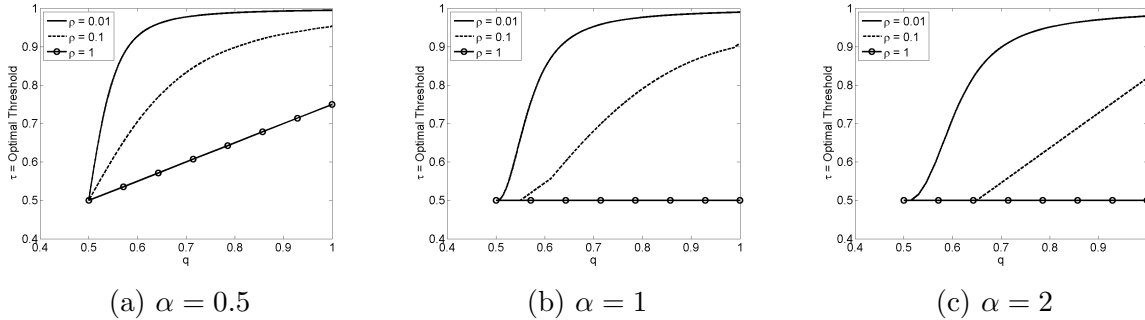(a) $\alpha = 0.5$ (b) $\alpha = 1$ (c) $\alpha = 2$

Figure 1: Engage threshold τ as a function of q for fixed combinations of $\rho \in \{0.01, 0.1, 1\}$ and $\alpha \in \{0.5, 1, 2\}$.

An interesting phenomenon relates to the expected number of tips received by the searcher when acting optimally. One would expect that this number will decrease as the informant becomes more reliable and therefore the searcher can reach the engage decision faster. Figure 2 demonstrates that this is not always the case. See Appendix E for the derivation of the expected number of tips. Assuming the search starts in the uniform state $p = (0.5, 0.5)$, Figures 2b and 2c show that if ρ is small (the inflow rate of tips is much larger than the attack rate) it is possible that the expected number of tips actually increases with q when the latter is small enough. This non-monotonicity results from two conflicting factors. On one hand, as q increases the threshold increases (see Figure 1), which suggests that the searcher may need more tips to reach the threshold. On the other hand, a larger q implies that the informant will point to the correct cell more frequently, which suggests that the searcher will reach the threshold following fewer tips. Specifically, for $q \approx 1$, the searcher will only need one tip. In general, the first or second factor may dominate depending upon the values of α, ρ, q. In most cases, when ρ is relatively large, the imminent attack dictates a swift action by the searcher, as shown in the dashed and $-\circ-$ curves, which are close to 0.

The jumps in Figure 2 occur when the optimal tip-differential changes by 1. For a fixed optimal tip-differential, the expected number of tips decreases as q increases because a more reliable informant will produce a stream of tips that reaches that tip-differential faster (probabilistically) than a less reliable informant.
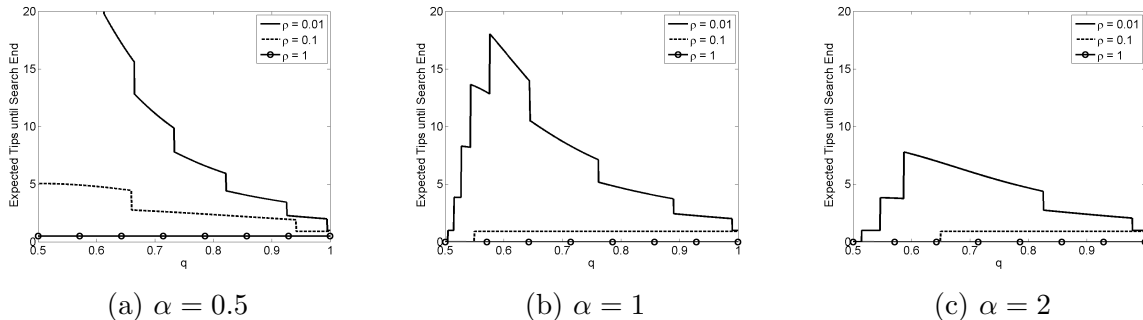

(a) $\alpha = 0.5$ (b) $\alpha = 1$ (c) $\alpha = 2$

Figure 2: Expected number of tips, starting from the uniform state $p = (0.5, 0.5)$, until the search ends as a function of q for fixed combinations of $\rho \in \{0.01, 0.1, 1\}$ and $\alpha \in \{0.5, 1, 2\}$. The search ends either when the searcher engages or when a mature attack occurs.

5   The Case of an Infinite Number of Cells

When n is very large and the cells are equally likely to contain the target, it is unlikely that the informant will point to the same incorrect cell twice. Thus, a second tip to the same cell should indicate that it is the correct one. In Appendix F.1 we make this argument more rigorous. If $n = \infty$ and the informant points twice to the same cell, then the searcher knows with certainty that this cell contains the target. We refer to the second tip to the same cell as the *confirming tip*. In Appendix F.2 we derive the optimal policy, which we summarize in the next Proposition.

Proposition 2. *The searcher will choose the lowest cost alternative among the following three options*

1. *Immediately engage any cell before receiving the first tip: cost is* 1

2. *Obtain one tip and engage the corresponding cell: cost is* $\frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-q)$;

13

3. *Wait for the confirming tip and then engage: cost is* $\alpha\left(1 - \left(\frac{q}{\rho+q}\right)^2\right)$.

Thus, the searcher should

$$choose\ option\ 1\ \ iff\ \ \alpha > 1 + \frac{q}{\rho},$$

$$choose\ option\ 3\ \ iff\ \ \alpha < 1 + \frac{q}{\rho}\frac{q - \rho^2 - \rho q - q^2}{\rho + 2q - q^2},$$

$$choose\ option\ 2\ \ otherwise.$$

The searcher causes collateral damage if he chooses option 1 because he engages the wrong cell. The cost for option 2 follows immediately from (11) as $p_1 = q$ after the tip. If the searcher chooses option 3, there is no collateral damage, but the target may execute the attack before confirming tip arrives.

Figure 3 illustrates what the searcher should do for different α, ρ pairs for $q \in \{0.1, 0.8\}$. The searcher chooses option 1 if the parameters lie above the solid curve, option 3 if the parameters lie below the dashed curve, and option 2 otherwise. The searcher is more likely to wait for the confirming tip for small α/ρ pairs and engage for large values. Not surprisingly the region in which option 2 is optimal increases as we increase q, as one tip provides significant information for larger values of q.
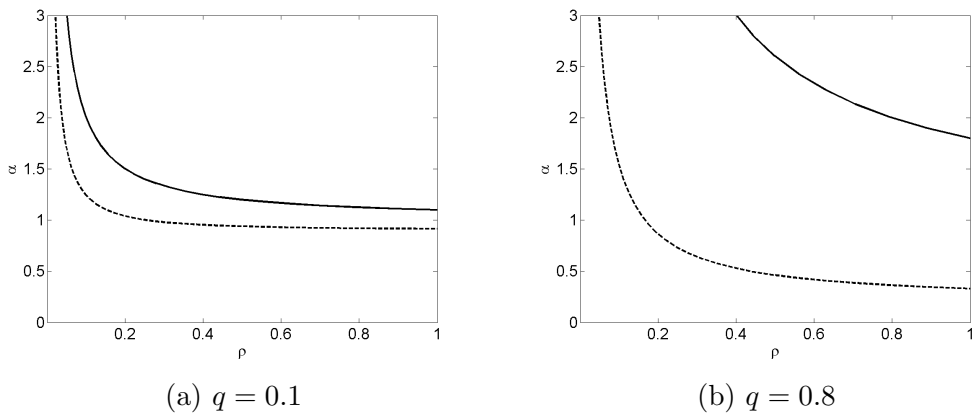

(a) $q = 0.1$ (b) $q = 0.8$

Figure 3: Searcher should engage for (ρ, α) lying above solid line, wait for the confirming tip if (ρ, α) lies below the dashed line, and engage after one tip for situations between the two curves.

The optimal strategy for the $n = \infty$ case suggests a heuristic for $n < \infty$, where the searcher chooses among the three options listed in Proposition 2. We compute the finite-n costs for these three options in Appendix F.3.1. Overall the heuristic performs very well and provides near optimal results in many situations, often even for small n. This heuristic generates a cost within 1% on average over many scenarios covering a variety of different parameter combinations. Unfortunately this heuristic only applies for the uniform state. Appendices F.3.2 and J.1 contain more details on the performance of this heuristic.

6 Policy for $2 < n < \infty$

Suppose that $q = 1$. In this case, the searcher either immediately engages cell 1, or he waits for the first tip and then engages the correct cell. In the former the expected cost is $(1 - p_1)$, and in the latter it is $\frac{\rho}{1+\rho}\alpha$. Thus, the searcher should engage now if and only if

$$p_1 \geq \frac{\rho}{1 + \rho}(1 - \alpha) + \frac{1}{1 + \rho}. \tag{15}$$

Condition (15) is sufficient to engage for any value of q. We derive this formally in Section 6.1. This observation leads to the following preliminary analysis for the case where $q < 1$ and the searcher has no prior information: $p_1 = \ldots = p_n = \frac{1}{n}$. In that case the searcher engages any cell before receiving a tip if $\frac{1}{n} \geq \frac{\rho}{1+\rho}(1 - \alpha) + \frac{1}{1+\rho}$. We call this situation a *blind engagement* because the searcher effectively shoots in the dark. If the searcher obtains one tip and engages the corresponding cell, then the initial state $p = (1/n, 1/n, \ldots 1/n)$ transitions to $p^{(+1)} = (q, (1 - q)/(n - 1), \ldots (1 - q)/(n - 1))$ (see equation (5)) and the expected cost is $\frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1 - q)$. Thus, if $(1 - 1/n) > \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1 - q)$ the searcher should wait. In

summary, we have

$$\text{if} \quad \frac{1}{n} > \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho} \quad \longrightarrow \quad \textit{blind engagement,} \tag{16}$$

$$\text{if} \quad \frac{1}{n} < \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q \quad \longrightarrow \quad \textit{wait.} \tag{17}$$

If $1/n$ falls between the two bounds, additional analysis is needed. Note the equivalence between condition (17) and the two-cell condition in (12). Conditions (16)–(17) suggest that if n is small, ρ is large (an imminent attack is likely), and α is large (damage from a mature attack exceeds the false positive cost), the searcher may optimally choose a cell uniformly at random before receiving any tips. Figure 4 presents the region in α, ρ space where the searcher chooses to wait rather than blindly engage (condition (17)) for different values of n and q. The wait region falls below the curves. For large n and a reliable informant, the searcher will wait for even reasonably large values of α and ρ. The curves look similar to those in Figure 3 for the $n = \infty$ case. The solid curve in Figure 3 corresponds to the thin dashed curve in in the northeastern portion of Figure 4, which represents the limiting case as $n \to \infty$.
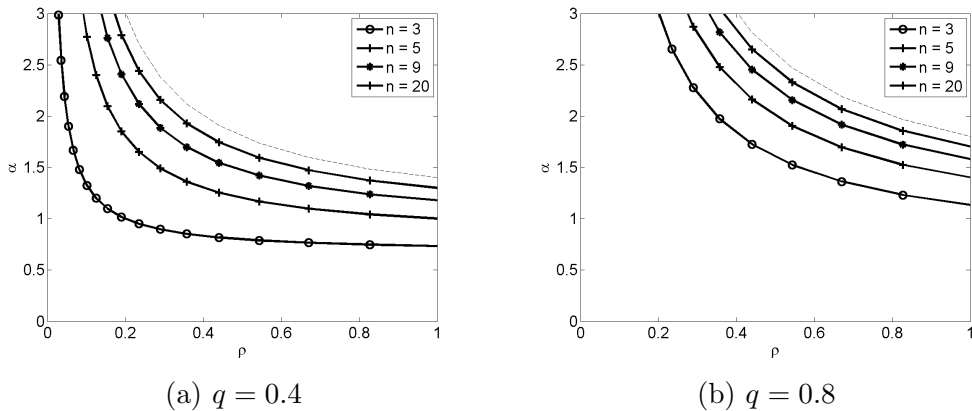

(a) $q = 0.4$ (b) $q = 0.8$

Figure 4: For the uniform state, the searcher should receive at least one tip if (ρ, α) lies below the curve.

We now turn to the general non-uniform state. Unlike the $n = 2$ case, there is no threshold policy for optimally responding to tips, as shown in the next example.

16

Example 2: let $q = 0.3$, $\alpha = 0.8$, $\rho = 1/9$. The searcher should engage in state $p = (0.316, 0.246, 0.246, 0.191)$ and should wait in state $\hat{p} = (0.366, 0.366, 0.134, 0.134)$. However, $0.316 = p_1 < \hat{p}_1 = 0.366$.

Example 2 suggests that the key factor driving the decision lies in the differential between the two cells with the highest probability. This type of result appears in many algorithms used for multinomial selection problems (Bechhofer et al., 1959; Ramey and Alam, 1979; Kim and Nelson, 2006). One might propose that the optimal policy takes a threshold form based on $p_1 - p_2$ or p_1/p_2. Unfortunately, the next example shows a threshold policy based on either of those two quantities is not optimal.

Example 3: Let $q = 0.42$, $\alpha = 0.5$, $\rho = 1$. The searcher should engage in state $p = (0.556, 0.384, 0.060)$ but should wait in state $\hat{p} = (0.512, 0.244, 0.244)$.

Our state space $\{p \mid p_1 \geq ... \geq p_n, \sum_{i=1}^{n} p_i = 1\}$ is an $n - 1$ dimensional closed convex set, and thus we should not be surprised that the optimal policy cannot be represented by a 1-dimensional subspace. As the optimal policy does not take on a simple form, we next present sufficient conditions to engage or wait. The searcher can use the conditions in this section as the basis for heuristic policies. We compare these heuristic policies to the optimal policy in Section 7.1 and Appendix J.

We derive the sufficient conditions by computing upper and lower bounds on the value of the second term of the cost function $C(p)$ in equation (11); the second term corresponds to the expected cost to wait. If the engage value $(1 - p_1)$ is less than or equal to this lower bound, then the searcher should engage in state p. If $(1 - p_1)$ exceeds the upper bound, then the searcher should wait in state p. If $(1 - p_1)$ lies between the lower bound and upper bound to wait, then we need to perform additional analysis or derive tighter bounds to determine the searcher's optimal decision.

We defer the construction of the upper and lower bounds to Appendix G. Rather than focus on the general structure of the bounds, we instead present several specific sufficient conditions to engage or wait in Sections 6.1 and 6.2, respectively. These conditions converge

to a necessary and sufficient condition to engage (see Proposition EC.7 in Appendix G). This allows us to theoretically approximate $C(p)$ to any desired precision and determine whether the searcher should engage or wait in state p. The computational feasibility depends upon ρ (see (EC.100)–(EC.101) in Appendix G). For $\rho \geq 0.1$, we can solve for $C(p)$ and the optimal decision in less than a second on most problems on a standard laptop for $n \sim 100$. However, for $\rho \leq 0.01$ the calculations can bog down or become intractable for $n \leq 10$.

6.1 Sufficient Conditions to Engage

In Appendix H we present several sufficient conditions to engage, including a family of conditions that converges to a necessary and sufficient condition. Here we focus on three conditions to engage that provide insight on the decision.

For our first bound we set $C(p^{(+i)}) = 0$ in (11). This assumes that the searcher knows the location of the target with certainty after receiving one tip. This best-case scenario produces a lower bound on the optimal cost $C(p)$ and yields condition (15), which we derived earlier by assuming $q = 1$. Combining Proposition 1 and condition (15) produces the following sufficient condition to engage

$$\text{engage if} \quad q \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}, \quad \text{for uniform prior and } |B(p)| = 1. \quad (18)$$

If condition (18) holds for the uniform prior case, then the searcher would receive at most one tip before engaging cell 1.

To derive a tighter, less conservative, sufficient condition to engage, we set $C(p^{(+i)}) = 0$ after two tips in (11) (rather than after one as assumed in (15)). In Appendix H.1 we show that if the following condition holds, then the searcher should engage cell 1.

$$p_1 \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}\left(\sum_{i=1}^{n} r_i(p)\left(\max\left(\max_j p_j^{(+i)}, \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}\right)\right)\right). \quad (19)$$

The right-hand side of (19), which depends now, through $r_i(p)$ and $p_j^{(+i)}$, on q is always smaller than the right-hand side of (15). We derive (15) from (11) by assuming $C(p^{(+i)}) = 0$, but we derive (19) from (11) by assuming

$$C(p^{(+i)}) = \min \left(1 - \max_j p_j^{(+i)}, \frac{\rho}{1+\rho} \alpha \right) \geq 0.$$

We conclude this subsection with a heuristic based on the threshold policy for the two-cell case, where cells $2, 3, \ldots, n$ are combined into an uber-cell. Accordingly, define a two-cell state $\hat{p}$ such that $\hat{p}_1 = p_1$ and $\hat{p}_2 = 1 - p_1 = \sum_{i=2}^n p_i$. If the searcher chooses to engage cell 1 when compared to the uber-cell, then the searcher should also engage cell 1 in the n-cell problem. We must modify q when moving from the n-cell problem to the two-cell problem to maintain the same γ, which captures informant effectiveness independent of n. Specifically, define $\hat{q} = \frac{\gamma}{1+\gamma}$, where γ applies to the original n-cell problem. If we denote $\tau(q, \alpha, \rho)$ as the optimal threshold for the two-cell problem, (see Proposition EC.4 of Appendix D), then we have the following condition:

$$\text{engage if} \quad p_1 \geq \tau(\hat{q}, \alpha, \rho). \tag{20}$$

6.2 Sufficient Conditions to Wait

Appendix I derives conditions to wait based on the common heuristic called the k-stage look-ahead rule. The searcher can receive at most k additional tips; after receiving the kth tip, the searcher must engage. Because the k-stage look-ahead rule restricts the searcher's strategy space, the policy will produce an upper bound on the cost function $C(p)$. Consequently, if the k-stage look-ahead policy recommends to wait, then the searcher should optimally wait. See Chapter 5.1 of Ferguson (2004) or 7.4 of Berger (1985) for more details on the k-stage look-ahead policy. This heuristic transforms the infinite horizon problem of solving for $C(p)$ in (11) to a finite horizon problem. For small values of k, backward induction provides a

computationally tractable approach. The k-stage look-ahead heuristic usually performs well in practice (Ferguson, 2004).

We now focus on a *myopic* policy where $k = 1$. In this case the searcher considers just two options: (1) engage cell 1, (2) wait for the next tip and then engage. Condition (17) corresponds to the myopic policy starting from the uniform state. More generally, if the searcher uses the myopic policy, he will engage cell 1 if

$$p_1 \geq \frac{\rho}{1 + \rho} (1 - \alpha) + \frac{1}{1 + \rho} \sum_{i=1}^{n} \max \left(q p_i, \frac{1 - q}{n - 1} p_1 \right). \tag{21}$$

See Appendix I.1 for the derivation of (21). If condition (21) does not hold, the searcher waits until the next tip and then repeats the comparison between the two options using the new information obtained from the tip. The myopic condition simplifies in two special cases that depend upon the max term in (21):

$$p_1 \geq \begin{cases} \frac{\rho}{1+\rho} (1 - \alpha) + \frac{1}{1+\rho} q & \text{if } p_1 \leq \gamma p_i \quad \forall i \\ \\ 1 - \alpha & \text{if } p_1 \geq \gamma p_i \quad \forall i > 1. \end{cases} \tag{22}$$

The first case in (22) occurs when the max expression in (21) always returns the first term. This situation corresponds to a "roughly uniform" state p; whatever cell the informant points to with the next tip will become a best candidate cell. The first case in (22) is similar to the condition for the optimal threshold in the 2-cell case exceeding 0.5 (see equation (12)). The second case in (22) corresponds to the case when the max in (21) always returns the second term. This occurs when cell 1 is a "strong" best candidate cell; even if the informant points to cell $i \neq 1$ with the next tip, cell 1 remains a best candidate cell.

If $\rho >> 1$ (i.e., the threat is imminent and tips are scarce) or we have a highly reliable informant (q close to 1), the myopic conditions to engage in (21)–(22) closely resemble the sufficient condition to engage in (15). In this case, the myopic policy produces nearly optimal recommendations.

The first part of condition (22) holds for the uniform state $\tilde{p} = (\frac{1}{n}, ..., \frac{1}{n})$ and corresponds to condition (17). Following one tip (pointing at cell 1) the system transitions from $\tilde{p}$ to the new state p, where $p_1 = q$ and $p_i = \frac{1-q}{n-1}$ for $i > 1$. Therefore the second part of condition (22) holds for state p. Consequently if $(1-q) < \alpha$ and the search starts with a uniform prior, the searcher obtains at most one tip before engaging if he follows the myopic policy. Specifically, the searcher engages cell 1 before obtaining any tips if

$$\frac{1}{n} \geq \frac{\rho}{1+\rho} (1-\alpha) + \frac{1}{1+\rho} q.$$

Otherwise the searcher engages the cell provided in the first tip since $p_1 = q > 1 - \alpha$.

7 Analysis

Looking at some representative scenarios, we next analyze results from Section 6. Subsection 7.1 examines the three-cell case and and in Subsection 7.2 we analyze the effect of number of cells on the expected cost.

7.1 Three-cell Case

Figure 5 illustrates the three-cell engage region in the $p_1 \times p_2$ plane for $p_1 \geq p_2 \geq p_3 = 1 - p_1 - p_2$. The thin dashed-line triangle outlines the feasible p_1, p_2 values. Each subfigure fixes values for α and ρ and contains four curves for $q \in \{0.35, 0.55, 0.75, 0.95\}$. The southeast area of the cone corresponds to the engage region of the state space. As discussed in the introduction of Section 6, a threshold policy may not be optimal. However, in many cases such a policy may perform well based on the vertical nature of the boundaries when, for example, α is relatively small or q is not too small.

Similarly to the two-cell case, the engage region decreases with the reliability of the informant because the benefit from additional tips increases. Larger values of α or ρ increase the size of the engage region because the cost or likelihood of an attack increases. For

larger value of ρ (Figures 5b and 5d) the boundaries for the various reliability values are closer together than for smaller ρ (Figures 5a and 5c). The informational value of tips for smaller ρ is greater than for larger ρ, and therefore the reliability has a greater impact. The wait region in Figure 5d is empty because this situation corresponds to a blind engagement scenario (see condition (16)), which implies the searcher will engage for any state for any informant reliability. We only consider $\rho \leq 1$ scenarios; larger values of ρ (imminent attack compared to the flow of tips) correspond to blind engagement scenarios for most values of α.

(a) $\alpha = 0.5$ and $\rho = 0.1$

(b) $\alpha = 0.5$ and $\rho = 1$

(c) $\alpha = 1.5$ and $\rho = 0.1$

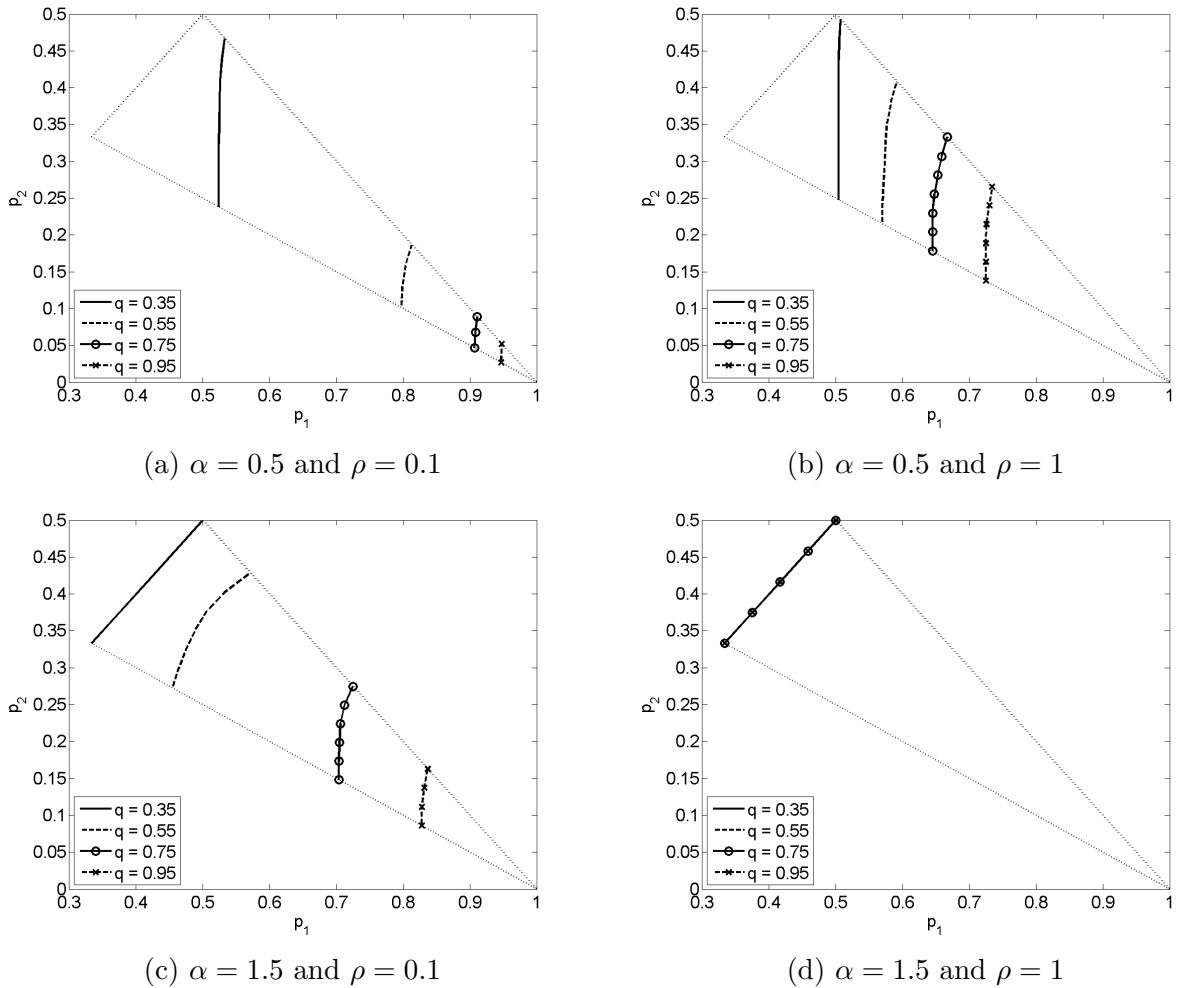(d) $\alpha = 1.5$ and $\rho = 1$

Figure 5: Engage region for $q \in \{0.35, 0.55, 0.75, 0.95\}$ and combinations of $\alpha \in \{0.5, 1.5\}$ and $\rho \in \{0.1, 1\}$. The engage region lies to the southeast of each curve.

In Section 6 we derive sufficient conditions to engage or wait that the searcher can use as heuristic policies. Figure 6, which has the same structure as Figure 5, illustrates the engage

regions generated by these heuristics. The smooth solid line represents the optimal *engage-wait* boundary. The other three (marked) solid lines correspond to heuristics based on the sufficient conditions to engage described in Section 6.1, as explained in the following:

- The sufficient condition to engage in (15), corresponding to perfect detection after one tip, is denoted *eng(1-tip)* and represented by the –∘– curve.

- Condition (19), corresponding to perfect detection after two tips, is denoted *eng(2-tips)* and represented by the the –×– curve. As discussed in Section 6.1, condition (19) is tighter than (15) and thus lies closer to the optimal curve.

- Condition (20), which we derive by combining cells 2 and 3 into an uber-cell and using the two-cell threshold policy, is denoted *eng(2-cell policy)* and corresponds to the –∇– curve.

Figure 6 also contains the myopic policy, which is associated with the *wait* conditions from Section 6.2. The condition appears in (21)–(22) and we denote it on the figure as *wait(myopic)* and it corresponds to the − − ∘ − − curve.

The *eng(1-tip)* heuristic (–∘–) performs poorly. This is not surprising considering it assumes 0 cost after one tip. The *eng(2-cell policy)* rule (–∇–) performs reasonably well overall. In situations with large α and ρ (Figure 6d), nearly all the heuristics produce optimal results.

The *wait(myopic)* heuristics performs very well except for small values of α and ρ (Figure 6a). In such "low-cost-of-attack, low-risk-of-attack" scenarios the searcher gains significant benefits from waiting for several additional tips, and *wait(myopic)* fails to account for this. "Murky" states with limited situational awareness lie at the northwest region of the state space, whereas "clear" states with a strong best candidate cell lie in the southeast. If *wait(myopic)* recommends to engage in a murky state, engaging usually is the optimal policy. However, this policy may produce the wrong decision in clear states for small values of ρ. For example consider the state $p = (0.70, 0.20, 0.10)$ in Figure 6a. Intuitively, engaging seems

like the right decision for this state as cell 1 is a strong candidate for the target's location. Indeed, *wait(myopic)* recommends to engage in this state. However, because ρ is small, the searcher can afford to collect several more tips to strengthen situational awarenesses and the optimal policy recognizes it: the optimal engage region lies significantly to the southeast of $p = (0.70, 0.20, 0.10)$ in Figure 6a.
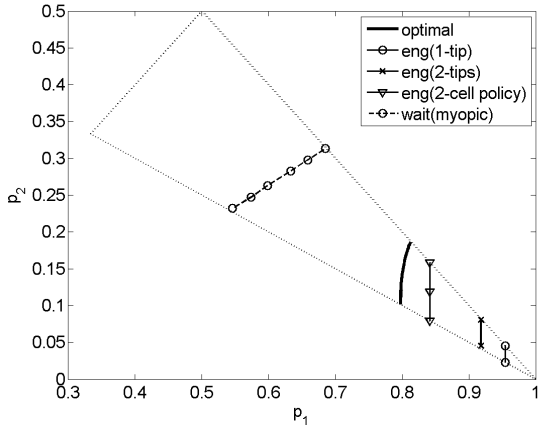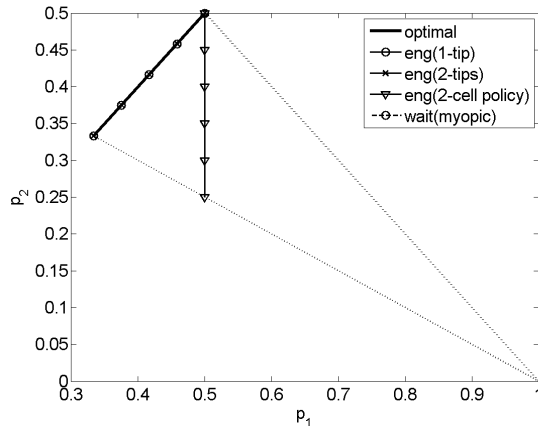

(a) $\alpha = 0.5$ and $\rho = 0.1$
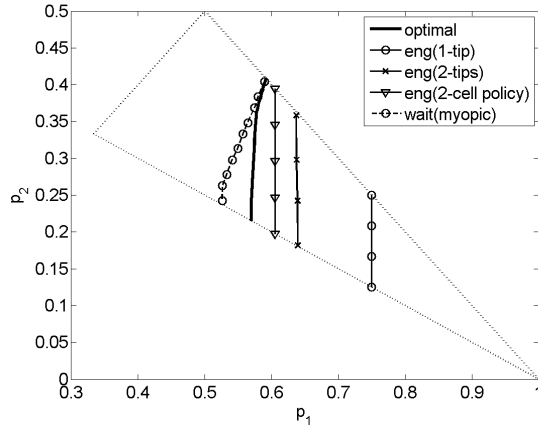
(b) $\alpha = 0.5$ and $\rho = 1$
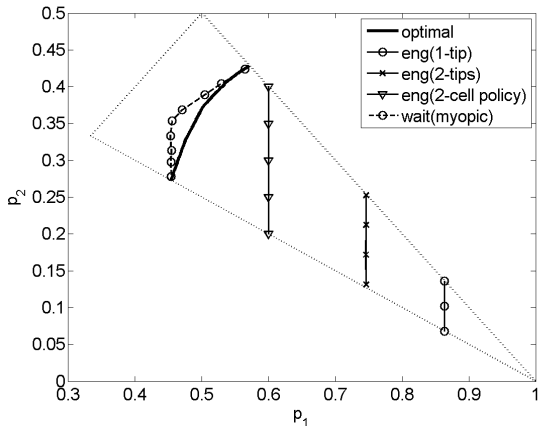
(c) $\alpha = 1.5$ and $\rho = 0.1$

(d) $\alpha = 1.5$ and $\rho = 1$

Figure 6: Engage region for various heuristic policies for $q = 0.55$ and combinations of $\alpha \in \{0.5, 1.5\}$ and $\rho \in \{0.1, 1\}$. The engage region lies to the southeast of each curve.

We also examine how much the cost increases using a heuristic instead of the optimal policy by generating 84000 scenarios representative of the examples in Figures 5 and 6 for $0.35 \leq q \leq 0.95$, $0.5 \leq \alpha \leq 1.5$, $0.1 \leq \rho \leq 1$, over the entire state space for p. The myopic policy performs very well; on average it is within 1% of optimal. Figure 6a illustrates when

the myopic policy can produce a cost significantly greater than optimal: small ρ and α and moderate q and p_1. There is no benefit to one additional tip, but reasonable cost reduction can occur through several additional tips. The strong performance of the myopic policy also occurs for $n > 3$ as long as ρ is not too small (i.e., $\rho > 0.1$). See Appendix J.1 for a more thorough analysis of several heuristics for both $n = 3$ and $n > 3$ scenarios. These results suggest that not only can the searcher confidently use the myopic policy operationally in most scenarios, but the policy may provide a rough estimate of the cost to wait, which is analytically difficult to compute. In practice, if the cost to wait is only slightly smaller than the cost to engage, the searcher may still choose to engage because of uncertainties associated with the model parameters or other frictions we do not account for in the model. In Appendix J.2 we explore this idea further.

7.2 Impact of Number of Cells

Following the discussion in Section 5, we observe that the situation seems to improve for the searcher as the number of cells n increases because it becomes less likely that incorrect tips will cluster on one particular cell, leading the searcher astray. Figure 7 displays the relationship between the optimal cost $C(p)$ and n for various values of q and two scenarios regarding an attack: (a) low-cost, low-risk (Figure 7a) and (b) high-cost, high-risk (Figure 7b). These figures illustrate that increasing n may generate only minor benefits, and the cost may actually increase in certain situations. The slope of the curve depends upon one of three possible policies taken by the searcher:

1. Blind engagement scenario: searcher engages a cell uniformly at random incurring cost of $(n - 1)/n$.

2. The searcher obtains one tip and engages the corresponding cell, which incurs cost $\frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1 - q)$.

3. The searcher obtains at least two tips.

For option 1 the searcher prefers a small n, the option 2 cost is independent of n, and intuitively the cost should decrease with n for option 3. In the high-cost, high-risk scenario in Figure 7b, the searcher chooses either option 1 (when the curves increase) or option 2 (when the curves flatten out). For small α and ρ (Figure 7a), the searcher chooses either option 2 or 3. Even though the cost is non-increasing with n in Figure 7a, the cost significantly decreases for only moderate values of q and the curves flatten out quickly.
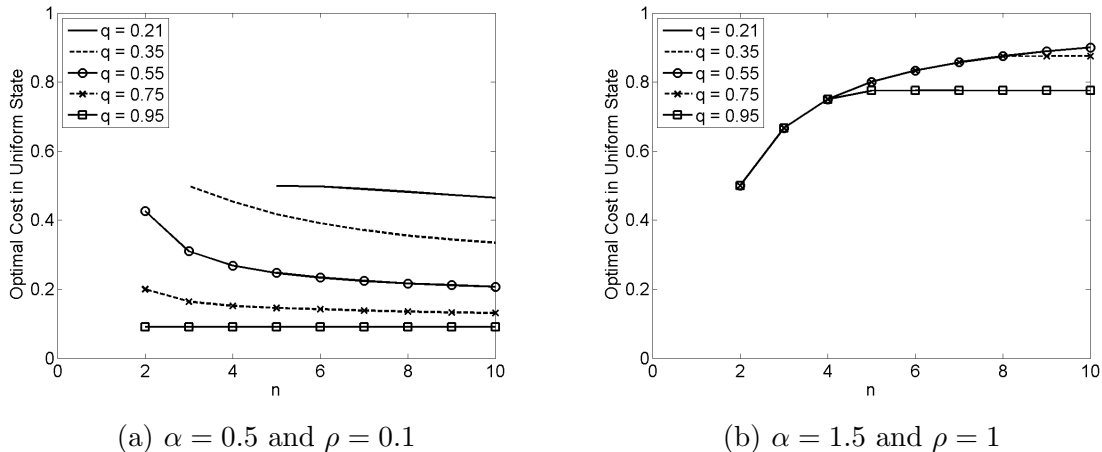

(a) $\alpha = 0.5$ and $\rho = 0.1$ (b) $\alpha = 1.5$ and $\rho = 1$

Figure 7: Optimal cost in the uniform state as a function of n for $q \in \{0.21, 0.35, 0.55, 0.75, 0.95\}$ for two combinations of α and ρ.

8   Extensions

In our model we make several assumptions that may not apply in reality. Our objective is to gain insight through analysis of a relatively simple setting. Several extensions are possible, and the key to handling them is to properly modify the cost function (11) such that most of the results from Section 3–6 generalize in a natural way. Due to space considerations, we only present one extension in this section. Appendix L considers several others. The main extension we analyze here focuses on the situation where the search continues if the searcher chooses the wrong cell. In this case, the target does not rush his attack if the searcher chooses the wrong cell and only executes a mature attack. In Appendix L we consider the situation where one source generates a stream of correlated tips. In that case future tips become less

valuable. We also examine the situation where there is no target and the searcher has the option to end the search before an engagement. Other extensions allow for multiple classes of informants and non-exponential distributions for the time until the target executes the attack.

8.1 Search Continues Following an Incorrect Engagement

In some situations, when the target is either oblivious of the searcher's failed attempt or determined to wait until the plot matures, the search may continue following the engagement of an empty cell. Because the target is static and detection is perfect, the searcher can discard evidently empty cells from future consideration. Specifically, $p_j = 0$ following an engagement of an empty cell j. The cost of engaging cell j incorrectly is c_j. Because we allow the false positive cost to vary by cell, the searcher may opt to engage cells with a small location probability if c_j is also small, in order to eliminate the cell from further consideration. Rather than use the cost-ratio α, in this subsection we include separate parameters for the the false positive cost (c_j) and the damage from a mature attack (d).

The system now has two types of state transitions. The first, as before, occurs when a tip points at cell i, in which case state p transitions to state $p^{(+i)}$. The second (new) type occurs when the searcher incorrectly engages cell j, and the state p transitions to state $p^{(-j)}$ in which $p_j = 0$. The set $A(p) = \{i : p_i > 0\}$ represents the "active" cells (i.e., cells that have not been incorrectly searched yet). The informant is aware of the searcher's failed engagements and therefore refrains from pointing at these cells in future tips. The probability mass associated with an evidently empty cell is proportionally redistributed among the active cells. That is,

$$\mathbf{P}[\text{informant says } i \mid p, \text{target in } k] = \begin{cases} \dfrac{q}{q + (|A(p)| - 1)\frac{1-q}{n-1}} & \text{if } i = k \\[3ex] \dfrac{\frac{1-q}{n-1}}{q + (|A(p)| - 1)\frac{1-q}{n-1}} & \text{if } i \neq k. \end{cases}$$

Under this reasonable assumption the ratio γ between the probabilities of correct and incorrect tips remains unchanged, and therefore $p^{(+i)}$ is computed as in equation (5). If cell i is searched and found empty then,

$$
p_j^{(-i)} =
\begin{cases}
0 & \text{if } j = i \\[2ex]
\dfrac{p_j}{\sum_{k \neq i} p_k} & \text{if } j \neq i.
\end{cases}
$$

Next, we slightly modify the the definition of $r_i(p)$ from (7) to ensure that $\sum_{i=1}^{n} r_i(p) = 1$. Specifically,

$$
r_i(p) =
\begin{cases}
\dfrac{q}{q + (|A(p)| - 1)\frac{1-q}{n-1}} p_i + \dfrac{\frac{1-q}{n-1}}{q + (|A(p)| - 1)\frac{1-q}{n-1}} (1 - p_i) & \text{if } i \in A(p) \\[3ex]
0 & \text{if } i \notin A(p).
\end{cases}
$$

While the expected cost to *wait* remains essentially the same as in the original model, the expected cost to engage becomes:

$$
\mathbf{E}[\text{Cost of engaging cell } j \mid p] = (1 - p_j)(c_j + C(p^{(-j)})).
$$

The updated cost function is:

$$
C(p) = \min \left(\min_{j \in A(p)} \left((1 - p_j)(c_j + C(p^{(-j)})) \right), \frac{\rho}{1 + \rho} d + \frac{1}{1 + \rho} \sum_{i \in A(p)} r_i(p) C(p^{(+i)}) \right). \tag{23}
$$

Obviously, if only one active cell remains ($|A(p)| = 1$), $C(p) = 0$ because the searcher knows the only remaining cell contains the target.

The analysis of the cost function and engage decision is similar to the analysis in Sections 3–7. First consider the case of imminent threat where the searcher does not wait for tips but continuously engages cells until he finds the target. This is the classical *whereabouts search* problem (Kadane, 1971; Stone, 1975) for which the optimal policy is to search the cells in

ascending order of the ratios $\frac{c_j}{p_j}, j = 1, ..., n$. Let $g(i)$ denote the index of the ith smallest value of $\frac{c_j}{p_j}$ in $A(p)$. So, $g(1)$ and $g(|A(p)|)$ are the indices of the cells with the smallest and largest ratios $\frac{c_j}{p_j}$, respectively. Let $K(p)$ denote the cost of this policy. In the Appendix K we show that

$$K(p) = \sum_{j=2}^{|A(p)|} p_{g(j)} \sum_{i=1}^{j-1} c_{g(i)}. \tag{24}$$

The searcher should engage a cell if $K(p) \leq \frac{\rho}{1+\rho}d$. If that engaged cell is empty, this condition may not hold in the next state. It is most reasonable (albeit, not proved) that the searcher should engage cell $g(1)$.

$K(p)$ also plays a crucial role in the sufficient condition to wait

$$\text{wait if} \quad \min_{j \in A(p)} c_j(1 - p_j) > \frac{\rho}{1+\rho}d + \frac{1}{1+\rho} \sum_{i \in A(p)} r_i(p)K(p^{(+i)}).$$

Note that computing $K(p^{(+i)})$ requires ranking according to $\frac{c_j}{p_j^{(+i)}}$, which depends on i.

9 Summary and Conclusions

In this paper we study a time-critical variant of the whereabout problem in search theory. This variant applies to many criminal, military, and homeland security situations where an investigation team must decide when to act on uncertain intelligence. Examples include counter-terror and counterinsurgency operations, which rely on human intelligence and intercepted communications. Unlike the original whereabout model that produces a *sequencing* rule, we consider here a *stopping* rule; rather than advising the searcher how to optimally sequence the search among the various cells, our model identifies the time when the information is sufficiently definitive to act upon. Either actions - engage or wait for additional information – incur costs. We analytically solve the two extremes: the two-cell case uses a threshold policy and the searcher chooses among three options in the infinite-cell case. We

also illustrate how the engage region of the state space varies with the model parameters for the three-cell case. For larger problems, we use a k–stage look-ahead approach to obtain sufficient conditions to engage or wait. We show that these conditions converge to a necessary and sufficient condition to engage as k increases. In particular for $k = 1$, the myopic policy provides nearly optimal results over a broad range of parameter values. The model clearly captures the tradeoffs among the various components of the threat: the mean time until the plot matures, the flow rate of tips, and the damages associated with failed searches and successful attacks. We present several variants of the model in Section 8 and Appendix L to capture alternative scenarios. These include the search continuing after an incorrect engagement, multiple types of informants, and non-exponential attack time distributions. Most of the analysis and methods discussed apply to these extensions.

Some of our main results are intuitive: the searcher is more likely to wait with a more reliable informant and is more likely to engage as the cost or likelihood of a mature attack increases. Less intuitive insights that emerge from our analysis include: (1) The optimal number of tips received by the searcher may not be monotone as a function of the informant reliability (see Section 4) and (2) In many cases there is little to no reduction in the optimal cost as we increase the number of cells (see Section 7.2).

Future work could model the reliability parameter q as a random variable (e.g., beta distributed), which updates as the searcher receives more information. This would be particularly appropriate in the situation where the target only executes his attack when it fully matures (see Section 8.1). In this case the searcher could search multiple cells and and thus verify the reliability of the informant. Another variant would capture strategic behavior of the target who trades off a more effective attack that needs longer planning time with the increased risk of detection by the searcher. Finally, one could examine another time-critical situation where the target may leave instead of executing an attack (e.g., a criminal or terrorist leader who moves around to avoid detection). In this case the searcher has three options: receive another tip, engage a cell, or call off the search because the target has likely left the

system. The modeling of this situation may include changepoint analysis (Carlstein et al., 1994) to handle the change in tip dynamics after the target departs.

Acknowledgment

The authors would like to thank their colleagues in the Operations Research department at the Naval Postgraduate School for valuable feedback. In particular, Susan Sanchez provided several useful references and Kyle Lin aided in the formulation of proof concepts. They would also like the thank the referees and associate editor for invaluable comments that significantly improved the paper.

References

An, Bo, Matthew Brown, Yevgeniy Vorobeychik, Milind Tambe. 2013. Security games with surveillance cost and optimal timing of attack execution. *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 223–230.

Bechhofer, Robert E., Salah Elmaghraby, Norman Morse. 1959. A single-sample multiple-decision procedure for selecting the multinomial event which has the highest probability. *The Annals of Mathematical Statistics* **30**(1) 102–119.

Berger, James. 1985. *Statistical Decision Theory and Bayesian Analysis*. Springer.

Boccio, John. 2012. *Gambler's Ruin Problem*. Physics Department, Swarthmore. `http://www.johnboccio.com/research/quantum/notes/ruin.pdf`.

Carlstein, Edward, Hans-Georg Müller, David Siegmund. 1994. *Change-point Problems*. Institute of Mathematical Statistics.

Chen, Pinyuen. 1988. An integrated formulation for selecting the most probable multinomial cell. *Annals of the Institute of Statistical Mathematics* **40**(3) 615–625.

Chow, Yuan Shih, Herbert Robbins, David Siegmund. 1971. *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin.

Dayanik, Savas, Angela J. Yu. 2013. Reward-rate maximization in sequential identification under a stochastic deadline. *SIAM Journal on Control and Optimization* **51**(4) 2922–2948.

Dennis, Martin J., Woo-Kyoung Ahn. 2001. Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition* **29**(1) 152–164.

Ferguson, Thomas S. 2004. *Optimal Stopping and Applications*. Mathematics Department, UCLA. `http://www.math.ucla.edu/~tom/Stopping/Contents.html`.

Frazier, Peter, Angela J. Yu. 2007. Sequential hypothesis testing under stochastic deadlines. J.C. Platt, D. Koller, Y. Singer, S.T. Roweis, eds., *Advances in Neural Information Processing Systems 20*.

Kadane, Joseph B. 1971. Optimal whereabouts search. *Operations Research* **19**(4) 894–904.

Kaplan, Edward H. 2010. Terror queues. *Operations research* **58** 773–784.

Kaplan, Edward H. 2012. Estimating the duration of jihadi terror plots in the United States. *Studies in Conflict & Terrorism* **35**(12) 880–894.

Kim, Seong-Hee, Barry L Nelson. 2006. Selecting the best system. Shane Henderson, Barry Nelson, eds., *Handbooks in Operations Research and Management Science: Simulation*. Elsevier, 501–534.

Komiya, Toru, Koji Iida, Ryusuke Hohzaki. 2006. An optimal investigation in two stage search with recognition errors. *Journal of the Operations Research Society of Japan* **49**(2) 130–143.

Kress, Moshe, Kyle Y. Lin, Roberto Szechtman. 2008. Optimal discrete search with imperfect specificity. *Mathematical Methods of Operations Research* **68**(3) 539–549.

Lange, Rutger-Jan. 2012. Brownian motion and multidimensional decision making. Ph.D. thesis, University of Cambridge.

Ramey, James T., Khursheed Alam. 1979. A sequential procedure for selecting the most probable multinomial event. *Biometrika* **66**(1) 171–173.

Shiryaev, Albert. 2007. *Optimal Stopping Rules*. Stochastic Modelling and Applied Probability, Springer.

Stone, Lawrence D. 1975. *Theory of Optimal Search*. Academic Press New York.

Szechtman, Roberto, Moshe Kress, Kyle Lin, Dolev Cfir. 2008. Models of sensor operations for border surveillance. *Naval Research Logistics (NRL)* **55**(1) 27–41.

Vieira, Helcio, Susan M. Sanchez, Paul J. Sanchez, Karl Kienitz, Mischel Belderrain. 2014. A restricted multinomial hybrid selection procedure. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* **24**(2).

Wald, Abraham, Jacob Wolfowitz. 1948. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics* **19**(3) 326–339.

Washburn, Alan. 2014. *Two-Person Zero-Sum Games*. 4th ed. Springer, New York.

Wilson, Kurt E., Roberto Szechtman, Michael P. Atkinson. 2011. A sequential perspective on searching for static targets. *European Journal of Operational Research* **215**(1) 218–226.

When is Information Sufficient for Action? Search with Unreliable Yet Informative Intelligence

Michael P. Atkinson, Moshe Kress, Rutger-Jan Lange

This electronic companion contains the proofs for the propositions in the main text and other supplementary information and technical details that could not be included in the main text due to space constraints.

A Proposition 1: Bound on $\max_j p_j$ with Unique Best Cell

Using equation (4) for the uniform prior case ($\tilde{p}_i = \frac{1}{n}$ for all i) produces

$$p_i = \frac{\gamma^{s_i}}{\sum_{j=1}^n \gamma^{s_j}}.$$

In the uniform prior case, the cell with the highest probability of containing the target corresponds to the cell with the largest number of tips. To bound $\max_j p_j$, we label the cells in descending order by tips received: $s_1 \geq s_2 \geq \ldots \geq s_n$. With this notation $p_1 = \max_j p_j$. We bound $\max_j p_j$ based on s_1 and s_2:

$$\max_j p_j = p_1 = \frac{\gamma^{s_1}}{\sum_{j=1}^n \gamma^{s_j}} \geq \frac{\gamma^{s_1}}{\gamma^{s_1} + (n-1)\gamma^{s_2}} = \frac{\gamma^{s_1-s_2}}{\gamma^{s_1-s_2} + (n-1)} = q\frac{\gamma^{s_1-s_2-1}}{q\gamma^{s_1-s_2-1} + (1-q)}. \quad \text{(EC.1)}$$

If $|B(p)| = 1$, then $s_1 - s_2 - 1 \geq 0$ and hence condition (EC.1) implies that $\max_j p_j \geq q$. ∎

B Characteristics of Cost Function $C(p)$

Proposition EC.1. $C(p)$ *is a concave function over the domain* $\{p \,:\, p \in [0,1]^n, \sum_{i=1}^{n} p_i = 1\}$ *and has a global maximum at* $p^* = (\frac{1}{n}, \frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n})$.

The proof appears in Appendix B.1. Using the concavity of $C(p)$ we next characterize the engage region.

Proposition EC.2. *The engage region is a convex subset of the polytope defined by* $p_1 \geq p_2... \geq p_n$.

The proof for this proposition as well as the next corollary appears in Appendix B.2.

Corollary EC.1. *If the searcher should engage in p, then the searcher should also engage in* $p^{(+1)}$.

The proofs of the next proposition and corollary are given in Appendix B.3.

Proposition EC.3. *A confirming tip – a tip pointing at cell 1 – produces a cost lower than any other tip:* $C(p^{(+1)}) \leq C(p^{(+k)})$ *for all k.*

Corollary EC.2. *A confirming tip reduces the cost function:* $C(p^{(+1)}) \leq C(p)$.

Since, intuitively, a non-confirming tip makes the situation more murky, one would expect that an analogous result to Corollary EC.2 exists, which states that a non-confirming tip leads to an increase in the expected cost. The following example illustrates that this is not necessarily the case.

Example 1: Let $q = 0.4$, $\alpha = 0.6$, $\rho = 1$. For the state $p = [0.182, 0.136, 0.136, 0.136, 0.136, 0.136, 0.136]$, we have $C(p) = 0.539$. However, after receiving an inconsistent tip that cell 2 contains the target, the state transitions to $p^{(+2)} = [0.174, 0.174, 0.130, 0.130, 0.130, 0.130, 0.130]$ with decreased cost $C(p^{(+2)}) = 0.537$.

The situational awareness is poor for both states in this example. The probability that cell 1 or 2 contains the target is higher in $p^{(+2)}$ (0.348) than in p (0.318), which leads to the slightly lower expected costs.

B.1 Proposition EC.1: Concavity of $C(p)$

Concavity is a common property for optimal stopping cost functions. See Lemma 1 of Frazier and Yu (2007), page 168 of Shiryaev (2007), page 4.11 of Ferguson (2004), or page 49 of Chow et al. (1971). We follow the approach of Ferguson (2004). Assume the searcher knows he faces either state p^A or state p^B. Before the searcher chooses to engage or wait, a coin flip occurs. It lands heads with probability β, and in this case the searcher faces state p^A. Otherwise the searcher faces state p^B if the coin lands tails. If the searcher observes the result of the coin toss, then he will act optimally for the specified state and incurs the following expected cost

$$\beta C(p^A) + (1 - \beta)C(p^B).$$

If the searcher cannot observe the result of the coin flip, then the searcher acts optimally for state $p = \beta p^A + (1 - \beta)p^B$ and achieves an expected cost of

$$C(p) = C\left(\beta p^A + (1 - \beta)p^B\right).$$

If the searcher observes the coin flip, he can always ignore that information and follow the optimal strategy without that knowledge. Thus the expected cost achieved with information must be less than the expected cost without information:

$$\beta C(p^A) + (1 - \beta)C(p^B) \leq C\left(\beta p^A + (1 - \beta)p^B\right),$$

which is the condition for concavity.

Intuitively the cost function $C(p)$ has a global maximum at $p^* = (\frac{1}{n}, \frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n})$ by symmetry. Formally, assume that $C(p)$ achieves its maximum value at $\hat{p} \neq p^*$ and $C(\hat{p}) > C(p^*)$. By symmetry there must be at least $n!$ maximizers because we can permute the elements of $\hat{p}$ without changing the cost. If we label these maximizers $\hat{p}^1, \hat{p}^2, \ldots \hat{p}^{n!}$, then $p^* = \frac{1}{n!} \sum_{i=1}^{n!} \hat{p}^i$. This follows because each element of $\hat{p}$ appears in the jth index of exactly

$(n-1)!$ of the $\hat{p}^k$ vectors. By concavity we have a contradiction: $C(p^*) \geq \frac{1}{n!} \sum_{i=1}^{n!} C(\hat{p}^i) = C(\hat{p}) = \max_p C(p)$. Hence p^* must be a global maximizer of $C(p)$. $\blacksquare$

B.2   Proposition EC.2 and Corollary EC.1 : Engage Region

Without loss of generality we can always restrict out state space to the polytope $A = \{p :$ $p \in [0,1]^n, \sum_{i=1}^n p_i = 1, \ p_1 \geq p_2 \geq p_3 \geq \ldots \geq p_n\}$. Define $p^X \in A$ and $p^Y \in A$ as two states where engage is the optimal policy. Next define $\beta \in [0,1]$ and $p = \beta p^X + (1-\beta)p^Y$. Note that $p \in A$ because A is a convex set. By concavity of $C(p)$ we have

$$C(p) \geq \beta C(p^X) + (1-\beta)C(p^Y) \tag{EC.2}$$

$$\geq \beta(1 - p_1^X) + (1-\beta)(1 - p_1^Y) \tag{EC.3}$$

$$\geq (1 - (\beta p_1^X + (1-\beta)p_1^Y)) \tag{EC.4}$$

$$\geq 1 - p_1. \tag{EC.5}$$

Condition (EC.3) follows because the searcher engages in states p^X and p^Y. Condition (EC.5) follows from the definition of p. We also have by definition of the cost function in (11)

$$C(p) \leq 1 - p_1. \tag{EC.6}$$

Combining conditions (EC.5) and (EC.6) produces $C(p) = (1 - p_1)$, and the searcher should engage in state p. Thus the engage region is convex.

The proof for Corollary EC.1 follows from the convexity of the engage region. Assume the searcher should engage in state p. Consider the standard unit vector $e_1 = [1, 0, 0 \ldots, 0]$. Clearly the searcher should engage in state e_1 because he achieves the minimum possible cost of 0 when taking this action. If we define $\beta = \frac{1}{\gamma p_1 + (1-p_1)}$ then

$$p^{(+1)} = \beta p + (1-\beta)e_1.$$

Because both p and e_1 lie in the engage region, by Proposition EC.2 so too must $p^{(+1)}$. Thus the searcher should engage in $p^{(+1)}$. ∎

B.3 Proposition EC.3 and Corollary EC.2: Cost Decreases After Confirming Tip

We prove $C(p^{(+1)}) \leq C(p^{(+k)})$ by utilizing the concavity of $C(\cdot)$. By symmetry we can permute the elements of $p^{(+1)}$ to construct $n!$ state vectors that all produce the same cost. Label these $n!$ vectors $\hat{p}^1$, $\hat{p}^2$, ..., $\hat{p}^{n!}$. By construction we have $\hat{p}^j = p^{(+1)}$ for some j and $C(\hat{p}^i) = C(\hat{p}^m) = C(p^{(+1)})$ for all $1 \leq i, m \leq n!$. In Lemma EC.1 we show that we can write $p^{(+k)}$ as a linear combination of the $\hat{p}^i$:

$$p^{(+k)} = \sum_{i=1}^{n!} \omega_i \hat{p}^i$$

$$\text{s.t.} \sum_{i=1}^{n!} \omega_i = 1, \ \omega_i \geq 0.$$

However by concavity of $C(\cdot)$ this implies

$$C(p^{(+k)}) \geq \sum_{i=1}^{n!} \omega_i C(\hat{p}^i) = C(p^{(+1)}),$$

which is the desired result.

To prove Corollary EC.2, we first assume the searcher should engage in state p. By Corollary EC.1 the searcher should also engage in state $p^{(+1)}$. This leads to the following relationship

$$C(p) = 1 - p_1 \geq 1 - \frac{\gamma}{\gamma p_1 + (1 - p_1)} p_1 = 1 - p_1^{(+1)} = C(p^{(+1)}). \tag{EC.7}$$

Thus $C(p^{(+1)}) \leq C(p)$ for states p in the engage region. Now we assume the searcher should

wait in state p:

$$C(p) = \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)C(p^{(+i)}) \tag{EC.8}$$

$$\geq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\min_i C(p^{(+i)}) \tag{EC.9}$$

$$= \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}C(p^{(+1)}) \tag{EC.10}$$

$$\geq \frac{\rho}{1+\rho}C(p^{(+1)}) + \frac{1}{1+\rho}C(p^{(+1)}) \tag{EC.11}$$

$$= C(p^{(+1)}). \tag{EC.12}$$

Condition (EC.10) follows from Proposition EC.3. Condition (EC.11) follows because $C(p) \leq \alpha$ for all valid states p because the searcher can achieve a cost of α by using the policy "always wait." Combining conditions (EC.7)–(EC.12), we have shown $C(p^{(+1)}) \leq C(p)$ for all p. ■

Lemma EC.1. *For $\hat{p}^i$ as defined above in the proof for Proposition EC.3, we can write $p^{(+k)}$ as a linear combination of the $\hat{p}^i$*

$$p^{(+k)} = \sum_{i=1}^{n!} \omega_i \hat{p}^i \tag{EC.13}$$

$$s.t. \sum_{i=1}^{n!} \omega_i = 1, \ \omega_i \geq 0. \tag{EC.14}$$

Proof. To prove this we rely on properties of linear programming (LP). We construct an LP with constraints given by (EC.13)–(EC.14). To prove the lemma we only need to show that the problem is feasible, and thus we can define an arbitrary objective function. Define $e = [1,1,1,\ldots,1]$ to be a column vector of all ones and $\hat{P}$ the matrix with column i set to

$\hat{p}^i$. We define the following primal linear program in standard form

$$\text{maximize} \quad e^T \omega \tag{EC.15}$$

$$\text{s.t.} \quad \begin{pmatrix} \hat{P} \\ -\hat{P} \\ e^T \\ -e^T \end{pmatrix} \omega \leq \begin{pmatrix} p^{(+k)} \\ -p^{(+k)} \\ 1 \\ -1 \end{pmatrix} \tag{EC.16}$$

$$\omega \geq 0. \tag{EC.17}$$

By showing this problem is feasible, we will complete the proof. If this linear program is feasible, then by construction its optimal value is 1. We can prove the primal defined by (EC.15)–(EC.17) is feasible by showing that its dual is bounded. We next present the corresponding dual. We separate the dual variable y into four components corresponding to the four constraints in (EC.16).

$$\text{minimize} \quad \begin{pmatrix} (p^{(+k)})^T & -(p^{(+k)})^T & 1 & -1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} \qquad \text{(EC.18)}$$

$$\text{s.t.} \begin{pmatrix} \hat{P}^T & -\hat{P}^T & e & -e \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} \geq e \qquad \text{(EC.19)}$$

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} \geq 0. \qquad \text{(EC.20)}$$

The dual has a feasible point at $y_1 = y_2 = [0, 0, 0, \ldots, 0]$, $y_3 = 1$, $y_4 = 0$ that produces an objective value of 1. We now show that 1 is the optimal value of the dual and hence the primal is feasible. First assume that some feasible solution y^* produces an objective value for the dual less than 1:

$$(p^{(+k)})^T(y_1^* - y_2^*) + y_3^* - y_4^* < 1. \qquad \text{(EC.21)}$$

We now show that if y^* satisfies (EC.21) then it will violate one of the constraints in (EC.19), and hence y^* is infeasible and we have a contradiction. For notational simplicity we define $\delta \equiv y_1^* - y_2^*$. By inspection of (EC.19) and (EC.21), to show infeasibility, it suffices to show that there exists some j such that

$$(\hat{p}^j)^T \delta \leq (p^{(+k)})^T \delta. \qquad \text{(EC.22)}$$

We will use a stochastic dominance argument to show (EC.22). We next define three random variables: X_1, X_2, X_3. Each random variable is a discrete random variable that takes on the same n values: $\delta_1, \delta_2, \delta_3, \ldots, \delta_n$. The random variables only differ by the probability assigned to each of those n values. The probability mass functions for X_1, X_2, and X_3 appear in Table 1. For X_1 we have $\mathbf{P}[X_1 = \delta_i] = p_i^{(+k)}$ and hence $\mathbf{E}[X_1] = (p^{(+k)})^T \delta$.

To define X_2 and X_3 we first sort the elements of δ in ascending order and denote this vector δ^S: δ_i corresponds to the ith element of δ, whereas δ_i^S is the ith smallest element of δ. Hence, $\min \delta = \delta_1^S \leq \delta_2^S \leq \ldots \leq \delta_n^S = \max \delta$. By defining an appropriate tie-breaking rule for duplicate values in δ (e.g., by original ordering), there is a one-to-one correspondence between the elements of δ and the elements of δ^S. With this correspondence, we can define the invertible function $g(\cdot)$ that maps elements of δ to elements of δ^S. By construction, this leads to the relationship $\delta_{g(i)}^S = \delta_i$. The g function effectively ranks the values in δ in ascending order. Before defining X_2 we must also define a sorted version of $p^{(+k)}$. The vectors p and $p^{(+1)}$ are sorted in decreasing order. However, $p^{(+k)}$ may not be sorted because in transitioning from p to $p^{(+k)}$ the kth value increases while the other values decrease. The kth element of $p^{(+k)}$ is the only value that is possibly out of order, so we only need to move that value to the correct index and shift the remaining values to produce a sorted version of $p^{(+k)}$. Assume that $p_k^{(+k)}$ is now the mth largest value of $p^{(+k)}$ for some $m \leq k$: $p_m^{(+k)} \leq p_k^{(+k)} < p_{m-1}^{(+k)}$. We define $\tilde{p}^{(+k)}$ as the sorted version of $p^{(+k)}$ with $p_k^{(+k)}$ moved to index m. The following lists p, $p^{(+1)}$, $p^{(+k)}$, and $\tilde{p}^{(+k)}$

$$p = \qquad (\ p_1, p_2, p_3, \ldots, p_{m-1}, p_m, p_{m+1}, \ldots, p_{k-1}, \ p_k, p_{k+1}, \ldots, p_n) \quad \text{(EC.23)}$$

$$p^{(+1)} = \frac{1}{\gamma p_1 + (1 - p_1)}(\gamma p_1, p_2, p_3, \ldots, p_{m-1}, p_m, p_{m+1}, \ldots, p_{k-1}, \ p_k, p_{k+1}, \ldots, p_n) \quad \text{(EC.24)}$$

$$p^{(+k)} = \frac{1}{\gamma p_k + (1 - p_k)}(\ p_1, p_2, p_3, \ldots, p_{m-1}, p_m, p_{m+1}, \ldots, p_{k-1}, \gamma p_k, p_{k+1}, \ldots, p_n) \quad \text{(EC.25)}$$

$$\tilde{p}^{(+k)} = \frac{1}{\gamma p_k + (1 - p_k)}(\ p_1, p_2, p_3, \ldots, p_{m-1}, \gamma p_k, p_m \ , \ldots, p_{k-2}, p_{k-1}, p_{k+1}, \ldots, p_n). \quad \text{(EC.26)}$$

We define X_2 by using $\tilde{p}^{(+k)}$: $\mathbf{P}[X_2 = \delta_i] = \tilde{p}_{g(i)}^{(+k)}$. This assigns the largest probabilities in

$p^{(+k)}$ to the smallest values of δ. Thus $\mathbf{E}[X_2] \leq \mathbf{E}[X_1]$.

Finally we define X_3 with the probability vector $p^{(+1)}$: $\mathbf{P}[X_3 = \delta_i] = p^{(+1)}_{g(i)}$. Table 2 presents the exceedance probabilities $\mathbf{P}[X_2 \geq \delta_i^S]$ and $\mathbf{P}[X_3 \geq \delta_i^S]$. Table 1 lists the values in terms of δ_i, whereas Table 2 lists the values in terms of δ_i^S so that the exceedance probabilities will be non-increasing as we move down the table. By the construction of the ranking function $g(\cdot)$, $\mathbf{P}[X_3 = \delta_i^S] = p^{(+1)}_i$ and therefore $\mathbf{P}[X_3 \geq \delta_i^S] = 1 - \sum_{j=1}^{i-1} p^{(+1)}_j$. We write out the values in Table 2 in terms of p using the expressions in (EC.24) and (EC.26). A comparison of the columns in Table 2 reveals that X_2 stochastically dominates X_3. To see this note that the denominators (numerators) in the X_2 column are always less (greater) than the corresponding denominators (numerators) in the X_3. This stochastic dominance implies $\mathbf{E}[X_3] \leq \mathbf{E}[X_2]$ and by transitivity $\mathbf{E}[X_3] \leq \mathbf{E}[X_1]$. Furthermore there exists some j such that $\hat{p}_i^j = p^{(+1)}_{g(i)}$ for all i, so that the probability mass function of X_3 corresponds to the jth column of $\hat{P}$. This allows us to write $\mathbf{E}[X_3] = (\hat{p}^j)^T \delta$ and we have derived condition (EC.22)

$$(\hat{p}^j)^T \delta = \mathbf{E}[X_3] \leq \mathbf{E}[X_1] = (p^{(+k)})^T \delta.$$

Therefore any solution y^* that produces an objective value less than 1 must be infeasible. Specifically we have shown the jth constraint of (EC.19) is infeasible for index j defined above. Since the dual has an optimal solution of 1, the primal defined in (EC.15)–(EC.17) is feasible and the proof is complete.

$\square$

Value	$\mathbf{P}[X_1 = \delta_i]$	$\mathbf{P}[X_2 = \delta_i]$	$\mathbf{P}[X_3 = \delta_i]$
δ_1	$p_1^{(+k)}$	$\tilde{p}_{g(1)}^{(+k)}$	$p_{g(1)}^{(+1)}$
δ_2	$p_2^{(+k)}$	$\tilde{p}_{g(2)}^{(+k)}$	$p_{g(2)}^{(+1)}$
δ_3	$p_3^{(+k)}$	$\tilde{p}_{g(3)}^{(+k)}$	$p_{g(3)}^{(+1)}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
δ_n	$p_n^{(+k)}$	$\tilde{p}_{g(n)}^{(+k)}$	$p_{g(n)}^{(+1)}$

Table 1: Probability mass function for X_1, X_2, and X_3

Value	$\mathbf{P}[X_2 \geq \delta_i^S]$	$\mathbf{P}[X_3 \geq \delta_i^S]$
δ_1^S	1	1
δ_2^S	$1 - \tilde{p}_1^{(+k)} = \frac{(\gamma-1)p_k+(1-p_1)}{\gamma p_k+(1-p_k)}$	$1 - p_1^{(+1)} = \frac{(1-p_1)}{\gamma p_1+(1-p_1)}$
δ_3^S	$1 - \tilde{p}_1^{(+k)} - \tilde{p}_2^{(+k)} = \frac{(\gamma-1)p_k+(1-p_1-p_2)}{\gamma p_k+(1-p_k)}$	$1 - p_1^{(+1)} - p_2^{(+1)} = \frac{(1-p_1-p_2)}{\gamma p_1+(1-p_1)}$
δ_4^S	$1 - \sum_{j=1}^{3} \tilde{p}_i^{(+k)} = \frac{(\gamma-1)p_k+(1-\sum_{j=1}^{3} p_j)}{\gamma p_k+(1-p_k)}$	$1 - \sum_{j=1}^{3} p_i^{(+1)} = \frac{(1-\sum_{j=1}^{3} p_j)}{\gamma p_1+(1-p_1)}$
$\vdots$	$\vdots$	$\vdots$
δ_m^S	$1 - \sum_{j=1}^{m-1} \tilde{p}_i^{(+k)} = \frac{(\gamma-1)p_k+(1-\sum_{j=1}^{m-1} p_j)}{\gamma p_k+(1-p_k)}$	$1 - \sum_{j=1}^{m-1} p_i^{(+1)} = \frac{(1-\sum_{j=1}^{m-1} p_j)}{\gamma p_1+(1-p_1)}$
δ_{m+1}^S	$1 - \sum_{j=1}^{m} \tilde{p}_i^{(+k)} = \frac{(1-\sum_{j=1}^{m} p_j)+p_m-p_k}{\gamma p_k+(1-p_k)}$	$1 - \sum_{j=1}^{m} p_i^{(+1)} = \frac{(1-\sum_{j=1}^{m} p_j)}{\gamma p_1+(1-p_1)}$
δ_{m+2}^S	$1 - \sum_{j=1}^{m+1} \tilde{p}_i^{(+k)} = \frac{(1-\sum_{j=1}^{m+1} p_j)+p_{m+1}-p_k}{\gamma p_k+(1-p_k)}$	$1 - \sum_{j=1}^{m+1} p_i^{(+1)} = \frac{(1-\sum_{j=1}^{m+1} p_j)}{\gamma p_1+(1-p_1)}$
$\vdots$	$\vdots$	$\vdots$
δ_k^S	$1 - \sum_{j=1}^{k-1} \tilde{p}_i^{(+k)} = \frac{(1-\sum_{j=1}^{k-1} p_j)+p_{k-1}-p_k}{\gamma p_k+(1-p_k)}$	$1 - \sum_{j=1}^{k-1} p_i^{(+1)} = \frac{(1-\sum_{j=1}^{k-1} p_j)}{\gamma p_1+(1-p_1)}$
δ_{k+1}^S	$1 - \sum_{j=1}^{k} \tilde{p}_i^{(+k)} = \frac{(1-\sum_{j=1}^{k} p_j)}{\gamma p_k+(1-p_k)}$	$1 - \sum_{j=1}^{k} p_i^{(+1)} = \frac{(1-\sum_{j=1}^{k} p_j)}{\gamma p_1+(1-p_1)}$
δ_{k+2}^S	$1 - \sum_{j=1}^{k+1} \tilde{p}_i^{(+k)} = \frac{(1-\sum_{j=1}^{k+1} p_j)}{\gamma p_k+(1-p_k)}$	$1 - \sum_{j=1}^{k+1} p_i^{(+1)} = \frac{(1-\sum_{j=1}^{k+1} p_j)}{\gamma p_1+(1-p_1)}$
$\vdots$	$\vdots$	$\vdots$
δ_n^S	$\tilde{p}_n^{(+k)} = \frac{p_n}{\gamma p_k+(1-p_k)}$	$p_n^{(+1)} = \frac{p_n}{\gamma p_1+(1-p_1)}$

Table 2: Exceedance probabilities for X_2 and X_3

C   Corollary 1: Threshold Policy for $n = 2$ Cells

The optimality of a threshold policy follows immediately from Proposition EC.2 in Appendix B, which states the engage region is convex. The searcher should engage if $p_1 = 1$ because he achieves the minimum cost of 0 and the cost to wait is at least $\frac{\rho}{1+\rho}\alpha$. A convex subset of the interval $[\frac{1}{2}, 1]$ that contains the point 1 must be a subinterval of the form $[\tau, 1]$. This completes the proof for Corollary 1.

The more interesting aspect is deriving an expression of τ as a function of the model parameters q, α and ρ. As the expression, and accompanying derivation, of τ is quite complicated we present it in the next section.

D Derivation of Optimal Two-cell Threshold τ

In this section we derive the threshold τ for the $n = 2$ case.

Proposition EC.4. *The searcher should engage if and only if $p_1 \geq \tau$ for the following τ:*

$$
\tau =
\begin{cases}
0.5 & \text{if } \quad \frac{1}{2} \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q \\[2mm]
\sup\{p_1 \mid 1 - p_1 = h(p_1, q, \alpha, \rho), \quad p_1 \in [0.5, 1]\} & \text{otherwise,}
\end{cases}
$$

where

$$
h(p_1, q, \alpha, \rho) = \alpha + G_U\left(\frac{1}{1+\rho}, q, 1, B^*(p_1)\right)\left(-\alpha p_1 - (\alpha - 1)(1 - p_1)\frac{1}{\gamma}\right)
$$

$$
+ G_L\left(\frac{1}{1+\rho}, q, 1, B^*(p_1)\right)\left(-\alpha(1 - p_1)\gamma^{B^*(p_1)} - (\alpha - 1)p_1\right) \quad \text{(EC.27)}
$$

$$
B^*(p_1) = -\left\lceil \frac{2\log\left(\frac{1-p_1}{p_1}\right)}{\log\gamma} \right\rceil \quad \text{(EC.28)}
$$

$$
G_U(x, q, A, B) = \frac{f_1(x, q)^B - f_2(x, q)^B}{f_1(x, q)^{A+B} - f_2(x, q)^{A+B}} \quad \text{(EC.29)}
$$

$$
G_L(x, q, A, B) = \left(\frac{1-q}{q}\right)^B \frac{f_1(x, q)^A - f_2(x, q)^A}{f_1(x, q)^{A+B} - f_2(x, q)^{A+B}} \quad \text{(EC.30)}
$$

$$
f_1(x, q) = \frac{1 + \sqrt{1 - 4q(1 - q)x^2}}{2qx} \quad \text{(EC.31)}
$$

$$
f_2(x, q) = \frac{1 - \sqrt{1 - 4q(1 - q)x^2}}{2qx}. \quad \text{(EC.32)}
$$

The functions $G_U(\cdot)$, $G_L(\cdot)$, $f_1(\cdot)$, $f_2(\cdot)$ in equations (EC.29)–(EC.32) all relate to calculations involving the gambler's ruin problem. To derive τ we view the searcher's stopping problem as choosing parameters of a gambler's ruin problem, and then we leverage off the rich results from that area.

We prove Proposition EC.4 in two parts by considering the two cases separately. In Appendix D.1 we present Lemma EC.2, which examines the case where $\tau = 0.5$; in this situation the searcher engages immediately for any state vector p. In Lemma EC.3, appearing in Appendix D.2, we consider the general case where $\tau > 0.5$.

D.1 $\tau = 0.5$ case: always engage

Lemma EC.2. $\tau = 0.5$ *if and only if* $\frac{1}{2} \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q$.

Proof. We complete the proof by considering the the uniform state vector $p = \left(\frac{1}{2}, \frac{1}{2}\right)$ and manipulating the cost equation in (11). In this special case, $C(p^{(+1)}) = C(p^{(+2)})$ by symmetry. Therefore we can write the cost function in state $p = \left(\frac{1}{2}, \frac{1}{2}\right)$ as:

$$
C(p) = \min\left(\frac{1}{2}, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}C(p^{(+1)})\right).
$$

Showing that the searcher should engage in the uniform case if and only if $\frac{1}{2} \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q$ completes the proof.

First we argue that if $\frac{1}{2} < \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q$, the searcher should wait and hence $\tau > 0.5$. To see this consider the myopic policy where the searcher either engages or receives one more tip and then engages. In the uniform case $p_1^{(+1)} = \frac{\gamma}{\gamma+1} = q$. Using the myopic policy, the searcher should wait if

$$
\frac{1}{2} > \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-q). \tag{EC.33}
$$

Rearranging reveals the desired condition:

$$
\text{wait if} \quad \frac{1}{2} < \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q. \tag{EC.34}
$$

If condition (EC.34) holds, then the myopic policy produces a lower cost than engaging. Since we have found a wait strategy (myopic policy) that produces a lower cost than engaging, the

optimal policy will also wait. Thus, if condition EC.34 holds $\tau > 0.5$.

Next we show that if $\frac{1}{2} \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q$, the searcher should engage and hence $\tau = 0.5$. First assume the condition holds:

$$\frac{1}{2} \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q. \tag{EC.35}$$

By (EC.33) and (EC.34), condition (EC.35) is equivalent to

$$\frac{1}{2} \leq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-q). \tag{EC.36}$$

Condition (EC.36) states that engaging now produces a lower cost than obtaining one more tip and then engaging. We now proceed to show that if condition (EC.36) (and hence (EC.35)) holds the searcher should engage. We do this via contradiction. Assume (EC.36) holds but the searcher should wait. Assume the searcher waits until he receives s_1 tips for cell 1 and s_2 tips for cell 2. Define $\tilde{p} = (\tilde{p}_1, \tilde{p}_2)$ as the new state after receiving these tips. In this case we can write the $\tilde{p}$ in terms of the tip-differential between s_1 and s_2

$$\tilde{p}_1 = \frac{\gamma^{s_1}}{\gamma^{s_1} + \gamma^{s_2}} = \frac{\gamma^{s_1-s_2}}{\gamma^{s_1-s_2}+1}\tilde{p}_2 \qquad = \frac{\gamma^{s_2}}{\gamma^{s_1} + \gamma^{s_2}} = \frac{\gamma^{s_2-s_1}}{\gamma^{s_2-s_1}+1}.$$

If we define $\Delta = |s_1 - s_2|$, then the cost to engage in state $\tilde{p}$ is $\frac{1}{\gamma^{\Delta}+1}$. We can now view our threshold policy τ in terms of a threshold on the tip-differential Δ. By the optimality of a threshold policy, if the searcher should wait in the uniform state, there exists some integer $M \geq 1$ such that the searcher should wait for $\Delta \leq M - 1$ but engage when $\Delta = M$. Define $p(m)$ to be the state when the tip-differential in favor of cell 1 is m:

$$p_1(m) = \frac{\gamma^m}{\gamma^m + 1}, \qquad p_2(m) = \frac{1}{\gamma^m + 1}.$$

By assumption we have that the searcher should wait in state $p(M-1)$ but engage in state $p(M)$. Consequently we have the following relationship (by definition of the cost function in

(11))

$$\frac{1}{\gamma^{M-1}+1} > \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(r_1(p(M-1))\frac{1}{\gamma^M+1} + (1-r_1(p(M-1)))C(p(M-2))). \quad \text{(EC.37)}$$

Because $C(p)$ is a concave function that achieves a maximum at $C(1/2)$ (see Proposition EC.1) and $C(1) = 0$, $C(p)$ decreases over the interval $p \in [1/2, 1]$. Hence $C(p(M-2)) \geq C(p(M)) = \frac{1}{\gamma^M+1}$. Substituting this relationship into (EC.37) produces

$$\frac{1}{\gamma^{M-1}+1} > \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\frac{1}{\gamma^M+1}. \quad \text{(EC.38)}$$

We now proceed to show that conditions (EC.36) and (EC.38) are incompatible. Let us rewrite these conditions to highlight their connection

$$\frac{1}{2} - \frac{1}{1+\rho}(1-q) \leq \frac{\rho}{1+\rho}\alpha \quad \text{(EC.39)}$$

$$\frac{1}{\gamma^{M-1}+1} - \frac{1}{1+\rho}\frac{1}{\gamma^M+1} > \frac{\rho}{1+\rho}\alpha, \quad \text{(EC.40)}$$

where (EC.39) corresponds to (EC.36) and (EC.40) corresponds to (EC.38). Condition (EC.39) is an assumption and condition (EC.40) follows because we assume that the searcher waits in the uniform state and engages when $\Delta = M \geq 1$. The left-hand sides of (EC.39) and (EC.40) have the form $g(i)$, where

$$g(i) = \frac{1}{\gamma^{i-1}+1} - \frac{1}{1+\rho}\frac{1}{\gamma^i+1}.$$

The left-hand side of (EC.39) corresponds to $g(1)$ and (EC.40) corresponds to $g(M)$. Because we assume that the myopic policy of receiving one tip and engaging (i.e., $M = 1$) produces a higher cost than engaging immediately (condition (EC.36)), we must have $M > 1$. Inspection

of $g(i)$ reveals it is a strictly decreasing function. Consequently

$$g(M) < g(1)$$

$$\frac{1}{\gamma^{M-1}+1} - \frac{1}{1+\rho}\frac{1}{\gamma^M+1} < \frac{1}{2} - \frac{1}{1+\rho}(1-q),$$

which implies an incompatibility between conditions (EC.39) (equivalently (EC.35) or (EC.36)) and (EC.40) (equivalently (EC.38)). Therefore if condition (EC.35) holds the searcher must engage, and $\tau = 0.5$.

$\square$

D.2 $\tau > 0.5$ case

In this section we consider the general case of Proposition EC.4. The proof is involved with many lemmas to break it into more manageable pieces. It relies on analyzing the tip-differential between cell 1 and cell 2 as a random walk. A threshold policy on p_1 can be viewed as defining absorbing boundaries on this random walk. We then leverage gambler's ruin machinery to derive the quantities of interest. Lemma EC.3 states the desired result.

Lemma EC.3. *If* $\frac{1}{2} < \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}q$, *the optimal threshold is given by*

$$\tau = \sup\{p_1 \mid 1 - p_1 = h(p_1, q, \alpha, \rho), \quad p_1 \in [0.5, 1]\},$$

where equation (EC.27) defines $h(\cdot)$.

Proof. To start we consider a general threshold policy given by parameter $T \in (1/2, 1)$. The searcher should wait until the first time that $p_1 \geq T$ and then the searcher engages. We use T to denote an arbitrary threshold policy and τ to denote the optimal threshold policy. In Lemma EC.4 we show that for a given p_1, the cost associated with using threshold policy T

is given by

$$
g(T, p_1, q, \alpha, \rho) = \begin{cases} 1 - p_1 & \text{if } T \leq p_1 \\ \alpha + G_U\left(\frac{1}{1+\rho}, q, A_T, B_T\right)\left(-\alpha p_1 - (\alpha - 1)(1 - p_1)\frac{1}{\gamma^{A_T}}\right) & \\ + G_L\left(\frac{1}{1+\rho}, q, A_T, B_T\right)\left(-\alpha(1 - p_1)\gamma^{B_T} - (\alpha - 1)p_1\right) & \text{if } T > p_1, \end{cases} \quad \text{(EC.41)}
$$

where $G_U(\cdot)$ is defined in (EC.29), $G_L(\cdot)$ is defined in (EC.30), and

$$
A_T = \left\lceil \frac{\log\frac{T(1-p_1)}{(1-T)p_1}}{\log\gamma} \right\rceil \quad \text{(EC.42)}
$$

$$
B_T = -\left\lfloor \frac{\log\frac{(1-T)(1-p_1)}{Tp_1}}{\log\gamma} \right\rfloor. \quad \text{(EC.43)}
$$

A_T is the minimum tip-differential in favor of cell 1 required for the updated posterior probability of cell 1 to exceed the threshold T. B_T is the minimum differential in favor of cell 2 required for the updated posterior probability of cell 2 to exceed the threshold T. $B_T \geq A_T$ because $p_1 \geq 1/2$. For a given p_1 we can find an optimal threshold by solving $\min_T g(T, p_1, q, \alpha, \rho)$. For any particular p_1 the optimal threshold is not unique because of the discrete nature of tips. A range of T (depending upon the p_1 of interest) around the true optimal threshold τ will produce optimal results.

Next comparing (EC.27) and (EC.41) we see that

$$
h(p_1, q, \alpha, \rho) = \lim_{T \downarrow p_1} g(T, p_1, q, \alpha, \rho).
$$

That is we can view $h(p_1, q, \alpha, \rho)$ as the expected cost when we use a threshold infinitesimally greater than p_1. Using this policy the searcher will engage with a tip-differential of one in favor of cell 1 or a tip-differential of $B^*(p_1)$ (see equation (EC.28)) in favor of cell 2. Recall that $B^*(p_1)$ is the minimum tip-differential in favor of cell 2 such that the updated posterior probability for cell 2 exceeds p_1. Using the parameter Δ defined in Lemma EC.4 as the tip-differential in favor of cell 1, then $h(p_1, q, \alpha, \rho)$ is the expected cost from using the policy

"wait until $\Delta = 1$ or $\Delta = -B^*(p_1)$". The way to think about the result for this proposition is to compare the cost to engage $(1 - p_1)$ to the cost of using the policy "wait until $\Delta = 1$ or $\Delta = -B^*(p_1)$" ($h(p_1, q, \alpha, \rho)$).

To complete the proof, we first show in Lemma EC.5 that the optimal threshold τ satisfies

$$\tau = \sup\{p_1 \mid 1 - p_1 > h(p_1, q, \alpha, \rho), \quad p_1 \in [0.5, 1]\}. \tag{EC.44}$$

In Lemma EC.6 we show that this τ satisfies $(1 - \tau) = h(\tau, q, \alpha, \rho)$, and finally in Lemma EC.7 we show that

$$1 - p_1 < h(p_1, q, \alpha, \rho) \quad \forall \tau < p_1 \leq 1.$$

Combining the pieces from Lemmas EC.5–EC.7 produces the desired result:

$$\tau = \sup\{p_1 \mid 1 - p_1 = h(p_1, q, \alpha, \rho), \quad p_1 \in [0.5, 1]\}.$$

Note the expression for $h(p_1, q, \alpha, \rho)$ simplifies significantly for $B^*(p_1) = 1$. As we approach $p_1 = 1/2$, $B^*(p_1)$ eventually equals 1 for any values of q, α, ρ. When $B^*(p_1) = 1$, the searcher uses the myopic policy: engage after the first tip. The cost in this case simplifies to:

$$\lim_{p_1 \downarrow 1/2} h(p_1, q, \alpha, \rho) = \left(1 - \left(\frac{\rho}{1 + \rho}(1 - \alpha) + \frac{1}{1 + \rho}q\right)\right). \tag{EC.45}$$

Condition (EC.45) is similar to the condition for whether $\tau = 0.5$ (e.g., see condition (EC.34)). This leads to the relationship

$$\lim_{p_1 \downarrow 1/2} h(p_1, q, \alpha, \rho) < \frac{1}{2} \iff \tau > 0.5.$$

Thus the conditions for the $\tau > 0.5$ case and the $\tau = 0.5$ are consistent.

$\square$

Lemma EC.4. *For a given p_1, the expected cost associated with using threshold policy T is given by the function $g(T, p_1, q, \alpha, \rho)$ defined by equation (EC.41).*

Proof. Recall that we can update the state probability vector p after receiving M tips if we know the number of tips for cell 1 (s_1) and the number of tips for cell 2 ($s_2 = M - s_1$). The updated probabilities only depend upon the tip-differential $s_1 - s_2$ (see equations (13)–(14)). If we define $\Delta = s_1 - s_2$, then we can transform the threshold policy from "engage as soon as $p_1 \geq T$" to a threshold policy on Δ: engage as soon as $\Delta \geq A_T$ or $\Delta \leq -B_T$. The specific formulas for A_T and B_T appear in equations (EC.42)–(EC.43); both B_T and A_T depend upon the threshold T and the initial state vector p. If $T \leq p_1$ then $A_T \leq 0$ and the searcher should engage immediately. If $T > p_1$ then $A_T > 0$ and the searcher should wait. Because both $p_1 \geq 1/2$ and $T > 1/2$, we have that $B_T > 0$. For the remainder of this Lemma, we will assume that $A_T > 0$ so that the searcher does not trivially engage immediately.

Given an initial state vector p and threshold T, we compute A_T and B_T from equations (EC.42)–(EC.43). The searcher will continue receiving tips until one of three events occur: the target attacks, the tip-differential reaches $\Delta = A_T$ (searcher engages cell 1), the tip-differential reaches $\Delta = -B_T$ (searcher engages cell 2). To compute the expected cost of using the threshold policy T, we need to determine the probabilities of each of these three events occurring and their associated costs.

If we condition on the location of the target, then Δ proceeds according to a random walk. If cell 1 (cell 2) contains the target, then Δ increases by 1 at each step with probability q $(1-q)$ and decreases by 1 with probability $1-q$ (q). We utilize random walk and gambler's ruin results to examine when the searcher will engage and whether the searcher engages when $\Delta = A_T$ or $\Delta = -B_T$. For the given threshold T, we can define a stopping time N_T when the searcher engages. Define Δ_i to be the value of the tip-differential Δ after the ith tip, so $\Delta_0 = 0$. We then define the stopping time N_T as

$$N_T = \inf\{i \mid \Delta_i \in \{-B_T, A_T\}\}.$$

If we further define Y_n as the cost incurred after receiving n tips, then the random cost incurred from using the stopping rule N_T is given by

$$Y_{N_T} = \alpha I(\text{attack occurs before tip } N_T)$$
$$+ I(\text{attack occurs after tip } N_T)I(\text{engage wrong cell}).$$

We now compute our desired quantity: the expected cost from using threshold policy T

$$\mathbf{E}[Y_{N_T}] = \alpha\mathbf{P}[\text{attack occurs before tip } N_T]$$
$$+ \mathbf{P}[\text{attack occurs after tip } N_T \text{ and choose wrong cell}]. \tag{EC.46}$$

The false positive cost in (EC.46) depends on the actual cell location of the target and whether the Δ random walk stops at $-B_T$ or A_T:

$$\mathbf{P}[\text{attack occurs after tip } N_T \text{ and choose wrong cell}]$$
$$= p_1\mathbf{P}[\text{attack occurs after tip } N_T, \Delta_{N_T} = -B_T \mid \text{ target in cell } 1]$$
$$+ (1 - p_1)\mathbf{P}[\text{attack occurs after tip } N_T, \Delta_{N_T} = A_T \mid \text{ target in cell } 2]. \tag{EC.47}$$

Focusing on the first term in (EC.47), we condition on the value of N_T:

$$\mathbf{P}[\text{attack occurs after tip } N_T, \Delta_{N_T} = -B_T \mid \text{ target in cell } 1]$$
$$= \sum_{i=0}^{\infty} \left(\frac{1}{1+\rho}\right)^i \mathbf{P}[N_T = i, \Delta_{N_T} = -B_T \mid \text{ target in cell } 1]. \tag{EC.48}$$

Fortunately, explicit expressions exist for the generating function term in (EC.48). Boccio (2012) presents several gambler's ruin results, including the generating functions relevant to our problem. Utilizing equation (45) of Boccio (2012) we can write

$$\mathbf{P}[\text{attack occurs after tip } N_T, \Delta_{N_T} = -B_T \mid \text{target in cell 1}] = G_L\left(\frac{1}{1+\rho}, q, A_T, B_T\right) \tag{EC.49}$$

$$\mathbf{P}[\text{attack occurs after tip } N_T, \Delta_{N_T} = -B_T \mid \text{target in cell 2}] = G_L\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right), \tag{EC.50}$$

where $G_L(\cdot)$ is defined in (EC.30). Furthermore, we can use equation (46) of Boccio (2012) to express

$$\mathbf{P}[\text{attack occurs after tip } N_T, \Delta_{N_T} = A_T \mid \text{target in cell 1}] = G_U\left(\frac{1}{1+\rho}, q, A_T, B_T\right) \tag{EC.51}$$

$$\mathbf{P}[\text{attack occurs after tip } N_T, \Delta_{N_T} = A_T \mid \text{target in cell 2}] = G_U\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right), \tag{EC.52}$$

where $G_U(\cdot)$ is defined in (EC.29). By combining equations (EC.49) (equivalently (EC.50)) and (EC.51) (equivalently (EC.52)) we produce the probability the attack occurs after engaging, conditioned on the location of the cell:

$$\mathbf{P}[\text{attack occurs after tip } N_T \mid \text{target in cell 1}] = G_L\left(\frac{1}{1+\rho}, q, A_T, B_T\right)$$
$$+ G_U\left(\frac{1}{1+\rho}, q, A_T, B_T\right) \tag{EC.53}$$

$$\mathbf{P}[\text{attack occurs after tip } N_T \mid \text{target in cell 2}] = G_L\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right).$$
$$+ G_U\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right) \tag{EC.54}$$

Substituting equations (EC.47), (EC.48), (EC.49), (EC.52), (EC.53), and (EC.54) into (EC.46)

produces the complete expression

$$\mathbf{E}[Y_{N_T}] =$$

$$p_1\left(\alpha\left(1 - \left(\underbrace{G_U\left(\frac{1}{1+\rho}, q, A_T, B_T\right)}_{\substack{\text{prob reach } A_T \text{ before attack.}\\ \text{(choose cell 1 correctly)}}} + \underbrace{G_L\left(\frac{1}{1+\rho}, q, A_T, B_T\right)}_{\substack{\text{prob reach } B_T \text{ before attack.}\\ \text{(choose cell 2 incorrectly)}}}\right)\right) + \underbrace{G_L\left(\frac{1}{1+\rho}, q, A_T, B_T\right)}_{\substack{\text{prob reach } B_T \text{ before attack.}\\ \text{(choose cell 2 incorrectly)}}}\right)$$

$$\underbrace{\phantom{G_U\left(\frac{1}{1+\rho}, q, A_T, B_T\right) + G_L\left(\frac{1}{1+\rho}, q, A_T, B_T\right)}}_{\text{prob engage before attack}}$$

$$+(1-p_1)\left(\alpha\left(1 - \left(\underbrace{G_U\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right)}_{\substack{\text{prob reach } A_T \text{ before attack.}\\ \text{(choose cell 1 incorrectly)}}} + \underbrace{G_L\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right)}_{\substack{\text{prob reach } B_T \text{ before attack.}\\ \text{(choose cell 2 correctly)}}}\right)\right)\right.$$

$$\underbrace{\phantom{G_U\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right) + G_L\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right)}}_{\text{prob engage before attack}}$$

$$\left.+\underbrace{G_U\left(\frac{1}{1+\rho}, 1-q, A_T, B_T\right)}_{\substack{\text{prob reach } A_T \text{ before attack.}\\ \text{(choose cell 1 incorrectly)}}}\right). \tag{EC.55}$$

The first line of (EC.55) conditions on the target being in cell 1 and the second and third lines condition on the target being in cell 2. Examination of equations (EC.29)–(EC.32) reveals the following relationships

$$G_U(x, q, A, B) = \gamma^A G_U(x, 1-q, A, B) \tag{EC.56}$$

$$G_L(x, q, A, B) = \frac{1}{\gamma^B} G_L(x, 1-q, A, B). \tag{EC.57}$$

Using (EC.56) and (EC.57) we simplify (EC.55) to our final expression

$$
\mathbf{E}[Y_{N_T}] = \begin{cases} 1 - p_1 & \text{if } T \le p_1 \quad \text{engage immediately} \\ \alpha + G_U\left(\frac{1}{1+\rho}, q, A_T, B_T\right)\left(-\alpha p_1 - (\alpha - 1)(1 - p_1)\frac{1}{\gamma^{A_T}}\right) \\ \quad + G_L\left(\frac{1}{1+\rho}, q, A_T, B_T\right)\left(-\alpha(1 - p_1)\gamma^{B_T} - (\alpha - 1)p_1\right) & \text{if } T > p_1. \end{cases}
$$

$$(EC.58)$$

Denoting $g(T, p_1, q, \alpha, \rho) \equiv \mathbf{E}[Y_{N_T}]$ completes the proof.

$\square$

Lemma EC.5. *The optimal threshold τ is defined by*

$$\tau = \sup\{p_1 \mid 1 - p_1 > h(p_1, q, \alpha, \rho), \quad p_1 \in [0.5, 1]\}.$$

Proof. Consider the set

$$D = \{x \mid 1 - x > h(x, q, \alpha, \rho), \quad x \in [0.5, 1]\}.$$

If $p_1 \in D$, then it is optimal to wait because we found a wait-policy that that performs better than engaging. Namely the policy that waits until $\Delta = 1$ or $\Delta = -B^*(p_1)$ achieves a lower cost $(h(p_1, q, \alpha, \rho))$ than engaging $(1 - p_1)$.

We next define

$$\tilde{\tau} = \sup D$$

and argue that $\tilde{\tau} \le \tau$. If $\tilde{\tau} \in D$, then as argued above it is optimal to wait for $p_1 = \tilde{\tau}$. Since the optimal policy is a threshold policy it must follow that $\tau > \tilde{\tau}$. If $\tilde{\tau} \notin D$, then by definition of the supremum, $\tilde{\tau}$ must be a limit point of D. In this case, if $\tilde{\tau} > \tau$, then there would exist some $\hat{p} \in (\tau, \tilde{\tau})$ and $\hat{p} \in D$ because $\tilde{\tau}$ is a limit point of D. This implies it is

optimal to wait for $p_1 = \hat{p} \in D$, but $\hat{p}$ is greater than the engage threshold τ by construction. This is a contradiction. Hence $\tilde{\tau} \leq \tau$

We next argue $\tilde{\tau} \geq \tau$. We proceed by contradiction and assume $\tilde{\tau} < \tau$. We will find a $\tilde{p} \in (\tilde{\tau}, \tau)$ such that the policy "wait until $\Delta = 1$ or $\Delta = -B^*(\tilde{p})$" is the optimal policy. Since $h(\tilde{p}, q, \alpha, \rho)$ is the cost associated with this policy, then $h(\tilde{p}, q, \alpha, \rho) < (1 - \tilde{p})$. That is if it is optimal to wait, then the optimal cost to wait must be strictly less than the cost to engage. If that is the case then $\tilde{p} \in D$. However, $\tilde{p} > \tilde{\tau} = \sup D$, which leads to a contradiction. Thus $\tilde{\tau} \geq \tau$. It remains to find such a $\tilde{p}$.

Any $\tilde{p}$ such that $\tilde{p} \in (L, \tau)$ where $L = \max(\tilde{\tau}, L_1, L_2, L_3)$ and

$$L_1 = \frac{\tau}{\gamma(1 - \tau) + \tau}, \tag{EC.59}$$

$$L_2 = \frac{1}{1 + \exp\left(-\frac{(B^*(\tau)-1)\log\gamma}{2}\right)}, \tag{EC.60}$$

$$L_3 = \frac{\gamma^{B^*(\tau)-1}(1 - \tau)}{\gamma^{B^*(\tau)-1}(1 - \tau) + \tau}, \tag{EC.61}$$

will suffice. By construction all are less than τ. $L_1 < \tau$ by inspection because $\gamma > 1$. L_1 ensures that one additional tip for cell 1 will push the updated maximum posterior probability above the optimal threshold τ, where it is optimal to engage. L_2 is constructed such that $B^*(L_2) = B^*(\tau) - 1$. This condition, coupled with inspection of (EC.28), reveals that $L_2 < \tau$. We require $L_2 < \tilde{p}$ because this ensures that $B^*(\tilde{p}) = B^*(\tau)$. Finally by definition, $B^*(\tau)$ (see (EC.28)) is the minimum tip-differential in favor of cell 2 such that the updated posterior probability for cell 2 exceeds τ. Therefore,

$$L_3 = \frac{\gamma^{B^*(\tau)-1}(1 - \tau)}{\gamma^{B^*(\tau)-1}(1 - \tau) + \tau} < \tau \leq \frac{\gamma^{B^*(\tau)}(1 - \tau)}{\gamma^{B^*(\tau)}(1 - \tau) + \tau}. \tag{EC.62}$$

Combining $L_2, L_3 < \tilde{p}$ we see that a tip-differential of $B^*(\tau) - 1 = B^*(\tilde{p}) - 1$ in favor of cell

2 is not enough to push the posterior for cell 2 above τ starting from $\tilde{p}$:

$$\frac{\gamma^{B^*(\tau)-1}(1-\tilde{p})}{\gamma^{B^*(\tau)-1}(1-\tilde{p})+\tilde{p}} = \frac{\gamma^{B^*(\tilde{p})-1}(1-\tilde{p})}{\gamma^{B^*(\tilde{p})-1}(1-\tilde{p})+\tilde{p}} \qquad \text{by } L_2 < \tilde{p} \tag{EC.63}$$

$$< \tau \qquad \text{by } L_3 < \tilde{p}. \tag{EC.64}$$

However, because $\tilde{p} < \tau$ and $B^*(\tilde{p}) = B^*(\tau)$ (by $L_2 < \tilde{p}$) a tip-differential of $B^*(\tilde{p}) = B^*(\tau)$ in favor of cell 2 does push the posterior for cell 2 above τ:

$$\tau < \frac{\gamma^{B^*(\tilde{p})}(1-\tilde{p})}{\gamma^{B^*(\tilde{p})}(1-\tilde{p})+\tilde{p}}. \tag{EC.65}$$

Combining conditions (EC.64) and (EC.65) produces

$$\frac{\gamma^{B^*(\tilde{p})-1}(1-\tilde{p})}{\gamma^{B^*(\tilde{p})-1}(1-\tilde{p})+\tilde{p}} < \tau < \frac{\gamma^{B^*(\tilde{p})}(1-\tilde{p})}{\gamma^{B^*(\tilde{p})}(1-\tilde{p})+\tilde{p}}. \tag{EC.66}$$

Condition (EC.66) states that, starting from $p_1 = \tilde{p}$, the updated posterior first exceeds the optimal threshold τ with a tip-differential in favor of cell 2 when $\Delta = -B^*(\tilde{p})$. Thus for $\Delta \in [-B^*(\tilde{p}) + 1, 0]$ the maximum posterior probability is below the optimal threshold τ and the searcher should wait. The searcher should engage when Δ first hits either $\Delta = 1$ or $\Delta = -B^*(\tilde{p})$. Engaging for $\Delta \in \{-B^*(\tilde{p}), 1\}$ produces a maximum posterior probability in excess of τ. We have completed the proof because we have found a $\tilde{p} \in (\tilde{\tau}, \tau)$ where it is optimal to wait and has optimal cost of $h(\tilde{p}, q, \alpha, \rho)$, which must be less than the cost to engage $(1 - \tilde{p})$. Therefore, we have produced a contradiction and so $\tilde{\tau} \geq \tau$.

We have shown that $\tilde{\tau} \leq \tau$ and $\tilde{\tau} \geq \tau$ and thus our proof is complete:

$$\tau = \tilde{\tau} = \sup D = \sup\{p_1 \mid 1 - p_1 > h(p_1, q, \alpha, \rho), \quad p_1 \in [0.5, 1]\}.$$

$\square$

Lemma EC.6. *The optimal τ defined by equation (EC.44) satisfies $(1 - \tau) = h(\tau, q, \alpha, \rho)$.*

Proof. In Lemma EC.5 we showed that the optimal threshold τ is defined by

$$\tau = \sup\{p_1 \mid 1 - p_1 > h(p_1, q, \alpha, \rho), \quad p_1 \in [0.5, 1]\}.$$

Now we show that the threshold satisfies $(1 - \tau) = h(\tau, q, \alpha, \rho)$. The quantity $h(\tau, q, \alpha, \rho)$ is the cost associated with the policy "wait until $\Delta = 1$ or $\Delta = -B^*(\tau)$". By definition it is optimal for the searcher to engage at $p_1 = \tau$. However, if the searcher chose to wait in an optimal fashion, then he would use the policy associated with $h(\tau, q, \alpha, \rho)$: wait until $\Delta = 1$ or $\Delta = -B^*(\tau)$. Thus $h(\tau, q, \alpha, \rho)$ is the optimal cost to wait with $p_1 = \tau$ (corresponding to the second term of (11)). At the threshold, the cost to engage must equal the optimal cost from waiting. This follows from the concavity of the cost function $C(p)$. If the cost to engage did not equal the optimal cost from waiting at the threshold, this would imply a discontinuity in the cost function (impossible by concavity) or that the threshold is incorrect. Therefore $(1 - \tau) = h(\tau, q, \alpha, \rho)$. To see this more rigorously consider the values L_1, L_2, L_3 from (EC.59)–(EC.61). As discussed in the proof for Lemma EC.5, for $p_1 \in (\max(L_1, L_2, L_3), \tau)$ the optimal cost is given by $h(p_1, q, \alpha, \rho)$. However, because $B^*(p_1) = B^*(\tau)$ in this range of interest, $h(p_1, q, \alpha, \rho)$ is just a linear function of p_1. Thus, we can write out the cost function $C(p)$ explicitly for $p_1 > \max(L_1, L_2, L_3)$ in terms of a piecewise linear function

$$C(p) = \begin{cases} \alpha + G_U\left(\frac{1}{1+\rho}, q, 1, B^*(\tau)\right)\left(-\alpha p_1 - (\alpha - 1)(1 - p_1)\frac{1}{\gamma}\right) & \\ \quad + G_L\left(\frac{1}{1+\rho}, q, 1, B^*(\tau)\right)\left(-\alpha(1 - p_1)\gamma^{B^*(\tau)} - (\alpha - 1)p_1\right) & \text{if } \max(L_1, L_2, L_3) < p_1 < \tau \\ 1 - p_1 & \text{if } p_1 \geq \tau. \end{cases}$$

We know $C(p)$ is concave and thus the two linear pieces must intersect at the threshold τ. This completes the proof: $(1 - \tau) = h(\tau, q, \alpha, \rho)$.

$\square$

Lemma EC.7. *For the optimal τ defined by equation (EC.44) we have the following condition*

$$1 - p_1 < h(p_1, q, \alpha, \rho), \quad \forall \tau < p_1 \leq 1.$$

Proof. $h(p_1, q, \alpha, \rho)$ is the cost associated with the policy "wait until $\Delta = 1$ or $\Delta = -B^*(p_1)$". This cost must be greater than the optimal cost to wait from our cost function $C(p)$ (from (11))

$$h(p_1, q, \alpha, \rho) \geq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{2} r_i(p)C(p^{(+i)}).$$

Thus to complete the proof it suffices to show that

$$1 - p_1 < \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{2} r_i(p)C(p^{(+i)}), \quad \forall \tau < p_1 \leq 1. \tag{EC.67}$$

For the region of interest it is always optimal to engage in state $p^{(+1)}$. Thus $C(p^{(+1)}) = (1 - p_1^{(+1)})$. The optimal policy for state $p^{(+2)}$ depends upon how close p_1 is to τ. Let us denote $K(p)$ as the optimal cost to wait and rewrite it

$$\begin{aligned}
K(p) &= \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{2} r_i(p)C(p^{(+i)}) \\
&= \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\left(r_1(p)(1 - p_1^{(+1)}) + r_2(p)(1 - p_1^{(+2)})\right) \\
&\quad - \frac{1}{1+\rho}r_2(p)\left((1 - p_1^{(+2)}) - C(p^{(+2)})\right).
\end{aligned} \tag{EC.68}$$

It is a straightforward exercise to show that

$$r_1(p)(1 - p_1^{(+1)}) + r_2(p)(1 - p_1^{(+2)}) = 1 - p_1. \tag{EC.69}$$

Substituting (EC.69) into (EC.68) yields

$$K(p) = \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-p_1) - \frac{1}{1+\rho}r_2(p)\left((1-p_1^{(+2)}) - C(p^{(+2)})\right).$$

Condition (EC.67) is equivalent to $(1-p_1) < K(p)$. Because τ is the optimal engage threshold, it must hold that $(1-p_1) \leq K(p)$ over the region of interest. Thus to complete the proof we need to show that $(1-p_1) \neq K(p)$ for $p_1 > \tau$. We show this condition holds by contradiction. Let us assume that condition (EC.67) does not hold. Therefore there exists some state $p = (p_1, (1-p_1))$, such that $p_1 > \tau$ and the cost to engage equals the cost to wait: $(1-p_1) = K(p)$. We will show this implies that for the state $\tilde{p} = (p_1 - \epsilon, (1-p_1+\epsilon))$, such that $\tilde{p}_1 = p_1 - \epsilon > \tau$, we have the relationship $(1-\tilde{p}_1) > K(\tilde{p})$. This implies it is optimal to wait in state $\tilde{p}$, which is a contradiction because $\tilde{p}_1$ exceeds the engage threshold: $\tilde{p}_1 = p_1 - \epsilon > \tau$. Let us examine $K(\tilde{p})$

$$
\begin{aligned}
K(\tilde{p}) =& \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-\tilde{p}_1) - \frac{1}{1+\rho}r_2(\tilde{p})\left((1-\tilde{p}_1^{(+2)}) - C(\tilde{p}^{(+2)})\right) \\
=& \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-p_1) - \frac{1}{1+\rho}r_2(\tilde{p})\left((1-\tilde{p}_1^{(+2)}) - C(\tilde{p}^{(+2)})\right) + \frac{1}{1+\rho}\epsilon \\
=& \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-p_1) - \frac{1}{1+\rho}r_2(p)\left((1-p_1^{(+2)}) - C(p^{(+2)})\right) \\
& + \frac{1}{1+\rho}\left(\epsilon - \left[r_2(\tilde{p})\left((1-\tilde{p}_1^{(+2)}) - C(\tilde{p}^{(+2)})\right) - r_2(p)\left((1-p_1^{(+2)}) - C(p^{(+2)})\right)\right]\right) \\
=& K(p) \\
& + \frac{1}{1+\rho}\left(\epsilon - \left[r_2(\tilde{p})\left((1-\tilde{p}_1^{(+2)}) - C(\tilde{p}^{(+2)})\right) - r_2(p)\left((1-p_1^{(+2)}) - C(p^{(+2)})\right)\right]\right). \quad \text{(EC.70)}
\end{aligned}
$$

The final key is showing the term is square brackets is non-negative. To do this first note that

$$r_2(\tilde{p}) > r_2(p) \quad\quad\quad\quad \text{(EC.71)}$$

because $\tilde{p}_1 < p_1$. This follows from the definition of $r_i(\cdot)$ in equation (7) and $q > 1/2$. Next

for any state $\hat{p}$,

$$1 - \hat{p}_1 \geq C(\hat{p}) \tag{EC.72}$$

because $C(\hat{p})$ is the optimal cost and $(1 - \hat{p}_1)$ is greater than or equal to the cost to engage (it only equals the cost to engage if $\hat{p}_1 \geq 1/2$). Furthermore, we have

$$1 - \hat{p}_1 = C(\hat{p}) \quad \text{for } \hat{p} \geq \tau, \tag{EC.73}$$

since it is optimal to engage for $\hat{p} \geq \tau$. Finally by concavity

$$C(\hat{p}) \text{ is concave as we increase } \hat{p}_1. \tag{EC.74}$$

Combining conditions (EC.72)–(EC.74) produces the following condition:

$$(1 - \hat{p}_1) - C(\hat{p}) \text{ is non-increasing as we increase } \hat{p}_1 \text{ over } [0, 1]. \tag{EC.75}$$

Because $\tilde{p}_1 < p_1$ (and hence $\tilde{p}_1^{(+2)} < p_1^{(+2)}$) we can combine conditions (EC.71) and (EC.75) to produce the desired result

$$r_2(\tilde{p}) \left((1 - \tilde{p}_1^{(+2)}) - C(\tilde{p}^{(+2)}) \right) - r_2(p) \left((1 - p_1^{(+2)}) - C(p^{(+2)}) \right) \geq 0. \tag{EC.76}$$

Substituting (EC.76) into (EC.70) yields

$$K(\tilde{p}) \leq K(p) + \frac{1}{1 + \rho} \epsilon.$$

By assumption $K(p) = (1 - p_1)$ and thus we have shown the desired contradictory condition:

$$K(\tilde{p}) \leq (1 - p_1) + \frac{1}{1 + \rho} \epsilon < (1 - p_1 + \epsilon) = (1 - \tilde{p}_1).$$

Consequently, condition (EC.67) must hold and the proof is complete:

$$1 - p_1 < h(p_1, q, \alpha, \rho), \quad \forall \tau < p_1 \le 1.$$

$\square$

E   Expected Number of Tips

The search ends either when the searcher engages or the target attacks. In this section we compute the expected number of tips received during the search, accounting for both possible ending scenarios. For a given set of parameters α, ρ and q we first compute τ from Proposition EC.4. Starting from the uniform state $p = (0.5, 0.5)$, we next examine when the tip-differential Δ will cause the posterior probability to exceed the threshold τ. Inspection of equation (13) reveals that the searcher will engage when the tip-differential Δ first reaches A_τ (and the searcher engages cell 1) or $-A_\tau$ (and the searcher engages cell 2), where

$$A_\tau = \left\lceil \frac{\log \frac{\tau}{1-\tau}}{\log \gamma} \right\rceil. \tag{EC.77}$$

We model Δ_i as the tip-differential after the ith tip with $\Delta_0 = 0$. If we condition on cell 1 containing the target, then Δ_i follows a random walk that increases by 1 with probability q and decreases by 1 with probability $(1-q)$. We define a Markov chain for Δ_i that incorporates the potential for the mature attack by the target. Δ_i proceeds as a random walk on the state space $\{-A_\tau, -(A_\tau - 1), \ldots, -1, 0, 1, \ldots, A_\tau - 1, A_\tau\}$, and we add a state X to denote the target attacking before the next tip. The Markov chain has three absorbing states: $\{-A_\tau, A_\tau, X\}$

and when Δ_i is in a non-absorbing state it transitions according to the following dynamics

$$
\Delta_{i+1} = \begin{cases}
\Delta_i + 1 & \text{with probability } q\frac{1}{1+\rho} \\
\Delta_i - 1 & \text{with probability } (1-q)\frac{1}{1+\rho} \\
X & \text{with probability } \frac{\rho}{1+\rho}.
\end{cases} \tag{EC.78}
$$

The search ends when the Markov chain transitions to one of the absorbing states. We can now define a system of equations to solve for the expected number of tips the searcher receives before the search ends. Denote T_j as the expected number of tips until the search ends when the system starts with $\Delta_0 = j$. We present T_0 in Figure 2. If $\tau = 0.5$ then $T_0 = 0$ because the searcher engages immediately. Otherwise the T_i satisfy the following system of equations

$$
T_i = 1 + q\frac{1}{1+\rho}T_{i-1} + (1-q)\frac{1}{1+\rho}T_{i-1} - \frac{\rho}{1+\rho}, \quad \text{for } -A_\tau < i < A_\tau \tag{EC.79}
$$

$$
T_{A_\tau} = T_{-A_\tau} = 0. \tag{EC.80}
$$

Solving the system of equations given by (EC.79)–(EC.80) is a straighforward linear algebra exercise. Note we subtract $\frac{\rho}{1+\rho}$ in equation (EC.79) because the searcher does not receive a tip if the target attacks. We can condition on cell 2 containing the target and compute a similar expected number of tips and then combine the two quantities to calculate the final unconditional expected number of tips. Not surprisingly the conditional expected number of tips is the same for the two locations.

F $n = \infty$

In Appendix F.1 we provide the mathematical motivation for the problem. We prove Proposition 2 in Appendix F.2. Finally in Appendix F.3 we compute the costs for the finite n heuristic and evaluate its performance.

F.1 Motivation

Assume we have n cells and consider a maximum of $m < n$ tips provided to the searcher. We will show in this section that for large enough m and n, the searcher will receive a second tip for some cell by tip m with high probability, and the searcher will engage that cell once he receives the repeat tip.

First we compute the probability that the informant provides no repeat tips in the first m tips. Define X_m as the number of tips pointing to cell 1, which is a binomial random variable with parameters m and q. The probability of no repeats in any cell in the first m tips is less than the probability of no repeats to cell 1 in the first m tips. Therefore

$$\mathbf{P}[\text{no repeats in } m \text{ tips}] \leq P[X_m \leq 1] = (1-q)^m + mq(1-q)^{m-1} \leq 2m(1-q)^{m-1} \to 0 \quad \text{(EC.81)}$$

Thus for any $\epsilon > 0$, we can use condition (EC.81) to find an m_0 such that for any $m_0 < m < n$ we have $\mathbf{P}[\text{no repeats in } m \text{ tips}] < \epsilon$.

Next we will show that for large enough n relative to our choice of $m > m_0$, the searcher will engage if some cell receives its second tip prior to tip m. Assume that tip m is the first time a cell receives a second tip. Thus the first $m-1$ tips point to distinct cell. Since we assume a uniform prior, the posterior of cell 1, which receives the second tip, is

$$p_1 = \frac{\gamma^2}{\gamma^2 + (m-2)\gamma + (n-m+1)} = \frac{\left(\frac{q}{1-q}\right)^2 (n-1)}{\left(\frac{q}{1-q}\right)^2 (n-1) + (m-2)\frac{q}{1-q} + \frac{n-m+1}{n-1}} \quad \text{(EC.82)}$$

This follows directly from equation (4). Clearly, the expression for p_1 in (EC.81) converges to 1 as $n \to \infty$ for fixed q and m. Consequently for some $n > n_0$, p_1 will exceed the condition to engage in (15) and the searcher will engage no later than the first time he receives a repeat tip as long as the first repeat tip occurs before m. Note of course, it may be optimal to engage before that.

Putting the pieces together, we have for any $m > m_0$ that there is less than probability

ϵ that we will receive more than m tips before our first repeat. We can then use equation (EC.82) to show that if and when that repeat occurs, the searcher will engage if $n > n_0, m$.

This motivates the infinite cell case: for large n with high probability the searcher will engage no later than when he receives the first repeat tip.

F.2 Proof of Proposition 2

We re-list the three possible optimal alternatives for convenience

1. Engage any cell: cost 1

2. Obtain one tip and engage the corresponding cell: cost $\frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-q)$

3. Engage the first cell to receive two tips: cost $\alpha\left(1 - \left(\frac{q}{\rho+q}\right)^2\right)$.

To complete the proof we must complete three tasks: (1) derive the costs provided for each option, (2) show that any other option is suboptimal, (3) show the minimal cost option corresponds to the conditions in Proposition 2. We perform task 1 in Appendix F.2.1, task 2 in Appendix F.2.2, and task 3 in Appendix F.2.3.

F.2.1 Cost Associated with Three Options

If the searcher engages immediately (option 1), he will not detect the correct cell and thus collateral damage is guaranteed, yielding a cost of 1.

If the searcher chooses option 2, then with probability $\frac{\rho}{1+\rho}$ the target attacks, which generates a cost of α. With probability $\frac{1}{1+\rho}$ the searcher receives the tip. The informant points to the correct cell with probability q. Thus if the the searcher engages after this tip, he will choose incorrectly with probability $(1-q)$, causing collateral damage. This completes the derivation of the cost for option 2. Another way to compute this cost is to start with finite n and uniform prior. After one tip, the state transitions from $p = (1/n, 1/n, \ldots, 1/n)$ to $p^{+1} = (q, (1-q)/(n-1), \ldots, (1-q)/(n-1))$. See equation (5) for the transition dynamics. Since $p_1^{+1} = q$ for any n, this will also hold in the limit as $n \to \infty$.

The derivation for option 3 is slightly more complicated. If the searcher waits until he receives a confirming tip, then he know with certainty that cell contains the target. The only cost that can occur for option 3 is if the target attacks before the searcher receives the confirming tip. Therefore, we proceed to compute the probability the searcher will receive the confirming tip before the target attacks.

We first compute the probability the searcher will receive the confirming tip on tip $m \geq 2$. For this to occur, the mth tip points to cell 1 (prob q), and in the first $m-1$ tips, $m-2$ tips point to a wrong cell (probability $(1-q)^{m-2}$) and one points to cell 1 (probability q). Since the first tip pointing to cell 1 can occur in any of the first $m-1$ tips, the probability that the confirming tip occurs on the mth tip is $(m-1)q(1-q)^{m-2}q$. Finally, the target's attack must not occur before the mth tip, which happens with probability $1/(1+\rho)^m$. We now sum over all possible values of m for the final probability:

$$
\begin{aligned}
\mathbf{P}[\text{confirming tip before attack}] &= \sum_{i=2}^{\infty}(i-1)q(1-q)^{i-2}q\left(\frac{1}{1+\rho}\right)^i \\
&= \frac{q^2}{1-q}\left(\frac{1}{1+\rho}\right)\sum_{i=2}^{\infty}(i-1)(1-q)^{i-1}\left(\frac{1}{1+\rho}\right)^{i-1} \\
&= \frac{q^2}{1-q}\left(\frac{1}{1+\rho}\right)\sum_{i=2}^{\infty}(i-1)\left((1-q)\frac{1}{1+\rho}\right)^{i-1} \\
&= \frac{q^2}{1-q}\left(\frac{1}{1+\rho}\right)\sum_{i=1}^{\infty}i\left((1-q)\frac{1}{1+\rho}\right)^i.
\end{aligned}
$$

Using geometric series results this simplifies to

$$
\begin{aligned}
\mathbf{P}[\text{confirming tip before attack}] &= \frac{q^2}{1-q}\frac{1}{1+\rho}\frac{(1-q)\frac{1}{1+\rho}}{\left(1-(1-q)\frac{1}{1+\rho}\right)^2} \\
&= \left(\frac{q}{\rho+q}\right)^2.
\end{aligned}
$$

Thus the expected cost for choosing option 3 is the desired result:

$$\mathbf{E}[\text{cost of option 3}] = \alpha\mathbf{P}[\text{attack occurs before confirming tip}]$$

$$= \alpha(1 - \mathbf{P}[\text{confirming tip before attack}])$$

$$= \alpha\left(1 - \left(\frac{q}{\rho + q}\right)^2\right).$$

F.2.2 Suboptimality of Other Options

The searcher will engage after receiving a confirming tip because that state generates 0 cost. Thus to determine the optimal policy we only need to consider the searcher engaging on or before the confirming tip. Proposition 2 lists two additional options: engage before receiving any tips or engage the cell pointed to by the first tip. We need to consider one final option: engage after the first tip but before the confirming tip. In this section we show this option is suboptimal and thus we only need to consider the three options listed in Proposition 2.

Let us define a related stopping problem where the searcher has received m tips about m different cells. The searcher has the option of waiting or engaging one of the m cells uniformly. We will use m as our state variable. Before defining the cost function we will derive some important probabilities. The searcher reaches state m (ignoring the possibility of attack) with probability $(1 - q)^m + mq(1 - q)^{m-1}$, which accounts for either m incorrect tips or 1 correct tip and $m - 1$ incorrect tips. This implies

$$\mathbf{P}[\text{first } m \text{ tips contain correct cell}] = \frac{mq}{(1 - q) + mq}, \qquad \text{(EC.83)}$$

which converges to 1 as m increases. If the searcher engages he will choose an incorrect cell if either (1) none of the m tips point to the correct cell, or (2) one of the m tips points to the correct cell, but the searcher does not engage that cell. The first situation occurs with

probability

$$\mathbf{P}[\text{first } m \text{ tips do not point to correct cell}] = \frac{1-q}{(1-q)+mq} \qquad \text{(EC.84)}$$

The probability the second situation occurs, given the first m tips contain the correct cell, is $(m-1)/m$. Putting everything together we have

$$\mathbf{P}[\text{searcher chooses incorrect cell}] = \frac{(1-q)+(m-1)q}{(1-q)+mq} = 1 - \frac{q}{(1-q)+mq}, \qquad \text{(EC.85)}$$

which monotonically approaches 1 as we increase m. Equation (EC.85) also represents the cost for the searcher to engage in state m as we normalize collateral damage to 1. If the searcher waits, then with the probability in (EC.84) the m tips do not contain the correct cell and thus the system will transition to state $m+1$ (provided an attack does not occur). If one of the m tips does contain the correct cell (probability in (EC.83)), then with probability q we receive a confirming tip and engage and with probability $(1-q)$ we transition to state $m+1$. This provides us with the probability of transitioning to state $m+1$

$$\mathbf{P}[\text{next tip not confirming}] = \frac{1-q}{(1-q)+mq} + \frac{mq}{(1-q)+mq}(1-q) = 1 - \frac{mq^2}{(1-q)+mq},$$

which monotonically decreases to $1-q$ as we increase m. We now present our cost function $\hat{C}(m)$, for $m > 0$, for this system:

$$\hat{C}(m) = \min\left(1 - \frac{q}{(1-q)+mq}, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\left(1 - \frac{mq^2}{(1-q)+mq}\right)\hat{C}(m+1)\right) \qquad \text{(EC.86)}$$

We will now show that if the searcher optimally engages in state $m > 1$, they he should engage in $m-1$ and thus the searcher should engage in state $m = 1$ (option 2). That is the searcher should engage the cell pointed to by the first tip, which is option 2 in Proposition 2. This will complete the proof that it is suboptimal to receive more than one tip and engage before receiving a confirming tip.

If the searcher engages for some $m > 1$, this implies

$$1 - \frac{q}{(1-q) + mq} \leq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\left(1 - \frac{mq^2}{(1-q)+mq}\right)\hat{C}(m+1) \qquad \text{(EC.87)}$$

The optimal cost in state $m + 1$, $\hat{C}(m + 1)$, must be less than the cost to engage in state $m + 1$, which leads to

$$1 - \frac{q}{(1-q)+mq} \leq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\left(1 - \frac{mq^2}{(1-q)+mq}\right)\left(1 - \frac{q}{(1-q)+(m+1)q}\right) \qquad \text{(EC.88)}$$

Working through the math simplifies the right-hand-side of (EC.88) significantly

$$1 - \frac{q}{(1-q)+mq} \leq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-q). \qquad \text{(EC.89)}$$

The right-hand-side of (EC.89) is independent of m. The left-hand-side of (EC.89) monotonically increases with m, which implies that condition (EC.89), and hence (EC.87), hold for $m - 1$. Consequently, the searcher should engage in state $m - 1$. Following this inductive logic implies the searcher should engage for state $m = 1$, which corresponds to option 2 in Proposition 2.

F.2.3 Optimal Decision for Different Parameters

The searcher optimally chooses the minimum-cost alternative among the three options. In this section, we show that the conditions presented in Proposition 2 generate the optimal decision.

When comparing options 1 and 2, the searcher prefers to engage immediately over obtaining one tip if

$$\alpha > 1 + \frac{q}{\rho} \qquad \text{(EC.90)}$$

When comparing options 1 and 3, the searcher prefers to engage immediately over waiting

for a repeat tip if

$$\alpha > \frac{1}{\left(1 - \left(\frac{q}{\rho+q}\right)^2\right)}$$

$$= \frac{(\rho+q)^2}{(\rho+q)^2 - q^2}$$

$$= 1 + \frac{q}{\rho}\frac{q}{\rho+2q} \qquad \text{(EC.91)}$$

The right-hand-side of condition (EC.90) exceeds the right-hand-side of condition (EC.91) by inspection. If the searcher prefers option 1 to option 2 (condition (EC.90)), he will also prefer option 1 to option 3 (condition (EC.91)). Thus, condition (EC.90) is a necessary and sufficient condition for the searcher to choose option 1.

When this condition does not hold, we compare option 2 and option 3. The searcher chooses to wait for the confirming tip if

$$\alpha \left(1 - \left(\frac{q}{\rho+q}\right)^2\right) < \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1-q)$$

$$\rightarrow \alpha \left(\frac{1}{1+\rho} - \left(\frac{q}{\rho+q}\right)^2\right) < \frac{1}{1+\rho}(1-q)$$

$$\rightarrow \alpha < \frac{1-q}{1 - (1+\rho)\left(\frac{q}{\rho+q}\right)^2}$$

$$\rightarrow \alpha < \frac{(1-q)(\rho+q)^2}{\rho^2 + 2\rho q - \rho q^2}$$

$$\rightarrow \alpha < 1 + \frac{q}{\rho}\frac{q - \rho^2 - \rho q - q^2}{\rho + 2q - q^2} \qquad \text{(EC.92)}$$

Clearly the right-hand-side of (EC.92) is smaller than the right-hand-side of (EC.90). It is also straightforward to derive that the right-hand-side of (EC.92) is less than the right-hand-side of (EC.91). If the searcher prefers option 3 to option 2 (condition (EC.92)), he will also prefer option 3 to option 1 (condition (EC.91) does not hold). Consequently, (EC.92) is a necessary and sufficient condition for the searcher to choose option 3.

F.3 Finite n Heuristic

We can use the $n = \infty$ case as a basis for a heuristic policy for the uniform state. In the heuristic the searcher has the same three options: engage before receiving tips, receive one tip and then engage the corresponding cell, or engage the first cell to receive two tips. We could use the costs associated with each option in Proposition 2 to determine which option the searcher chooses. However, we can compute explicitly the cost for each option in the finite n case, and thus we use those costs as the basis for the heuristic.

We add one more option: wait forever. In this situation the searcher essentially allows the target to execute the attack, which causes a cost of α. If we do not allow for this option, then the heuristic can generate relatively large costs for small α, ρ and q situations by restricting the searcher to choose only one of the original three options.

In Appendix F.3.1 we compute the finite n cost for the first three options and in Appendix F.3.2 we evaluate how well this heuristic performs in practice for finite n.

F.3.1 Costs

For the *large n heuristic*, the searcher should choose the minimum cost alternative among the following four options

1. Engage any cell: cost $1 - 1/n$

2. Obtain one tip and engage the corresponding cell: cost $\frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}(1 - q)$

3. Engage the first cell to receive 2 tips: cost appears in equation (EC.97).

4. Always wait: cost α

Before proceeding, we note that this heuristic generates a sufficient condition to continue. If the searcher chooses not to engage (i.e., option 2, 3, or 4), then the searcher should optimally wait in the uniform state. The details appear below for each of the first three options.

Option 1

If the searcher chooses a cell before receiving a tip, he will choose correctly with probability $1/n$. The result follows immediately.

Option 2

If the searcher obtains one tip, the state transitions such that $p_1 = q$. This holds for any n. Thus the cost for option 2 in the finite n case is the same as in the $n = \infty$ case. See Appendix F.2.1 for more details.

Option 3

This case is much more complicated. In the $n = \infty$ case, collateral damage cannot occur for option 3. However, in the finite n case, we have to account for the possibility that the informant points to the wrong cell twice, which would cause collateral damage if the searcher chooses option 3. Furthermore, the searcher will obtain at most $n + 1$ tips before a repeat occurs.

We define p_d as the probability the searcher successfully detects the target before the attack occurs, p_c as the probability collateral damage occurs, and p_a as the probability the target executes his attack before the searcher engages a cell. Using this notation, the cost associated with option 3 is $p_c + p_a \alpha$. We first derive p_d.

To compute p_d we follow similar steps to our approach in Appendix F.2.1. We first compute the probability that the searcher will point to cell 1 twice before providing a repeat tip of an incorrect cell and the second tip occurs on tip m, $2 \leq m \leq n + 1$. For this to occur, the mth tip points to cell 1 (prob q), and in the first $m - 1$ tips, $m - 2$ tips point to $m - 2$ distinct wrong cells and one points to cell 1 (probability q). This analysis differs from Appendix F.2.1 because of how we account for the incorrect tips. The probability the informant points to $m - 2$ different incorrect cells is

$$\mathbf{P}[m - 2 \text{ different incorrect cells}] = (1 - q)^{m-2} \prod_{j=1}^{m-2} \frac{n - j}{n - 1} \qquad \text{(EC.93)}$$

We now can calculate p_d using the same logic as in Appendix F.2.1.

$$p_d = \sum_{i=2}^{n+1} \frac{(i-1)q^2}{(1+\rho)^i}(1-q)^{i-2} \prod_{j=1}^{i-2} \frac{n-j}{n-1}. \tag{EC.94}$$

No nicer form to (EC.94) exists, but we can easily compute p_d.

We next calculate p_c. We first focus on cell 2 for concreteness and examine the probability the searcher will engage cell 2 using option 3. The final value of p_c is just the probability of collateral damage in cell 2 multiplied by $n-1$ to cover all possible incorrect cells. The expression is complicated so we present it below and then explain each component. The probability of collateral damage in cell 2 is.

$$\frac{p_c}{n-1} = \sum_{i=3}^{n+1}(i-1)\left(\frac{1-q}{n-1}\right)^2 (i-2)q\frac{1}{(1+\rho)^i}(1-q)^{i-3}\prod_{j=2}^{i-2}\frac{n-j}{n-1}$$
$$+ \sum_{i=2}^{n}(i-1)\left(\frac{1-q}{n-1}\right)^2 \frac{1}{(1+\rho)^i}(1-q)^{i-2}\prod_{j=2}^{i-1}\frac{n-j}{n-1} \tag{EC.95}$$

The i in the summation corresponds to the second tip to cell 2 occurring on tip i. We have two slightly different terms because we must account for whether one of the other $i-2$ tips points to cell 1. The first line in (EC.95) represents the situation where one of the other $i-2$ tips points to cell 1. The searcher points twice to cell 2 (probability $\left(\frac{1-q}{n-1}\right)^2$) and the first tip to cell 2 can occur in any of the first $i-1$ tips. The tip to cell 1 can occur in any of the remaining first $i-2$ tips (the $(i-2)q$ term). This leaves $(i-3)$ tips that must point to $(i-3)$ different incorrect cells other than cell 2. The product expression is a slight modification of (EC.93). For the second line, the searcher does not point to cell 1. Thus the second line does not contain the $(i-2)q$ term and the searcher points to $(i-2)$ distinct incorrect cells other than cell 2. The first line starts the sum at $i=3$ because for $i=2$ a tip cannot point to cell 1. Finally the first line sums to $n+1$, whereas the second line only sums to n. When the second tip to cell 2 occurs on tip $i = n+1$, the searcher must point to cell 1 in one of the first n tips.

Finally

$$p_a = 1 - p_d - p_c. \tag{EC.96}$$

Combining the values from equations (EC.94)–(EC.96) produces the final cost for option 3:

$$\text{option 3 cost} = p_c + p_a \alpha \tag{EC.97}$$

F.3.2 Heuristic Performance

In Figure 8 we present the performance of the heuristic for various values of n, q, α, and ρ. The y-axis corresponds to $\frac{\text{heuristic cost}}{\text{optimal cost}} - 1$. In Appendix J.1 we examine the performance of several heuristics. Overall this large n heuristic performs very well.

In all cases, the heuristic is optimal for larger q. In these cases the searcher obtains one tip and engages the corresponding cell.

For small q the heuristic also performs well. For small α (Figures 8a and 8b), the heuristic chooses option 4, which is nearly optimal. In these cases the tips are not informative, but the searcher optimally receives many of them rather than engaging because collateral damage is significant. For larger α (Figures 8c and 8d), the heuristic never chooses option 4. In these situations, the heuristic chooses either option 1 or 2, which is usually the optimal decision. The searcher may improve his situation slightly by obtaining one tip from a poor informant, but receiving more than one tip from a poor informant does not provide the searcher with any benefit.

The interesting behavior in Figure 8 occurs for moderate values of q. The non-monotonic relationship is particularly strong for small values of ρ (Figures 8a and 8c). For small ρ, the searcher can afford to receive many tips. With small α (Figure 8a), the risk of receiving many tips is particularly small, but our heuristic puts a cap on the number of tips the searcher can receive. For moderate q there is great marginal value in receiving more than one tip, and option 3 does not necessarily represent the optimal policy well in these moderate q and

small α and ρ situations. In the worst case scenario in Figure 8 the heuristic cost exceeds the optimal cost by over 25%.

One surprising aspect of this analysis is that the the performance for smaller n scenarios is not noticeably worse than the large n scenarios.

(a) $\alpha = 0.5$ and $\rho = 0.1$

(b) $\alpha = 0.5$ and $\rho = 1$

(c) $\alpha = 1$ and $\rho = 0.1$
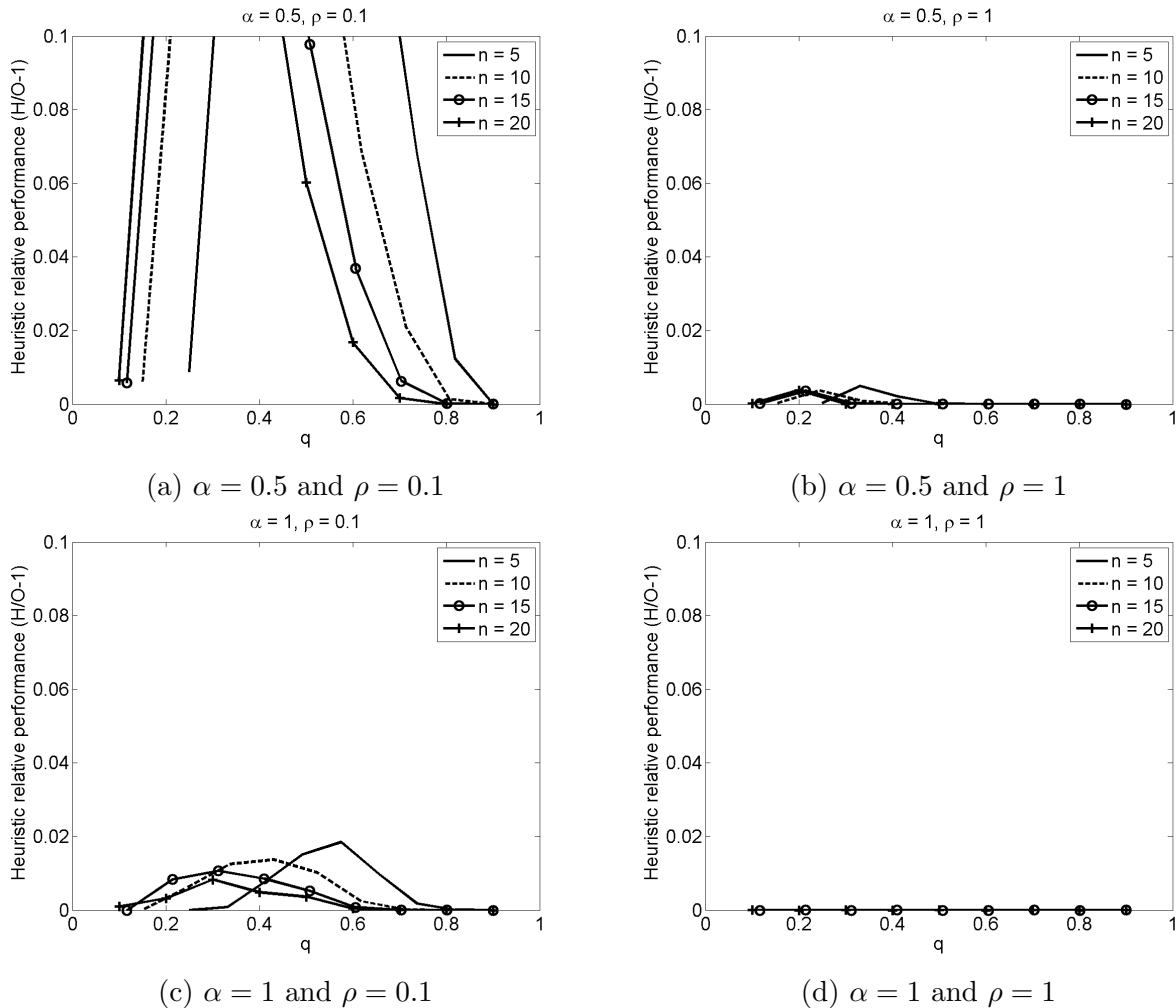
(d) $\alpha = 1$ and $\rho = 1$

Figure 8: Heuristic performance, measured in terms of how much the ratio of heuristic cost to optimal cost exceeds 1, as q varies for combinations of $\alpha \in \{0.5, 1\}$, $\rho \in \{0.1, 1\}$, and $n \in \{5, 10, 15, 20\}$.

G Bounds on the Cost Function in (11)

To derive upper and lower bounds on $C(p)$, we truncate the problem following $k \geq 0$ additional tips. For the lower bound, we assume a best-case scenario in which the searcher

correctly identifies the target's cell after the kth tip and incurs zero cost. Since $C(p) \geq 0$, we obtain the lower bound for $C(p)$ as the solution of a finite stage dynamic program:

$$L(p, k) = \begin{cases} \min\left(1 - p_1, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)L(p^{(+i)}, k-1)\right) & \text{if } k > 0 \\ 0 & \text{if } k = 0. \end{cases} \quad \text{(EC.98)}$$

To obtain an upper bound, we assume that the searcher engages after receiving k more tips:

$$U(p, k) = \begin{cases} \min\left(1 - p_1, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)U(p^{(+i)}, k-1)\right) & \text{if } k > 0 \\ 1 - p_1 & \text{if } k = 0. \end{cases} \quad \text{(EC.99)}$$

The only difference between $L(p, k)$ and $U(p, k)$ occurs in the terminal condition when $k = 0$. In Appendices G.1 – G.3 we prove the following propositions.

Proposition EC.5. *For $C(p)$ defined by (11), $L(p, k)$ defined by (EC.98), and $U(p, k)$ defined by (EC.99), we have the following relationship: $L(p, k) \leq C(p) \leq U(p, k)$ for all states p and $k \geq 0$.*

The following monotonicity result states the bounds become tighter as we increase k.

Proposition EC.6. *For $L(p, k)$ defined by (EC.98) and $U(p, k)$ defined by (EC.99), we have the following relationships: $L(p, k) \geq L(p, k-1)$ and $U(p, k) \leq U(p, k-1)$ for all states p and $k \geq 0$.*

Finally, the next proposition states the expected result that $L(p, k)$ and $U(p, k)$ both converge to $C(p)$.

Proposition EC.7. *For $C(p)$ defined by (11), $L(p, k)$ defined by (EC.98), and $U(p, k)$ defined by (EC.99), we have $L(p, k) \to C(p)$ and $U(p, k) \to C(p)$ as $k \to \infty$ for all states p.*

Specifically,

$$C(p) - L(p, k) \leq (1 + \alpha) \left(\frac{1}{1 + \rho} \right)^k \tag{EC.100}$$

$$U(p, k) - C(p) \leq (1 + \alpha) \left(\frac{1}{1 + \rho} \right)^k. \tag{EC.101}$$

Proposition EC.7 allows us to theoretically approximate $C(p)$ to any desired precision and determine whether the searcher should engage or wait in state p.

G.1 Proposition EC.5: Upper and Lower Bounds on the Cost Function

We first show that $L(p, k) \leq C(p)$ via an inductive argument. Obviously $L(p, 0) = 0 \leq C(p)$ for all p and thus we have proven the base case. We next turn to the inductive step and assume $L(p, k - 1) \leq C(p)$ for all p and $k > 1$. The engage cost for $L(p, k)$ in (EC.98) equals the engage cost for $C(p)$ in (11) for $k > 1$. By our induction assumption

$$\sum_{i=1}^{n} r_i(p) L(p^{(+i)}, k - 1) \leq \sum_{i=1}^{n} r_i(p) C(p^{(+i)}),$$

and thus the cost to wait for $L(p, k)$ in (EC.98) is less than the cost to wait for $C(p)$ in (11). Since both the engage and wait cases for $L(p, k)$ are less than or equal to their corresponding cases for $C(p)$, we have proven that $L(p, k) \leq C(p)$.

The proof for $U(p, k) \geq C(p)$ follows in nearly the same fashion. By inspection of (EC.99) we see that

$$U(p, 0) = 1 - p_1 \geq \min \left(1 - p_1, \frac{\rho}{1 + \rho} \alpha + \frac{1}{1 + \rho} \sum_{i=1}^{n} r_i(p) C(p^{(+i)}) \right) = C(p).$$

We have proven the base case, and the induction argument proceeds in the same fashion as for $L(p, k)$. ■

G.2 Proposition EC.6: Monotonicity of $L(p, k)$ and $U(p, k)$

We first show that $L(p, k) \geq L(p, k-1)$ via an inductive argument. For the base case, $k = 1$, $L(p, 1) \geq L(p, 0) = 0$ follows by inspection of equation (EC.98). We now proceed to the induction step and assume $L(p, k-1) \geq L(p, k-2)$ for $k > 1$. The engage costs for $L(p, k)$ and $L(p, k-1)$ both equal $(1 - p_1)$. To compare the wait cases, we utilize the induction assumption $L(p, k-1) \geq L(p, k-2)$ to derive

$$\sum_{i=1}^{n} r_i(p)L(p^{(+i)}, k-1) \geq \sum_{i=1}^{n} r_i(p)L(p^{(+i)}, k-2). \tag{EC.102}$$

Utilizing (EC.102) and inspecting (EC.98), we observe that the cost to wait for $L(p, k)$ exceeds the cost to wait for $L(p, k-1)$. Combining the results for the engage and wait options, we have proven $L(p, k) \geq L(p, k-1)$.

The proof for $U(p, k) \leq U(p, k-1)$ follows in nearly the same fashion. By inspection of (EC.99) we see that the base case holds for $k = 1$:

$$U(p, 1) = \min \left(1 - p_1, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho} \sum_{i=1}^{n} r_i(p)U(p^{(+i)}, 0) \right) \leq 1 - p_1 = U(p, 0).$$

The induction argument proceeds in the same fashion as for $L(p, k)$. ∎

G.3 Proposition EC.7: Convergence of $L(p, k) \to C(p)$ and $U(p, k) \to$ $C(p)$

For most of the paper we examine the searcher's problem through the lens of $C(p)$ and its associated dynamic programming formulation. See equations (8) and (11). However, in this section we examine the optimal stopping problem from first principles. We adapt much of the notation and approach from Chapter 1 of Ferguson (2004).

Preliminaries

First assume that cell 1 contains the target. For any state p, define X_i as the cell number specified by the ith tip in the future. Thus $\mathbf{P}[X_i = 1] = q$ and $\mathbf{P}[X_i = j] = \frac{1-q}{n-1}$ for $j = 2, 3, \ldots, n$. Next define a sequence of cost functions $y_i(\cdot)$ that represents the total expected cost incurred by the searcher moving forward if he decides to collect exactly i additional tips and then engage. We denote the specific value taken by the random variable X_i with a lowercase representation: $X_1 = x_1, X_2 = x_2, X_3 = x_3 \ldots$. Before defining the explicit form of y_i, we denote $\hat{p}(p, x_1, x_2, x_3, \ldots, x_i)$ as the updated posterior probability vector after receiving i additional tips taking on values $x_1, x_2, x_3, \ldots, x_i$ with prior p. The expected cost y_i depends on both the current state p and the values of the i future tips.

$$y_i(p, x_1, x_2, \ldots, x_i) = \frac{\rho}{1+\rho}\alpha \sum_{k=0}^{i-1} \left(\frac{1}{1+\rho}\right)^k + \left(\frac{1}{1+\rho}\right)^i \left(1 - \max_j \hat{p}_j(p, x_1, x_2, x_3, \ldots, x_i)\right)$$

$$= \alpha + \left(\frac{1}{1+\rho}\right)^i \left(\left(1 - \max_j \hat{p}_j(p, x_1, x_2, x_3, \ldots, x_i)\right) - \alpha\right). \quad \text{(EC.103)}$$

As expected the cost $y_i(p, x_1, x_2, \ldots, x_i)$ in (EC.103) approaches α for large i because eventually the plot matures and the target attacks.

One can define a stopping rule as a sequence of stopping functions $\phi_i(\cdot)$. After receiving i additional tips, $X_1 = x_1, X_2 = x_2, \ldots, X_i = x_i$, the searcher engages if the stopping function $\phi_i(p, x_1, x_2, \ldots, x_i) = 1$ and waits if $\phi_i(p, x_1, x_2, \ldots, x_i) = 0$. Technically the stopping function could output a number between zero and one that represents the probability the searcher engages given the sequence of tips. However, we will limit ourselves to deterministic stopping rules. A stopping rule ϕ comprises a collection of stopping functions: $\phi = \{\phi_0(p), \phi_1(p, x_1), \phi_2(p, x_1, x_2) \ldots\}$. The stopping rule specifies when the searcher should engage for every possible stream of future tips. If the searcher decides to engage whenever $(1 - \max_j \hat{p}_j(p, x_1, x_2, x_3, \ldots, x_i)) > 0.5$, we can map this condition to stopping functions ϕ_i and hence a stopping rule ϕ.

For a given stopping rule ϕ and random stream of future tips $X_1, X_2, X_3, \ldots$, we have

a random stopping time N. It is possible that $N = \infty$ if ϕ never specifies to engage for the given sequence of tips. We can theoretically solve our optimal stopping problem by determining the optimal stopping rule. We now rewrite $C(p)$ in terms of the stream of future tips X_i, the cost function y_i, the stopping rule ϕ, and the stopping time N:

$$C(p) = \min_{\phi} \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)]. \tag{EC.104}$$

Of course, the stopping time N is a function of the stopping rule ϕ. Refer to Chapter 1 of Ferguson (2004) for more details on this approach.

Using this framework and equation (EC.104) will not aid us in determining whether the searcher should engage or wait in state p. However, it provides the path to showing that $L(p, k) \to C(p)$ and $U(p, k) \to C(p)$.

Proof that $L(p, k) \to C(p)$

First by Proposition EC.6 we have $C(p) - L(p, k) \geq 0$. Next denote ϕ^C as an optimal stopping rule for $C(p)$ with corresponding stopping time N_C and ϕ^L as an optimal stopping rule for $L(p, k)$ with corresponding stopping time N_L. By definition $C(p) = \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C})]$. Note further that by construction $N_L \leq k$. We next modify the cost function in (EC.103) to define the cost function $y_i^L(p, x_1, x_2, \ldots, x_i)$ associated with $L(p, k)$

$$y_i^L(p, x_1, x_2, \ldots, x_i) = \begin{cases} y_i(p, x_1, x_2, \ldots, x_i) & \text{if } i < k \\[2mm] \alpha \left(1 - \left(\frac{1}{1+\rho} \right)^k \right) & \text{if } i = k. \end{cases} \tag{EC.105}$$

The special condition for $i = k$ in (EC.105) corresponds to the searcher receiving k additional tips and then obtaining the base case cost of 0 in equation (EC.98). We next create a new stopping rule ϕ by merging stopping rules ϕ^C and ϕ^L. The searcher uses ϕ^L until the kth

EC49

tip, and then the searcher uses $\boldsymbol{\phi}^C$ after the kth tip:

$$\phi_i(p, x_1, x_2, \ldots, x_i) = \begin{cases} \phi_i^L(p, x_1, x_2, \ldots, x_i) & \text{if } i < k \\ \phi_i^C(p, x_1, x_2, \ldots, x_i) & \text{if } i \geq k. \end{cases} \tag{EC.106}$$

We denote as N the stopping time associated with the stopping rule ϕ defined by (EC.106). Because N_C corresponds to an optimal stopping time for $C(p)$, we must have

$$C(p) = \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C})] \leq \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)].$$

Putting the pieces together yields

$$0 \leq C(p) - L(p, k)$$

$$= \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C})] - \mathbf{E}[y_{N_L}^L(p, X_1, X_2, \ldots, X_{N_L})]$$

$$\leq \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)] - \mathbf{E}[y_{N_L}^L(p, X_1, X_2, \ldots, X_{N_L})]$$

$$= \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)I(N < k)] + \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)I(N \geq k)]$$

$$- \mathbf{E}[y_{N_L}(p, X_1, X_2, \ldots, X_{N_L})I(N_L < k)] - \mathbf{E}[y_k^L(p, X_1, X_2, \ldots, X_k)I(N_L = k)]. \tag{EC.107}$$

By construction $N = i$ if and only if $N_L = i$ for $i < k$ (see equation (EC.106)). Thus the $I(N < k)$ term and the $I(N_L < k)$ cancel yielding

$$0 \leq C(p) - L(p, k) \leq \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)I(N \geq k)] - \mathbf{E}[y_k^L(p, X_1, X_2, \ldots, X_k)I(N_L = k)]. \tag{EC.108}$$

However, substituting (EC.103) and (EC.105) into (EC.108) produces

$$0 \le C(p) - L(p, k)$$

$$\le \mathbf{E}\left[\left(\frac{1}{1+\rho}\right)^N \left((1 - \max_j \hat{p}_j(p, X_1, X_2, \ldots, X_N)) - \alpha\right) I(N \ge k)\right] + \mathbf{E}\left[\alpha\left(\frac{1}{1+\rho}\right)^k I(N_L = k)\right]$$

$$\le \mathbf{E}\left[\left(\frac{1}{1+\rho}\right)^N I(N \ge k)\right] + \alpha\left(\frac{1}{1+\rho}\right)^k$$

$$\le \left(\frac{1}{1+\rho}\right)^k \mathbf{P}\left[N \ge k\right] + \alpha\left(\frac{1}{1+\rho}\right)^k$$

$$\le (1+\alpha)\left(\frac{1}{1+\rho}\right)^k \to 0.$$

Thus $L(p, k) \to C(p)$. ∎

Proof that $U(p, k) \to C(p)$

We prove $U(p, k)$ in a similar fashion. Denote $\boldsymbol{\phi}^U$ as an optimal stopping rule for $U(p, k)$ with corresponding stopping time N_U. By construction $N_U \le k$. Unlike with $L(p, k)$ we do not need to define a separate cost function $y_i^U(p, x_1, x_2, \ldots, x_i)$ as in (EC.105). We only need to consider $y_i(p, x_1, x_2, \ldots, x_i)$ in equation (EC.103). As with $L(p, k)$ we define a modified stopping time. In this case we merely truncate N_C

$$N = \min(N_C, k). \tag{EC.109}$$

By construction $U(p, k) = \mathbf{E}[y_{N_U}(p, X_1, X_2, \ldots, X_{N_U})] \le \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)]$. Using this fact, we can derive a sequence on inequalities as with $L(p, k)$

$$0 \le U(p, k) - C(p)$$

$$= \mathbf{E}[y_{N_U}(p, X_1, X_2, \ldots, X_{N_U})] - \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C})]$$

$$\le \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)] - \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C})]$$

$$= \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)I(N < k)] + \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)I(N = k)]$$

$$- \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C}))I(N_C < k)] - \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C})I(N_C \ge k)]. \quad (\text{EC.110})$$

The $I(N < k)$ and $I(N_C < k)$ terms in (EC.110) cancel due to the definition of N in (EC.109). This produces

$$0 \le U(p, k) - C(p)$$

$$\le \mathbf{E}[y_N(p, X_1, X_2, \ldots, X_N)I(N = k)] - \mathbf{E}[y_{N_C}(p, X_1, X_2, \ldots, X_{N_C})I(N_C \ge k)]$$

$$= \left(\frac{1}{1+\rho}\right)^k \mathbf{E}\left[\left((1 - \max_j \hat{p}_j(p, X_1, X_2, \ldots, X_k)) - \alpha\right) I(N = k)\right]$$

$$- \mathbf{E}\left[\left(\frac{1}{1+\rho}\right)^{N_C} \left((1 - \max_j \hat{p}_j(p, X_1, X_2, \ldots, X_{N_C})) - \alpha\right) I(N_C \ge k)\right]$$

$$\le \left(\frac{1}{1+\rho}\right)^k + \alpha \mathbf{E}\left[\left(\frac{1}{1+\rho}\right)^{N_C} I(N_C \ge k)\right]$$

$$\le \left(\frac{1}{1+\rho}\right)^k + \alpha \left(\frac{1}{1+\rho}\right)^k \mathbf{P}[N_C \ge k]$$

$$= (1 + \alpha) \left(\frac{1}{1+\rho}\right)^k \to 0.$$

Thus $U(p, k) \to C(p)$ and $U(p, k) - L(p, k) \to 0$. $\blacksquare$

H Conditions to Engage

A lower bound on the expected cost to wait in equation (11) is obtained by substituting $L(p^{(+i)}, k-1)$ from $(EC.98)$ for $C(p^{(+i)})$ in equation (11) for any value $k > 0$. This produces the following condition: If

$$1 - p_1 \leq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)L(p^{(+i)}, k-1) \tag{EC.111}$$

for some $k > 0$, then the searcher should engage cell 1.

Note that by Proposition EC.6, the right-hand side of (EC.111) is monotonically non-decreasing in k. Therefore, if condition (EC.111) holds for some $k = K$, it will also hold for any $k > K$. As we would expect, larger k values produce tighter conditions to engage.

In Appendix H.1 we analyze the $k = 2$ situation. In Appendix H.2 we consider a condition to engage unrelated to $L(p, k)$ and condition (EC.111) by examining the "best case" future stream of tips.

H.1 Condition (19)

To derive condition (19) we substitute $k = 2$ into condition (EC.111)

$$1 - p_1 \leq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)L(p^{(+i)}, 1). \tag{EC.112}$$

Furthermore by inspection of (EC.98),

$$L(p, 1) = \min\left(1 - p_1, \frac{\rho}{1+\rho}\alpha\right). \tag{EC.113}$$

Substituting (EC.113) into (EC.112) produces

$$1 - p_1 \leq \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)\min\left(1 - \max_j p_j^{(+i)}, \frac{\rho}{1+\rho}\alpha\right). \tag{EC.114}$$

Going through the algebra transforms (EC.114) into (19).

H.2 Best Case Future Stream of Tips: Informant Points only to Cell 1

We derive another condition to engage by constructing, once again, a lower bound on the cost-to-wait term of $C(p)$ in (11). If the searcher engages cell 1 after exactly k additional tips, then clearly the best case scenario – the smallest possible expected cost – occurs if all k tips point at cell 1, the cell with the highest probability. Using this approach, we derive that the searcher should engage cell 1 if

$$p_1 \geq 1 - \left(\alpha + \min_{k \in \mathbb{Z}^+} \left(\frac{1}{1+\rho} \right)^k \left(1 - \alpha - \frac{\gamma^k p_1}{\gamma^k p_1 + (1 - p_1)} \right) \right). \tag{EC.115}$$

where γ is defined by equation (3). The derivation for condition (EC.115) appears in Appendix H.2.1

For a fixed value of k, neither condition (EC.111) nor condition (EC.115) dominates the other. We derive condition (EC.111) by truncating the problem after a finite number of periods, whereas we perform no such truncation in the derivation of condition (EC.115). However, condition (EC.115) assumes the best case future. For any particular state, if we increase k in (EC.111), then eventually (EC.111) will dominate condition (EC.115). For most realistic values of n (~ 100), we should be able to increase k enough such that we only need to consider (EC.111). However, if the number of cells is large ($n \sim 1000$), computing the right-hand side of (EC.111) may be computationally intractable, even for small k. As an example consider $n = 100$ and $k = 5$. Solving for $L(p, k)$ in (EC.98) requires backward induction. To compute the size of the state space, we must determine the number of ways to place k indistinguishable balls in n distinguishable bins. Using standard combinatorial techniques, the size equals $9 \cdot 10^7$ for $n = 100$ and $k = 5$. However, for $n = 1000$ and $k = 5$ the state space balloons to size $8 \cdot 10^{12}$, which makes computations very difficult on standard

computers. Computing the the right-hand side of (EC.115) poses no such problems. In practice, any modern computer can perform the calculation in a trivial amount of time for any state p.

H.2.1 Derivation of Condition EC.115

If the searcher receives k additional confirming tips that cell 1 contains the target, then the state transitions from p to $\tilde{p}$

$$
\tilde{p}_j = \begin{cases} \frac{\gamma^k p_1}{\gamma^k p_1 + (1-p_1)} & \text{if } j = 1 \\[2em] \frac{p_j}{\gamma^k p_1 + (1-p_1)} & \text{if } j \neq 1. \end{cases}
$$

Consequently

$$
\tilde{p}_1 = \frac{\gamma^k p_1}{\gamma^k p_1 + (1 - p_1)}
$$

If the searcher engages cell 1 before receiving any tips, the searcher incurs an expected cost of $(1 - p_1)$. If the searcher instead receives k additional tips and then chooses a cell, the best case scenario occurs when the informant provides k confirming tips that cell 1 contains the target. In this best case, the searcher incurs the following expected cost after k additional tips

$$
\begin{aligned}
BC_k(p) &= \frac{\rho}{1+\rho}\alpha \sum_{i=0}^{k-1} \left(\frac{1}{1+\rho}\right)^i + \left(\frac{1}{1+\rho}\right)^k (1 - \tilde{p}_1) \\
&= \frac{\rho}{1+\rho}\alpha \times \frac{1 - \left(\frac{1}{1+\rho}\right)^k}{\frac{\rho}{1+\rho}} + \left(\frac{1}{1+\rho}\right)^k (1 - \tilde{p}_1) \\
&= \alpha + \left(\frac{1}{1+\rho}\right)^k \left(-\alpha + \left(1 - \frac{\gamma^k p_1}{\gamma^k p_1 + (1 - p_1)}\right)\right).
\end{aligned}
$$

Taking the minimum over all $k > 0$ produces a lower bound on the expected cost to wait in

state p. Thus the searcher should engage if $(1 - p_1) \leq \min_k BC_k(p)$, which corresponds to condition (EC.115). In practice, we only need to calculate $BC_k(p)$ for a finite number of k values to determine $\min_k BC_k(p)$. To see why, note that

$$f(k) = \left(-\alpha + \left(1 - \frac{\gamma^k p_1}{\gamma^k p_1 + (1 - p_1)} \right) \right)$$

is a decreasing function of k and $\lim_{k \to \infty} f(k) = -\alpha < 0$. Thus, we can find a $\hat{K}$ such that $f(k) < 0$ for all $k \geq \hat{K}$. Furthermore, since $f(k)$ is bounded, $\lim_{k \to \infty} \left(\frac{1}{1+\rho} \right)^k f(k) = 0$. We can then find a $K^* \geq \hat{K}$ such that $\left(\frac{1}{1+\rho} \right)^k f(k) \geq \left(\frac{1}{1+\rho} \right)^{\hat{K}} f(\hat{K})$ for all $k \geq K^*$. To compute $\min_{k \in \mathbb{Z}^+} BC_k(p)$, we only need to calculate $\min_{k \in [1, K^*]} BC_k(p)$, and we can determine whether condition (EC.115) holds in practice in a straightforward manner.

I Conditions to Wait

Using a similar approach as in Appendix H we use $U(p, k)$ from (EC.99) to derive sufficient conditions to wait for more tips. Specifically, if

$$1 - p_1 > \frac{\rho}{1 + \rho} \alpha + \frac{1}{1 + \rho} \sum_{i=1}^{n} r_i(p) U(p^{(+i)}, k - 1) \tag{EC.116}$$

for some $k > 0$, then the searcher should wait.

Note that besides the obvious reverse in the direction of the two inequalities, the only difference between (EC.116) and (EC.111) is that $U(p^{(+i)}, k - 1)$ appears in the right-hand side of (EC.116) instead of $L(p^{(+i)}, k - 1)$. By Proposition EC.7, the right-hand sides of both (EC.116) and (EC.111) converge to the same value as $k \to \infty$.

In Appendix I.1 we derive condition 21 for the myopic policy.

I.1 Condition (21): Myopic Policy

To derive condition (21) we substitute $k = 1$ into the complement of condition (EC.116)

$$1 - p_1 \leq \frac{\rho}{1 + \rho} \alpha + \frac{1}{1 + \rho} \sum_{i=1}^{n} r_i(p) U(p^{(+i)}, 0). \tag{EC.117}$$

Substituting $U(p, 0) = (1 - \max_j p_j)$ into (EC.117) yields

$$1 - p_1 \leq \frac{\rho}{1 + \rho} \alpha + \frac{1}{1 + \rho} \sum_{i=1}^{n} r_i(p)(1 - \max_j p_j^{(+i)}).$$

Rearranging produces

$$p_1 \geq \frac{\rho}{1 + \rho} (1 - \alpha) + \frac{1}{1 + \rho} \sum_{i=1}^{n} r_i(p) \max_j p_j^{(+i)}. \tag{EC.118}$$

Using (5), (3), and (7), we rewrite $p_j^{(+i)}$ in terms of $r_i(p)$

$$p_j^{(+i)} = \begin{cases} \frac{q p_i}{r_i(p)} & \text{if } j = i \\[2ex] \frac{\frac{1-q}{n-1} p_j}{r_i(p)} & \text{if } j \neq i. \end{cases}$$

Next define $\tilde{p}^{(+i)}$ by multiplying the vector $p^{(+i)}$ by the scalar $r_i(p)$

$$\tilde{p}_j^{(+i)} = r_i(p) \cdot p_j^{(+i)} = \begin{cases} q p_i & \text{if } j = i \\[2ex] \frac{1-q}{n-1} p_j & \text{if } j \neq i. \end{cases} \tag{EC.119}$$

We can then rewrite condition (EC.118) in terms of $\tilde{p}^{(+i)}$:

$$p_1 \geq \frac{\rho}{1 + \rho} (1 - \alpha) + \frac{1}{1 + \rho} \sum_{i=1}^{n} \max_j \tilde{p}_j^{(+i)}. \tag{EC.120}$$

The maximum of $\tilde{p}^{(+i)}$ occurs at either cell i or cell 1:

$$\max_j \tilde{p}_j^{(+i)} = \max\left(qp_i, \frac{1-q}{n-1}p_1\right). \tag{EC.121}$$

Substituting (EC.121) into (EC.120) produces (21)

$$p_1 \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho}\sum_{i=1}^{n}\max\left(qp_i, \frac{1-q}{n-1}p_1\right). \tag{EC.122}$$

If $p_1 \leq \gamma p_i$ for all i, then the max term in (EC.122) returns qp_i for all i and the sum evaluates to q. This leads to the first case in (22). If $p_1 \geq \gamma p_i$ for all $i > 1$, then we can simplify the sum in the right-hand side of (EC.122)

$$\sum_{i=1}^{n}\max\left(qp_i, \frac{1-q}{n-1}p_1\right) = qp_1 + \sum_{2=1}^{n}\frac{1-q}{n-1}p_1 = p_1. \tag{EC.123}$$

Substituting (EC.123) into (EC.122) and simplifying produces the second case in (22).

J   Heuristic Policies

We first present the performance of several heuristic policies across many different scenarios in Appendix J.1. This analysis reveals a fairly strong performance of the myopic policy. In Appendix J.2 we examine how well the myopic policy provides the searcher with an approximation for the benefits of waiting. In this section, we do include the large n heuristic discussed in Section 5 and Appendix F.3. See Appendix F.3.2 for additional analysis on the performance of that heuristic

J.1   Performance of Heuristic Policies

We examine how much the cost increases by using a heuristic rather than the optimal policy. We initially focus on the $n = 3$ situation and vary $0.35 \leq q \leq 0.95$, $0.5 \leq \alpha \leq 1.5$,

$0.1 \leq \rho \leq 1$ over the entire state space for p. This creates 84000 scenarios. We summarize our results in Table 3. Each row corresponds to a heuristic and we present 3 columns to measure heuristic performance: fraction of scenarios where the heuristic recommends the optimal action, maximum relative difference between the heuristic cost and optimal cost, and mean relative difference between heuristic cost and optimal cost. We define the relative difference presented in the tables as $\frac{\text{heuristic cost}}{\text{optimal cost}} - 1$. For reference the average optimal cost over all scenarios is 0.34. If the heuristic and optimal both recommend to wait, they will not necessarily generate the same cost because the heuristic may deviate from optimal for future states. Also recall the heuristic is re-evaluated at each state. For example if the searcher uses the myopic policy and it recommends to wait, the searcher does not have to engage in the next state. After receiving the next tip, the searcher performs the myopic comparison to determine whether to engage or wait in the new state.

Heuristic	Fraction Of Scenarios Heuristic Recommends Optimal Action	Maximum Relative Cost Difference	Average Relative Cost Difference
eng(1-tip)	0.69	6.99	0.28
eng(2-tips)	0.86	3.38	0.10
eng(2-cell policy)	0.85	1.77	0.07
wait(myopic)	0.97	1.02	0.01

Table 3: Heuristic performance over 84000 scenarios for $n = 3$ for parameters in $0.35 \leq q \leq 0.95$, $0.5 \leq \alpha \leq 1.5$, and $0.1 \leq \rho \leq 1$

The first three heuristics perform poorly when they wait in situations they should engage. For *eng(1-tip)* and *eng(2-tips)*, situations with small q, large α, small ρ and reasonably large p_1 lead to large cost differentials. The searcher should engage, but these heuristics recommend to wait in the current state and many future states, which yields a cost near α. The *eng(2-cell policy)* underperforms for small q, large α and ρ, and small p_1. This heuristic will only engage for $p_1 > 0.5$, which can cause the searcher to wait in situations where there is a high

probability of a significant attack. The *wait(myopic)* policy is the clear winner, providing on average nearly optimal results. The worst case for the myopic policy appears in Figure 6a: small ρ and α and moderate q and p_1. There is not a benefit to one additional tip, but reasonable cost reduction can occur through several additional tips. Even though the relative difference yields nearly double the optimal cost in the worst case in Table 3, the actual cost is still small using the myopic policy (optimal cost = 0.13 vs. heuristic cost of 0.27).

We next perform a similar analysis for larger n. We consider 1875 scenarios with $5 \leq n \leq 30$, $0.5 \leq \alpha \leq 1.5$, $0.1 \leq \rho \leq 1$, and vary q from just above $1/n$ to 0.9. We only examine the uniform state situations so we can include the large n heuristic described in Section 5 and Appendix F.3 (denoted here as *cont(largeN)* as it represents a sufficient condition to wait). The average optimal cost over the 1875 scenarios is 0.59. The large n heuristic performs very well; there is no noticeable degradation in performance even for $n = 5$. The large n heuristic performs poorly in situations with moderate q and small α and ρ (see also Figure 8a in Appendix F.3.2). In these situations the searcher benefits from many tips to uncover the target's location, but the heuristic limits the searcher to at most two tips in a cell. We discuss the performance of the *cont(largeN)* heuristic in more detail in Appendix F.3.2. Unfortunately, this heuristic only applies for the uniform prior.

Heuristic	Fraction Of Scenarios Heuristic Recommends Optimal Action	Maximum Relative Cost Difference	Average Relative Cost Difference
eng(1-tip)	0.88	0.87	0.10
eng(2-tips)	0.95	0.87	0.04
eng(2-cell policy)	0.88	0.87	0.09
wait(myopic)	1	1.60	0.07
wait(largeN)	1	0.29	0.01

Table 4: Heuristic performance for the uniform state over 1875 scenarios for parameters in $5 \leq n \leq 30$, $0.5 \leq \alpha \leq 1.5$, and $0.1 \leq \rho \leq 1$

The myopic results in Table 4 are much worse than the $n = 3$ results in Table 3. A deeper analysis of the numbers reveals that the myopic policy performs quite well for most parameter combinations, and we can characterize the situations where the myopic policy will perform poorly. If we focus on the $\rho \geq 0.3$ scenarios In Table 4, the myopic performance is within 1% of optimal on average. However, for the $\rho = 0.1$ scenarios, the average performance is nearly 30% greater than optimal. For murky states (e.g., the uniform state) and small ρ, there are tremendous benfits to waiting for many additional tips to significantly improve the situational awareness. The myopic policy, by definition, does not account for the additional benefits after the next tip. We also examine scenarios for $n > 3$ and non-uniform states. The main conclusion is that the myopic policy produces nearly optimal results over a large range of parameter combinations, except for small ρ and murky states.

J.2 Using the Myopic Policy to Approximate the Benefit of Waiting

If the cost to wait is only slightly smaller than the cost to engage, then in practice the searcher may choose to engage because of uncertainties associated with the model parameters or other frictions we do not account for in the model. Computing the cost to engage is trivial, but computing the optimal cost to wait analytically is difficult. Thus, it would be useful to have a closed form metric to provide the searcher with a rough idea about when waiting provides a significant benefit over engaging. One logical starting point is moving from p to $p^{(+1)}$ (see equation (5)) as this captures the best case enhancement in the searcher's situational awareness. If $p_1^{(+1)}$ is much greater than p_1, this suggests the searcher will likely see a significant cost reduction if he waits. While this has merit, it ignores the impact of ρ and α. Even if the difference between $p_1^{(+1)}$ and p_1 is small, the searcher may still benefit greatly from waiting for additional tips if ρ is also small.

Instead we propose a *modified myopic* policy to approximate the cost to wait (the second term in (11)). In the modified myopic policy the searcher has two non-engage options:

obtain one additional tip and then engage or wait forever, effectively surrendering to the target, which yields a cost of α. The searcher chooses the option that produces the lower cost. We can easily compute the cost of obtaining one additional tip and then engaging cell 1 (see equation 21 and Appendix I.1). Denoting $\tilde{C}(p)$ the cost from utilizing this modified myopic policy yields

$$\tilde{C}(p) = \min\left(\alpha, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\left(1 - \sum_{i=1}^{n}\max\left(qp_i, \frac{1-q}{n-1}p_1\right)\right)\right), \qquad \text{(EC.124)}$$

We expect that if $(1 - p_1) - \tilde{C}(p)$ is large, the searcher will benefit significantly from waiting. Similarly if $(1 - p_1) - \tilde{C}(p)$ is small or negative, the searcher may choose to engage because the potential benefits of waiting are small compared to the uncertainty in the parameters. In Figure 9 we plot $(1 - p_1) - C(p)$, which represents the true benefit of waiting over engaging, vs. our approximate benefit using the modified myopic policy, $(1 - p_1) - \tilde{C}(p)$. We fix $n = 5$, and for each value of q, α, and ρ, we choose several states p such that $0.21 \leq p_1 \leq 0.8$ and the remaining probability is spread equally over the other four cells. The dashed diagonal line represents when $C(p) = \tilde{C}(p)$: the modified myopic policy produces the optimal cost. The closer the curves lie to the dashed line, the better our approximation is. Small values of p_1 lie on the northeast part of the curves; engaging in these situations is very costly and thus there is great benefit to waiting.

For moderate values of α and ρ (Figure 9b) the approximation performs very well, and it will only improve for larger values of α and ρ. For smaller values of α and ρ (Figure 9a) the approximation is not as accurate, but we can still use $(1 - p_1) - \tilde{C}(p)$ as a rough metric for the potential benefits to waiting in most cases. It is interesting that the approximation in Figure 9a performs the worst for the intermediate q value. For small q, the searcher either optimally engages (for large p_1) or receives a large number of tips (for small p_1), which produces an optimal cost near α. In either case the optimal cost $C(p)$ is close to our approximation $\tilde{C}(p)$. For large q, one tip provides a significant benefit as $p_1^{(+1)} > q$ for any initial p_1, and

thus the modified myopic policy captures a reasonable portion of the benefits from waiting in most cases. However, for intermediate q and small p_1 the searcher must receive multiple tips to generate most of the benefits from waiting. Consequently, the intermediate q curve lies further away from the modified myopic, which only accounts for one additional tip. To summarize our modified myopic policy performs well as a proxy for the benefits from waiting. However, for small values of α and ρ and intermediate values of q and p_1 there may be significant benefits to waiting even if our modified myopic policy suggests otherwise. This is not surprising as our modified myopic policy only captures the benefit of one additional tip, whereas in these aforementioned scenarios the benefits accrue from several additional tips. These results are consistent with our findings in Appendix J.1, where we found the myopic policy performs very well except when ρ is small.

(a) $\alpha = 0.5$ and $\rho = 0.1$ (b) $\alpha = 0.7$ and $\rho = 0.3$
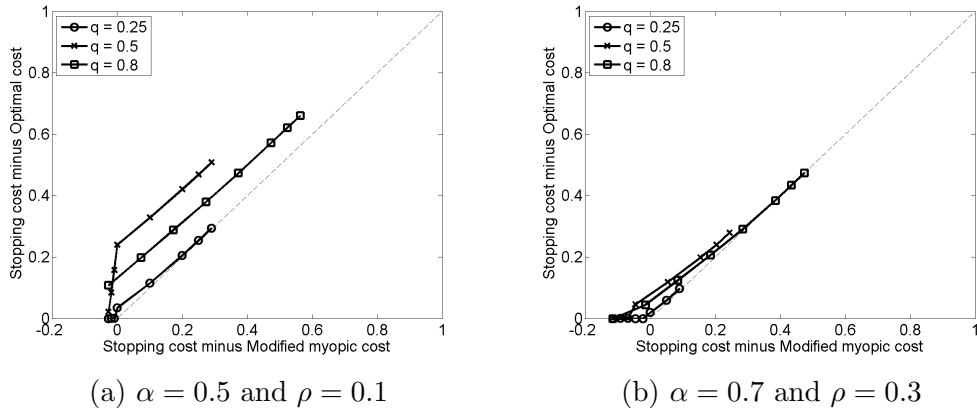
Figure 9: $(1 - p_1) - C(p)$ vs. $(1 - p_1) - \tilde{C}(p)$ for $n = 5$ $q \in \{0.25, 0.5, 0.8\}$, $0.21 \leq p_1 \leq 0.8$ and different combinations of ρ and α.

K   Search Continues if Searcher Chooses Wrong Cell

We take an approach similar to Section 6 and Appendices G–I and define lower and upper bounds on our cost function $C(p)$ akin to equation (EC.98)–(EC.99)

$$
L(p,k) = \begin{cases} \min\left(\min_{j\in A(p)}\left((1-p_j)(c_j + L(p^{(-j)}, k))\right), \dfrac{\rho}{1+\rho}d + \dfrac{1}{1+\rho}\sum_{i\in A(p)} r_i(p)L(p^{(+i)}, k-1)\right) & \text{if } k > 0 \\[3em] & \text{if } k = 0 \\[1em] 0 & \text{or } |A(p)| = 1. \end{cases}
$$

$$\text{(EC.125)}$$

and

$$
U(p,k) = \begin{cases} \min\left(\min_{j\in A(p)}\left((1-p_j)(c_j + U(p^{(-j)}, k))\right), \dfrac{\rho}{1+\rho}d + \dfrac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)U(p^{(+i)}, k-1)\right) & \text{if } k > 0 \\[2em] \min_{j\in A(p)}\left((1-p_j)(c_j + U(p^{(-j)}, k))\right), & \text{if } k = 0 \\[2em] 0, & \text{if } |A(p)| = 1. \end{cases}
$$

$$\text{(EC.126)}$$

$U(p,0)$ corresponds to $K(p)$ in equation (24). We derive its closed form expression in Section K.1. As in the base case $L(p,k) \leq C(p) \leq U(p,k)$, $L(p,k)$ monotonically non-decreases in k, and $U(p,k)$ monotonically non-increases in k. The proofs, which we omit, follow a similar structure to the corresponding proofs for the base case. We perform a simultaneous induction argument in both k and $|A(p)|$, which is why we do not decrement k in the engage component of $L(p,k)$ and $U(p,k)$. We next present conditions to engage or wait for this model.

K.1   Sufficient Condition to Engage

As in section 6.1 and Appendix H we derive a sufficient condition to engage by comparing an upper bound on the cost to engage in (23) to a lower bound on the cost to wait in (23).

The searcher should engage if

$$\min_{j \in A(p)} \left((1 - p_j)(c_j + U(p^{(-j)}, k_s)) \right) \leq \frac{\rho}{1 + \rho} d + \frac{1}{1 + \rho} \sum_{i \in A(p)} r_i(p) L(p^{(+i)}, k_c - 1), \quad \text{(EC.127)}$$

for some $k_s \geq 0$ and $k_c > 0$. The simplest condition occurs when $k_s = 0$ and $k_c = 1$. In this case the right-hand side of (EC.127) yields a lower bound on the cost to wait of $\frac{\rho}{1+\rho} d$

By setting $k_s = 0$, the left-hand side of (EC.127) simplifies to $U(p, 0)$. To compute $U(p, 0)$ we focus on the number of wrong choices made by the searcher, henceforth called collateral events, and the corresponding costs of those collateral events. $U(p, 0)$ describes the situation where the searcher repeatedly chooses cells without receiving additional tips until he successfully finds the target. We assume that the searcher can engage each cell quickly enough so that the target will not execute the attack before the searcher eventually finds the target. Thus to compute $U(p, 0)$, we calculate the expected cost generated from the collateral events. We first determine the order the searcher should engage the cells. Choose an arbitrary ordering and define $p_{(i)}$ as the probability the target is in the ith cell engaged by the searcher. Similarly, we define $c_{(i)}$ as the false positive cost associated with the ith cell engaged by the searcher. Under this ordering, with probability $p_{(1)}$ the searcher chooses correctly on his first engagement, which leads to no collateral damage. With probability $p_{(2)}$ the searcher chooses correctly on his second engagement and incurs collateral damage $c_{(1)}$ from his incorrect first engagement. With probability $p_{(3)}$ the searcher chooses correctly on his third engagement and incurs collateral damage $c_{(1)} + c_{(2)}$ from his incorrect first two engagements. Continuing in this fashion, we can calculate the expected collateral damage caused by this sequential search:

$$\sum_{j=2}^{|A(p)|} p_{(j)} \sum_{i=1}^{j-1} c_{(i)}. \quad \text{(EC.128)}$$

To compute $U(p, 0)$, we need to determine the ordering that minimizes (EC.128). The optimal ordering is ascending by $\frac{c_j}{p_j}$. To see this consider any other ordering and choose a k such that

EC65

$\frac{c_{(k)}}{p_{(k)}} > \frac{c_{(k+1)}}{p_{(k+1)}}$. We now show that if we swap the indices k and $k+1$ we will generate an ordering with a lower cost. Define K_1 as the cost associated with the collateral events using the original ordering and K_2 as the cost with indices k and $k+1$ swapped.

$$K_1 = \sum_{j=2}^{k-1} p_{(j)} \sum_{i=1}^{j-1} c_{(i)} + \sum_{j=k+2}^{|A(p)|} p_{(j)} \sum_{i=1}^{j-1} c_{(i)} + p_{(k)} \sum_{i=1}^{k-1} c_{(i)} + p_{(k+1)} \sum_{i=1}^{k} c_{(i)},$$

$$K_2 = \sum_{j=2}^{k-1} p_{(j)} \sum_{i=1}^{j-1} c_{(i)} + \sum_{j=k+2}^{|A(p)|} p_{(j)} \sum_{i=1}^{j-1} c_{(i)} + p_{(k+1)} \sum_{i=1}^{k-1} c_{(i)} + p_{(k)} \left(\sum_{i=1}^{k-1} c_{(i)} + c_{(k+1)} \right).$$

Most of the terms cancel when we take the difference between K_2 and K_1:

$$K_2 - K_1 = \left(p_{(k+1)} \sum_{i=1}^{k-1} c_{(i)} + p_{(k)} \left(\sum_{i=1}^{k-1} c_{(i)} + c_{(k+1)} \right) \right) - \left(p_{(k)} \sum_{i=1}^{k-1} c_{(i)} + p_{(k+1)} \sum_{i=1}^{k} c_{(i)} \right)$$

$$= p_{(k)} c_{(k+1)} - p_{(k+1)} c_{(k)} = p_{(k)} p_{(k+1)} \left(\frac{c_{(k+1)}}{p_{(k+1)}} - \frac{c_{(k)}}{p_{(k)}} \right) < 0.$$

The last condition follows from our assumption $\frac{c_{(k)}}{p_{(k)}} > \frac{c_{(k+1)}}{p_{(k+1)}}$. Thus the searcher should engage the cells in the order determined by the ratio $\frac{c_j}{p_j}$. Let us define $g(i)$ as the mapping to indices ranked by $\frac{c_j}{p_j}$. So $g(1)$ corresponds to the index j with the smallest ratio $\frac{c_j}{p_j}$, $g(2)$ corresponds to the index j with the second smallest ratio $\frac{c_j}{p_j}$, and $g(A(p))$ corresponds to the index j with the largest ratio $\frac{c_j}{p_j}$. We now define $U(p,0)$ in closed form:

$$K(p) \equiv U(p,0) = \sum_{j=2}^{|A(p)|} p_{g(j)} \sum_{i=1}^{j-1} c_{g(i)}. \tag{EC.129}$$

Putting the pieces together produces our sufficient condition to engage :

$$K(p) = \sum_{j=2}^{|A(p)|} p_{g(j)} \sum_{i=1}^{j-1} c_{g(i)} \leq \frac{\rho}{1+\rho} d.$$

K.2 Sufficient Condition to Wait

We next proceed as in Section 6.2 and Appendix I to define a sufficient condition to wait. We compare a lower bound on the cost to engage in (23) to an upper bound on the cost to wait in (23). The search should wait if

$$\min_{j \in A(p)} \left((1 - p_j)(c_j + L(p^{(-j)}, k_s)) \right) > \frac{\rho}{1 + \rho} d + \frac{1}{1 + \rho} \sum_{i \in A(p)} r_i(p) U(p^{(+i)}, k_c - 1), \quad \text{(EC.130)}$$

for some $k_s \geq 0$ and $k_c > 0$. The simplest condition occurs when $k_s = 0$ and $k_c = 1$. In this case the left-hand side of (EC.130) yields a lower bound to engage of $\min_{j \in A(p)} c_j (1 - p_j)$.

For $k_c = 1$, we can use the expression for $U(p, 0) = K(p)$ from equation (EC.129) in Appendix K.1 to compute the right-hand side of (EC.130). This produces the following condition to wait

$$\min_{j \in A(p)} c_j (1 - p_j) > \frac{\rho}{1 + \rho} d + \frac{1}{1 + \rho} \sum_{i \in A(p)} r_i(p) K(p^{(+i)}).$$

L Other Model Extensions

In this section we present several additional extensions. These extensions relate to the dependence structure of the stream of tips, multiple classes of informants, non-exponential distributions for the time until the attack occurs, the possibility that the system contains no target, and a finite tip scenario.

L.1 Correlated Tips

In the base model we assume independent tips. This assumption seems reasonable if the tips arrive from distinct independent sources. However, it is unlikely that this assumption will hold if there is a single or just a few sources of tips. Dependency in tips would decrease the value of this information and hence the benefits of waiting. One could incorporate dependency

in the stream of tips in a myriad of ways; in this section we present three possibilities.

L.1.1   Single Source

Suppose an informant either provides a new independent tip or provides a tip based on a constant repetitive source (which may be correct or wrong). Assume cell 1 contains the target. If we denote X_j as the jth tip received, then

$$X_j = \begin{cases} Y_j & \text{with probability } \beta \\ Z & \text{with probability } 1 - \beta, \end{cases}$$

where, as in the base case, Y_j are IID random variables such that $Y_j = 1$ with probability q, and $Y_j = i \neq 1$ with probability $\frac{1-q}{n-1}$. Z is a repetitive tip representing a common source of information. The searcher does not know whether X_j corresponds to Y_j or to Z. The random variable Z, with an arbitrary prior distribution, introduces a general dependency to the stream of tips. If $\beta = 1$ the problem simplifies to the base model. If $\beta = 0$ then the informants repeat the same cell every time, and the searcher gains all the information he can after just one tip. We assume the searcher knows the value of β. More realistically we could model β through a probability distribution (e.g., beta) rather than a known parameter. The searcher would update the distribution for β as he collects intelligence.

It no longer suffices for our state probability vector p to just specify the probability that a cell contains a target. We also need to track the likely value of Z. Thus the vector p is a now a joint distribution where

$$p_{ij} = \mathbf{P}[\text{cell } i \text{ contains target, } Z = j \mid \text{given all tips}].$$

As a result, we now must specify as input to the model the prior on the joint distribution for the target's location and the value of Z.

We define the marginal distribution t_i to represent the probability cell i contains the

target and the marginal distribution z_j to represent the probability that $Z = j$:

$$t_i = \sum_{\ell=1}^{n} p_{i\ell}, \tag{EC.131}$$

$$z_j = \sum_{\ell=1}^{n} p_{\ell j}. \tag{EC.132}$$

We next present the updated probability $p_{ij}^{(+k)}$ following a tip indicating cell k

$$p_{ij}^{(+k)} = \begin{cases} \dfrac{((1-\beta)+\beta q)p_{kk}}{(1-\beta)z_k+\beta\left(qt_k+\frac{1-q}{n-1}(1-t_k)\right)} & \text{if } i = j = k \\[3ex] \dfrac{((1-\beta)+\beta\frac{1-q}{n-1})p_{ik}}{(1-\beta)z_k+\beta\left(qt_k+\frac{1-q}{n-1}(1-t_k)\right)} & \text{if } i \neq j = k \\[3ex] \dfrac{\beta q p_{kj}}{(1-\beta)z_k+\beta\left(qt_k+\frac{1-q}{n-1}(1-t_k)\right)} & \text{if } i = k \neq j \\[3ex] \dfrac{\beta\frac{1-q}{n-1}p_{ij}}{(1-\beta)z_k+\beta\left(qt_k+\frac{1-q}{n-1}(1-t_k)\right)} & \text{if } i \neq k, j \neq k. \end{cases} \tag{EC.133}$$

When $\beta = 1$, we see that $t_i^{(+k)} = \sum_{\ell=1}^{n} p_{i\ell}^{(+k)}$ equals the updated $p_i^{(+k)}$ defined by (5) in the base model. When $\beta = 0$, we only need to consider the first tip. In this case if the informant first specifies cell k, then

$$p_{ij}^{(+k)} = \begin{cases} \dfrac{p_{ik}}{z_k} & \text{if } j = k \\[3ex] 0 & \text{if } j \neq k \end{cases},$$

$z_j^{(+k)} = I(j = k)$ and $t_i^{(+k)} = \frac{p_{ik}}{z_k}$. Any follow-up tip provides no new additional information as it would repeat cell k. Indeed for future tips: $p_{ij}^{(+k)} = p_{ik}I(j = k)$, $z_j^{(+k)} = I(j = k)$, and $t_i = p_{ik}$.

To extend $r(p)$ requires conditioning on whether the next tip X_j corresponds to an independent tip Y_j or the single source Z. We compute $r_k(p)$ in the denominator of (EC.133):

$$r_i(p) = (1 - \beta)z_i + \beta\left(qt_i + \frac{1-q}{n-1}(1 - t_i)\right). \tag{EC.134}$$

As expected, when $\beta = 1$ equation (EC.134) simplifies to the original definition of $r_i(p)$ in (7), and when $\beta = 0$ and the informant has provided at least one tip specifying cell k, $r_i(p) = z_i = I(i = k)$.

The new cost function has a similar form to (11). However, we now define p_{ij} by (EC.133), t_i by (EC.131), and $r_i(p)$ by (EC.134).

$$C(p) = \min\left(1 - t_1, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)C(p^{(+i)})\right).$$

The results and conditions from Section 6 remain valid. The only expression that will change is the myopic condition given by (21). We will need to define an analogous expression to (EC.119) in Appendix I.1 for the update function $p_{ij}^{(+k)}$ in (EC.133). The final myopic condition will have a more notationally cumbersome form than the current condition in (21).

Obviously, for the $\beta = 0$ case the searcher needs to accept at most one tip. If the informant specifies cell k, then the searcher either chooses to engage the cell with the largest posterior probability, incurring cost $(1 - \max_j t_j^{(+k)})$, or allows the target to eventually attack inflicting cost α.

L.1.2 Repeat Previous Tip

In this situation we assume that the informant either provides a new independent tip or merely repeats the previous tip. As in Appendix L.1.1, we assume cell 1 contains the target, define $\{Y_j, j \geq 1\}$ as a stream of IID tips, and denote X_j as the jth tip received. Then $X_1 = Y_1$ and for $j > 1$

$$X_j = \begin{cases} Y_j & \text{with probability } \beta \\ X_{j-1} & \text{with probability } 1 - \beta. \end{cases}$$

If $\beta = 1$ then we simplify back to the base model. If $\beta = 0$ then the informant repeats the same cell in every tip, and the searcher gains all the information he can after just one tip.

Here we need to expand our state space from the base model to (p, k). The vector p remains as in the base model with p_i specifying the probability that cell i contains the target given all previous tips. The auxiliary variable k represents the last tip received from the informant. We define the update function $p^{(k,+i)}$ to account for both the new tip i and the previous tip k. If $i \neq k$, then we know the informant provides a new independent tip and the update corresponds to update in the base case (5). When $i = k$, we must account for the possibility that the informant repeats the previous tip. Combining these cases leads to the following probabilities:

$$p_j^{(k,+i)} = \begin{cases} \frac{((1-\beta)+\beta q)p_i}{(1-\beta)+\beta(qp_i+\frac{1-q}{n-1}(1-p_i))} & \text{if } j = i = k \\[3mm] \frac{((1-\beta)+\beta\frac{1-q}{n-1})p_j}{(1-\beta)+\beta(qp_i+\frac{1-q}{n-1}(1-p_i))} & \text{if } j \neq i = k \\[3mm] \frac{\gamma p_i}{\gamma p_i+(1-p_i)} & \text{if } j = i \neq k \\[3mm] \frac{p_j}{\gamma p_i+(1-p_i)} & \text{if } j \neq i \neq k \end{cases} \qquad (\text{EC.135})$$

$$r_i(p, k) = \beta\left(qp_i + \frac{1-q}{n-1}(1-p_i)\right) + (1-\beta)I(i = k).$$

We next rewrite the cost function in (11) as a function of the state (p, k).

$$C(p, k) = \min\left(1 - p_1, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p, k)C(p^{(k,+i)}, i)\right).$$

We can derive conditions to engage or wait similar to those found in Section 6. As described in Appendix L.1.1, the only non-trivial extension of the conditions in Section 6 involves the myopic condition given by (21). We would need to change the expression in (21) to account for the specific update function in (EC.135)

EC71

L.1.3 Biased Incorrect Tips

In the base model we assume that when the informant provides a wrong tip, he specifies an incorrect cell uniformly at random. In this subsection we allow for a bias to exist based on past tips. We maintain the assumption that the informant identifies the correct cell with probability q, but the informant specifies a wrong cell with probability proportional to the current posterior probability of that cell. Formally:

$$\mathbf{P}[\text{informant says cell } i \mid \text{cell } j \text{ contains target}] = \begin{cases} q & \text{if } i = j \\ (1-q)\frac{p_i}{1-p_j} & \text{if } i \neq j. \end{cases} \tag{EC.136}$$

The logic behind (EC.136) derives from research that individuals have primacy biases where they tend to put higher weight on information obtained earlier (Dennis and Ahn, 2001). In our case an informant may provide tips based on heavier weighting of earlier information. The informant may also have reasons to bias his tips toward the searcher's belief.

We next define the probabilities $p^{(+i)}$ and $r(p)$ for this case

$$p_j^{(+i)} = \begin{cases} \dfrac{q}{q+(1-q)\sum_{k\neq i}\frac{p_k}{1-p_k}} & \text{if } j = i \\[2ex] \dfrac{(1-q)\frac{p_j}{1-p_j}}{q+(1-q)\sum_{k\neq i}\frac{p_k}{1-p_k}} & \text{if } j \neq i \end{cases} \tag{EC.137}$$

$$r_i(p) = qp_i + (1-q)\sum_{k\neq i}\frac{p_i}{1-p_k}p_k. \tag{EC.138}$$

Note that $r_i(p)$ is proportional to p_i. For small q, situations may occur where this confirming bias will cascade and the informant may state the same incorrect cell frequently. However, the searcher correctly accounts for this in his update in (EC.137) and does not necessarily put a high probability that the cell frequently stated by the informant actually contains the target.

We use the same cost function as in the base case (Equation (11)). However now we define $p^{(+i)}$ by (EC.137) and $r_i(p)$ by (EC.138). Similarly to Appendices L.1.1 and L.1.2 , the results, conditions, and analysis from Section 6 remain valid with little to no modifications.

L.2   Variable Reliability Among Informants

The base model assumes that all informants are equally active and reliable. This assumption is relaxed here; informant m, $m = 1, ..., M$, provide tips at rate λ_m and reliability q_m. In the base model we define the intensity-ratio $\rho = \mu/\lambda$ to capture the tradeoff between the frequency of tips and the time until a mature attack occurs. In this subsection we include separate parameters for tip rate (λ_m) and the expected time until a mature attack occurs $(1/\mu)$. The stream of tips from the various informants are independent, and successive tips by an informant are also independent.

We need to modify the probabilities $p^{(+i)}$ and $r(p)$ in equations (5) and (7), respectively, to reflect the identity of the informant:

$$p_j^{(m,+i)} = \begin{cases} \frac{\gamma_m p_i}{\gamma_m p_i + (1-p_i)} & \text{if } j = i \text{ and tip from informant } m \\ \\ \frac{p_j}{\gamma_m p_i + (1-p_i)} & \text{if } j \neq i \text{ and tip from informant } m, \end{cases} \quad \text{(EC.139)}$$

where $\gamma_m = \frac{q_m}{\frac{1-q_m}{n-1}}$. To generalize $r(p)$, we only need to modify equation (7) slightly by adding a subscript specifying the identity of the informant:

$$r_{mi}(p) \equiv \mathbf{P}[\text{informant } m \text{ says } i \mid p] = q_m p_i + \frac{1 - q_m}{n - 1}(1 - p_i).$$

Finally, extending the cost function in (11) requires additional bookkeeping to account for the identity of the informant providing the next tip in the *wait* case

$$C(p) = \min\left(1 - p_1, \frac{\mu}{\sum_{k=1}^{M} \lambda_k + \mu}\alpha + \sum_{m=1}^{M} \frac{\lambda_m}{\sum_{k=1}^{M} \lambda_k + \mu} \sum_{i=1}^{n} r_{mi}(p)C(p^{(m,+i)})\right). \quad \text{(EC.140)}$$

Where (EC.139) defines $p^{(m,+i)}$. We can now examine (EC.140) using much of the same analysis from Sections 3–7. The *engage* conditions (15) and (19) generalize to

$$p_1 \geq \frac{\mu}{\sum_{k=1}^{M} \lambda_k + \mu} (1 - \alpha) + \frac{\sum_{k=1}^{M} \lambda_k}{\sum_{k=1}^{M} \lambda_k + \mu} \qquad \text{and}$$

$$p_1 \geq \frac{\mu}{\sum_{k=1}^{M} \lambda_k + \mu} (1 - \alpha)$$

$$+ \sum_{m=1}^{M} \frac{\lambda_m}{\sum_{k=1}^{M} \lambda_k + \mu} \sum_{i=1}^{n} r_{mi}(p) \left(\max \left(\max_j p_j^{(m,+i)}, \frac{\mu}{\sum_{k=1}^{M} \lambda_k + \mu} (1 - \alpha) + \frac{\sum_{k=1}^{M} \lambda_k}{\sum_{k=1}^{M} \lambda_k + \mu} \right) \right),$$

respectively. The myopic condition (21) holds if we add appropriate summations:

$$p_1 \geq \frac{\mu}{\sum_{k=1}^{M} \lambda_k + \mu} (1 - \alpha) + \sum_{m=1}^{M} \frac{\lambda_m}{\sum_{k=1}^{M} \lambda_k + \mu} \sum_{i=1}^{n} \max \left(q_m p_i, \frac{1 - q_m}{n - 1} p_1 \right).$$

L.3   Non-Constant Attack Rate

In the base model, we assume a constant hazard function for the time until the plot matures and the target executes the attack. In reality, this function may increase – if planning of the attack follows some well structured stages – or decrease – if prolonged planning indicates difficulties in executing the attack. Suppose that the informant provides a finite number of tips (T) and does so at deterministic intervals. The target attacks at time period t with probability ν_t, $t = 0, 1, 2 \ldots T$. At period t the searcher first has the opportunity to engage cell 1. If he chooses to wait, then the target executes the attack with a certain probability. If the target does not attack, then the searcher receives the next tip and the process transitions to period $t + 1$. The searcher then has the opportunity to engage cell 1, and the process continues. Let $C(p, t)$ be the expected cost after t tips.

$$C(p, t) = \begin{cases} \min \left(1 - p_1, \frac{\nu_t}{\sum_{k=t}^{T} \nu_k} \alpha + \frac{\sum_{k=t+1}^{T} \nu_k}{\sum_{k=t}^{T} \nu_k} \sum_{i=1}^{n} r_i(p) C(p^{(+i)}, t + 1) \right) & \text{if } t \leq T \\ \alpha & \text{if } t > T. \end{cases} \quad \text{(EC.141)}$$

Solving (EC.141) via backward induction is a straightforward process.

L.4 System Contains No Target

In practice investigators may be pursuing a non-existent plot. At some point the investigation team realizes this and cancels the investigation. In this section we consider the possibility that no target exists, and the system is empty. The searcher now has three options after receiving a tip: engage cell 1, call off the investigation, or wait for another tip. To make the problem more realistic and mathematically interesting, we add a cost η associated with receiving each tip. This cost could be the actual amount paid to the informant or more abstractly represent the expected cost required to maintain an active investigation between tips. This situation occurs in practice: an investigative agency may decide to close the investigation and shift resources to other areas because it deems the likelihood the plot actually exists to be low.

We still label the cells $1, 2, 3, \ldots, n$. For notational convenience we introduce cell 0 to represent the no-target situation. We will assume the informant only specifies actual cells in his tips. That is he will never say the system is empty, although we could easily modify the problem to allow for this. We use the same state vector p, with an additional element. We define p_i as the probability cell $i > 0$ contains the target and p_0 denotes the probability that there is no target. As in the base model, if the target truly resides in cell $i > 0$, then the informant will point to cell i with probability q and will point to each wrong cell uniformly with probability $\frac{1-q}{n-1}$. If the system is empty, we assume the informant will point to each cell uniformly with probability $\frac{1}{n}$. Unlike in the base case, the target's location is not independent of whether a mature attack occurs before the next tip. If the system contains no target, then the searcher will receive the next tip with certainty. If ρ is very large and the searcher receives a tip, it is likely the system is empty. We must modify the transition dynamics slightly to account for this dependence. We redefine $r_i(p)$ as the probability the searcher receives a tip for cell i if the searcher chooses to wait

$$ r_i(p) = \frac{1}{1+\rho} \left(qp_i + \frac{1-q}{n-1}(1 - p_i - p_0) \right) + \frac{p_0}{n}. $$

Note that the sum over $r_i(p)$ for all i is not 1; it is the probability the searcher receives a tip (i.e., a mature attack does not occur):

$$\sum_{i=1}^{n} r_i(p) = \frac{1}{1+\rho}(1-p_0) + p_0 = 1 - \frac{\rho}{1+\rho}(1-p_0).$$

With this redefined $r_i(p)$, we can also modify $p_j^{(+i)}$

$$p_j^{(+i)} = \begin{cases} \frac{1}{1+\rho}\frac{qp_i}{r_i(p)} & \text{if } j = i > 0 \\[2mm] \frac{1}{1+\rho}\frac{\frac{1-q}{n-1}p_j}{r_i(p)} & \text{if } j \neq i, 0 \\[2mm] \frac{\frac{p_0}{n}}{r_i(p)} & \text{if } j = 0. \end{cases}$$

We modify our cost function to incorporate the new option of closing the investigation:

$$C(p) = \min\left(1 - p_1, \alpha(1-p_0), \frac{\rho}{1+\rho}(1-p_0)\alpha + \left(\frac{1}{1+\rho}(1-p_0) + p_0\right)\eta + \sum_{i=1}^{n} r_i(p)C(p^{(+i)})\right).$$

If $\eta = 0$ it is straightforward to show that calling off the investigation now produces the same expected cost as receiving one additional tip and then calling off the investigation. Without some additional cost for collecting a tip, the searcher has no incentive to call off the investigation.

We can now derive similar conditions to the original problem. For example, the searcher should call off the investigation if

$$\alpha(1-p_0) < \min\left(\frac{\rho}{1+\rho}(1-p_0)\alpha + \left(\frac{1}{1+\rho}(1-p_0) + p_0\right)\eta, 1 - p_1\right),$$

and the searcher should engage cell 1 if

$$1 - p_1 < \min\left(\frac{\rho}{1+\rho}(1-p_0)\alpha + \left(\frac{1}{1+\rho}(1-p_0) + p_0\right)\eta, \alpha(1-p_0)\right).$$

L.5 Finite Number of Tips

In some situations (e.g., when the searcher knows that the informant will soon disappear) the informant only provides a finite number of tips. In this case, the decision to engage must take into account the fact that the "inventory" of tips continuously decreases.

If after the maximum number of tips, say M, the searcher does not engage a cell then the target will eventually execute his plan inflicting damage α. We modify the cost function $C(p)$ in (11) to track the number of tips received. We define $C(p, m)$ as the expected cost when the searcher has received m tips:

$$
C(p, m) = \begin{cases} \min\left(1 - p_1, \frac{\rho}{1+\rho}\alpha + \frac{1}{1+\rho}\sum_{i=1}^{n} r_i(p)C(p^{(+i)}, m+1)\right) & \text{if } m \leq M \\ \alpha & \text{if } m > M. \end{cases}
\tag{EC.142}
$$

Note the similarities between $C(p, m)$ in (EC.142) and $L(p, k)$ in (EC.98) and $U(p, k)$ in (EC.99). The terminal case in (EC.142) differs from those in (EC.98) or (EC.99), but the general structure of the cost function is identical in all three cases. The structure is also very similar to the non-constant attack rate extension defined in equation (EC.141).

In the remainder of this subsection we assume a uniform prior. In the no-informant case ($M = 0$), the searcher must weigh the false positive cost and small chance of choosing correctly vs. the guaranteed damage if the searcher does nothing. In this case the searcher should engage a randomly selected cell if and only if $\frac{1}{n} \geq 1 - \alpha$. Obviously, if $\alpha > 1$ then the searcher should always engage.

If the informant only provides one tip ($M = 1$) then the searcher should engage a random cell before receiving the tip if and only if

$$
\frac{1}{n} \geq \frac{\rho}{1 + \rho}\left(1 - \alpha\right) + \frac{1}{1 + \rho}\max\left(q, 1 - \alpha\right).
\tag{EC.143}
$$

If the searcher chooses to wait for the tip, then the searcher should engage the cell corre-

sponding to the tip if

$$q \geq 1 - \alpha, \tag{EC.144}$$

and otherwise do nothing and accept the target's attack and the associated cost α. Examining (EC.143) we observe that, as expected, the searcher is less likely to immediately and randomly engage a cell when $M = 1$ than in the $M = 0$ case. Furthermore, since $q > \frac{1}{n}$, condition (EC.144) reveals that the searcher is more likely to choose a cell after receiving one tip than randomly engaging one when $M = 0$.

Finally, note that the $M = 1$ case has a similar form to the myopic policy described in Section 6.2 (compare (EC.143) to (22)). They differ slightly because the $M = 1$ situation allows the searcher to never act, whereas the myopic policy assumes the searcher will either act in the current period or the next period.