



Operations Research

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

When Is Information Sufficient for Action? Search with Unreliable yet Informative Intelligence

Michael Atkinson, Moshe Kress, Rutger-Jan Lange

To cite this article:

Michael Atkinson, Moshe Kress, Rutger-Jan Lange (2016) When Is Information Sufficient for Action? Search with Unreliable yet Informative Intelligence. Operations Research 64(2):315-328. <http://dx.doi.org/10.1287/opre.2016.1488>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2016, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

When Is Information Sufficient for Action? Search with Unreliable yet Informative Intelligence

Michael Atkinson, Moshe Kress

Operations Research Department, Naval Postgraduate School, Monterey, California 93943
{mpatkins@nps.edu, mkress@nps.edu}

Rutger-Jan Lange

Department of Finance, VU University Amsterdam, 1081 HV Amsterdam, The Netherlands; and
Erasmus School of Economics, Erasmus University Rotterdam, 3062 PA Rotterdam, The Netherlands, rutger-jan.lange@cantab.net

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.

Subject classifications: military: search/surveillance; decision analysis: sequential; dynamic programming: applications.

Area of review: Military and Homeland Security.

History: Received May 2014; revisions received February 2015, September 2015; accepted December 2015. Published online in *Articles in Advance* March 30, 2016.

1. Introduction

Operation Neptune Spear led to the capture and elimination of Osama bin Laden by the United States in 2011. Although U.S. 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 trade-off 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 §2 and formulate the mathematical model in §3. The cases of $n = 2$ and $n = \infty$ appear in §§4 and 5, respectively. Section 6 examines the optimal

strategy when $2 < n < \infty$. We present numerical illustrations in §7. Section 8 discusses extensions. All proofs appear in the online appendix (available as supplemental material at <http://dx.doi.org/10.1287/opre.2016.1488>).

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). Although 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 §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 *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, $q > 1/n$. 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, 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 §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. Although 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, 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, which specifies the likelihood that the next tip points at cell i . In §3.1 we compute these probabilities, and in §3.2 we use these probabilities to derive the expected costs.

3.1. Location and Pointing Probabilities

Let $p = (p_1, \dots, p_n)$ 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 $s = (s_1, \dots, s_n)$. In this subsection we assume that $s_1 \geq \dots \geq s_n$. The location probability of cell i given s is

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

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

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

where

$$\gamma = \frac{q}{(1-q)/(n-1)}. \quad (3)$$

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 equally likely to be pointed at. When we substitute (2) back into (1), most terms cancel, and the location probability simplifies to

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

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 Δ_j , $j = 1, \dots, n-1$.

Although s or Δ are natural state vectors, it is simpler to use the location probabilities $p = (p_1, \dots, p_n)$ as the state vector for most of the mathematical analysis in §§4–6. Specifically, if the next tip points to cell i , then the updated probability $p_j^{(+i)}$ for cell j is

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

Recall that according to our assumption $q > 1/n$ 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:

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

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 :

$$\begin{aligned} r_i(p) &\equiv \mathbf{P}[\text{informant says } i \mid p] \\ &= \sum_{k=1}^n \mathbf{P}[\text{informant says } i \mid p, \text{target in } k] \\ &\quad \cdot \mathbf{P}[\text{target in } k \mid p] \\ &= qp_i + \frac{1-q}{n-1} \sum_{k \neq i} p_k = qp_i + \frac{1-q}{n-1} (1-p_i). \end{aligned} \quad (7)$$

Inspection of (7) reveals that $r_i(p) \in [(1-q)/(n-1), q]$. Thus, a tip may point at a cell other than i , even if p_i is close to 1, if $q \ll 1$. 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,

$$C(p) = \min \left(\begin{array}{l} \text{expected cost if the searcher engages a cell,} \\ \frac{\rho}{1+\rho} \alpha + \frac{1}{1+\rho} \text{expected cost after} \\ \text{receiving the next tip} \end{array} \right). \quad (8)$$

If the searcher decides to wait, the target may attack before the searcher receives the next tip. In that case, which happens with probability $\rho/(1+\rho)$, 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 $p_1 \geq \dots \geq p_n$. Therefore, $B(p)$ contains cell 1 and

$$\begin{aligned} \mathbf{E}[\text{Cost if searcher decides to engage} \mid p] \\ = 1 - \max_j p_j = 1 - p_1. \end{aligned} \quad (9)$$

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

$$\begin{aligned} \mathbf{E}[\text{Cost if waiting for and receiving the next tip} \mid p] \\ = \sum_{i=1}^n \mathbf{P}[\text{informant says } i \mid p] C(p^{(+)}) \\ = \sum_{i=1}^n r_i(p) C(p^{(+)}) \end{aligned} \quad (10)$$

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

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

If the searcher is indifferent between engaging and waiting, he engages. In Appendix B we present characteristics of $C(p)$, such as its concavity. Because 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

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

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 prior state before the arrival of the $s_1 + s_2$ tips. Using Equation (4) we transform \tilde{p} to p :

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

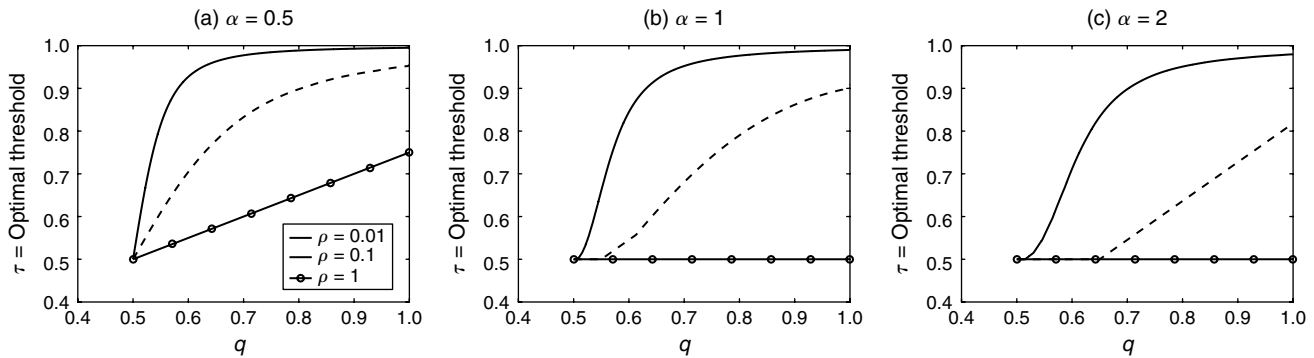
$$p_2 = \frac{\gamma^{s_2} (1 - \tilde{p}_1)}{\gamma^{s_1} \tilde{p}_1 + \gamma^{s_2} (1 - \tilde{p}_1)} = \frac{1 - \tilde{p}_1}{\gamma^{s_1 - s_2} \tilde{p}_1 + (1 - \tilde{p}_1)}. \quad (14)$$

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 nonnegative 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 Figures 1(a) to 1(c), 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 nondecreasing 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 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 .

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

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



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 2(b) and 2(c) 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 nonmonotonicity 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 , and 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 $-o-$ curves, which are close to zero.

The jumps in Figure 2 occur when the optimal tip-differential changes by one. 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.

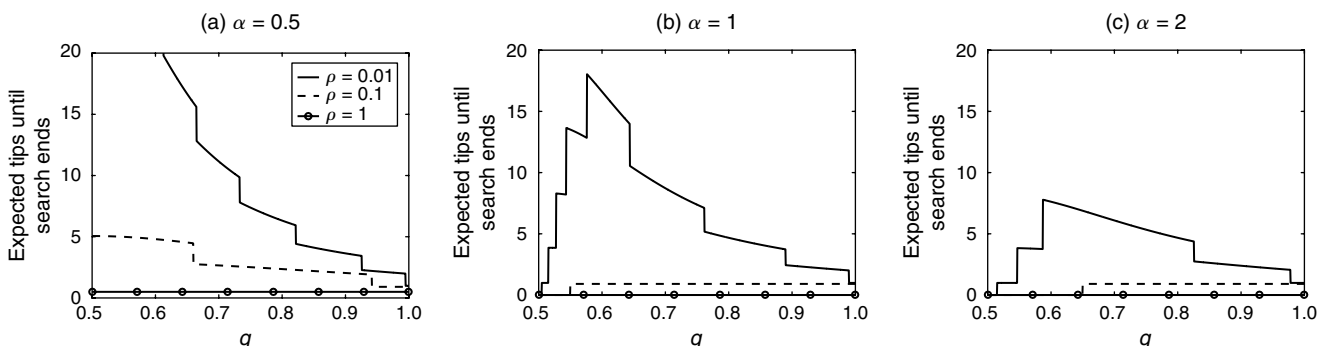
5. The Case of an Infinite Number of Cells

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 $(\rho/(1 + \rho))\alpha + (1/(1 + \rho))(1 - q)$;*
3. *Wait for the confirming tip and then engage: cost is $\alpha(1 - (q/(\rho + q))^2)$.*

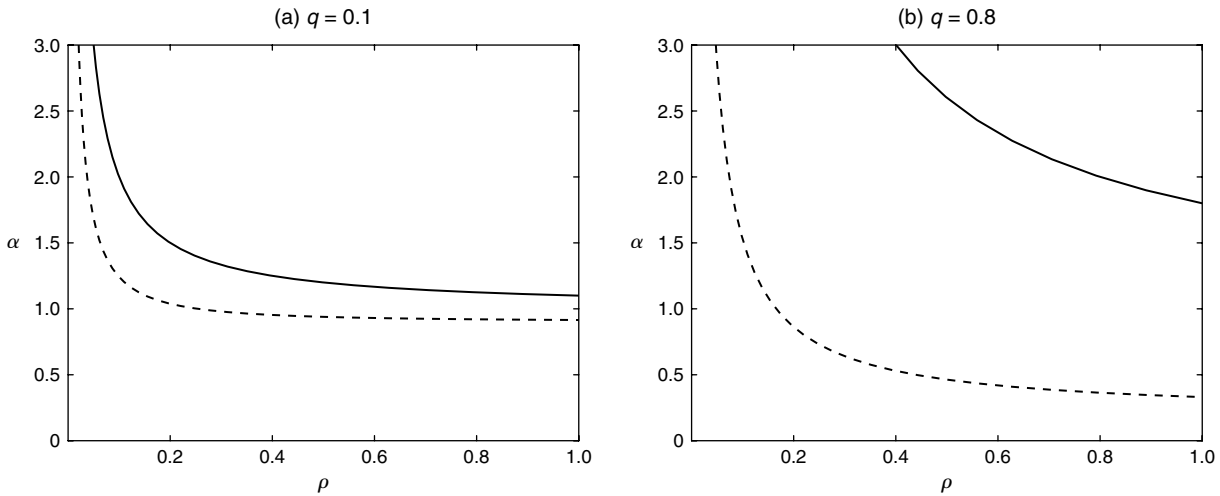
Figure 2. Expected number of tips, starting from the uniform state $p = (0.5, 0.5)$ until the search ends as a function of q for fixed combinations of $\rho \in \{0.01, 0.1, 1\}$ and $\alpha \in \{0.5, 1, 2\}$.



Note. The search ends either when the searcher engages or when a mature attack occurs.

Downloaded from informs.org by [205.155.65.226] on 21 September 2016, at 12:49. For personal use only, all rights reserved.

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.



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) because $p_1 = q$ after the tip. If the searcher chooses option 3, there is no collateral damage, but the target may execute the attack before the 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 because one tip provides significant information for larger values of q .

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 $(\rho/(1 + \rho))\alpha$. Thus, the searcher should engage now if and only if

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

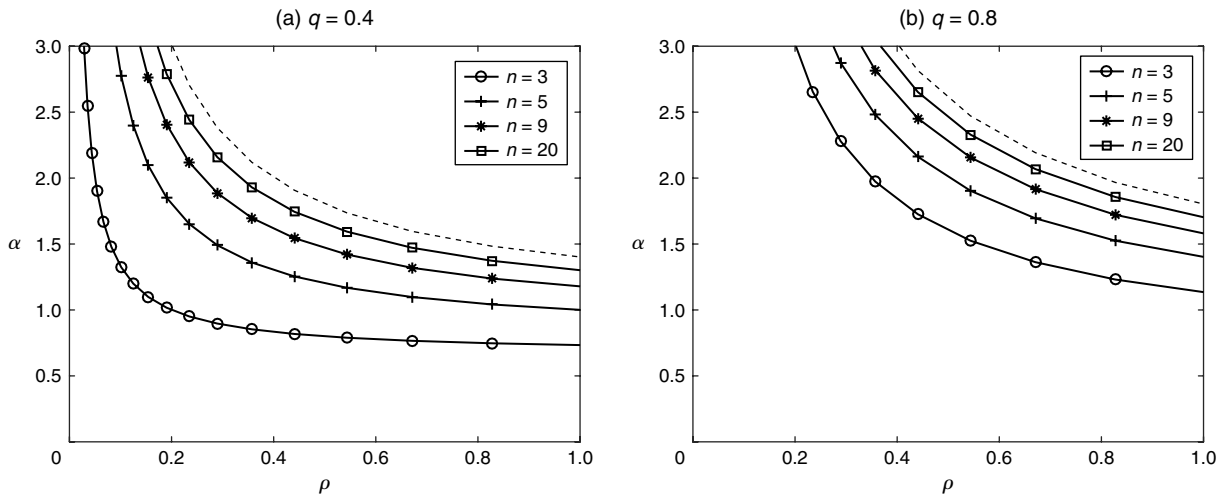
Condition (15) is sufficient to engage for any value of q . We derive this formally in §6.1. This observation leads to the following preliminary analysis for the case where $q < 1$ and the searcher has no prior information: $p_1 = \dots = p_n = 1/n$. In that case the searcher engages any cell before receiving a tip if $1/n \geq (\rho/(1 + \rho))(1 - \alpha) + 1/(1 + \rho)$. 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 $p = (1/n, 1/n, \dots, 1/n)$ transitions to $p^{(+1)} = (q, (1 - q)/(n - 1), \dots, (1 - q)/(n - 1))$ (see Equation (5)) and the expected cost is $(\rho/(1 + \rho))\alpha + (1/(1 + \rho))(1 - q)$. Thus, if $1 - 1/n > (\rho/(1 + \rho))\alpha + (1/(1 + \rho))(1 - q)$, the searcher should wait. In summary, we have

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

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

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), then 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))

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



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 the northeastern portion of Figure 4, which represents the limiting case as $n \rightarrow \infty$.

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

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 $\{p \mid p_1 \geq \dots \geq p_n, \sum_{i=1}^n p_i = 1\}$ is an $n - 1$ dimensional closed convex set, and thus we should not be surprised that the optimal policy cannot be represented by a one-dimensional subspace. Because 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 §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 §§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:

$$\text{engage if } q \geq \frac{\rho}{1 + \rho}(1 - \alpha) + \frac{1}{1 + \rho},$$

for uniform prior and $|B(p)| = 1$. (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.

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

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, ..., 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 $\hat{q} = \gamma/(1+\gamma)$, 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:

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

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 k th 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 or (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

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

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):

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

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 two-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 $\rho \gg 1$ (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 $\tilde{p} = (1/n, \dots, 1/n)$ 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 $p_i = (1-q)/(n-1)$ 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

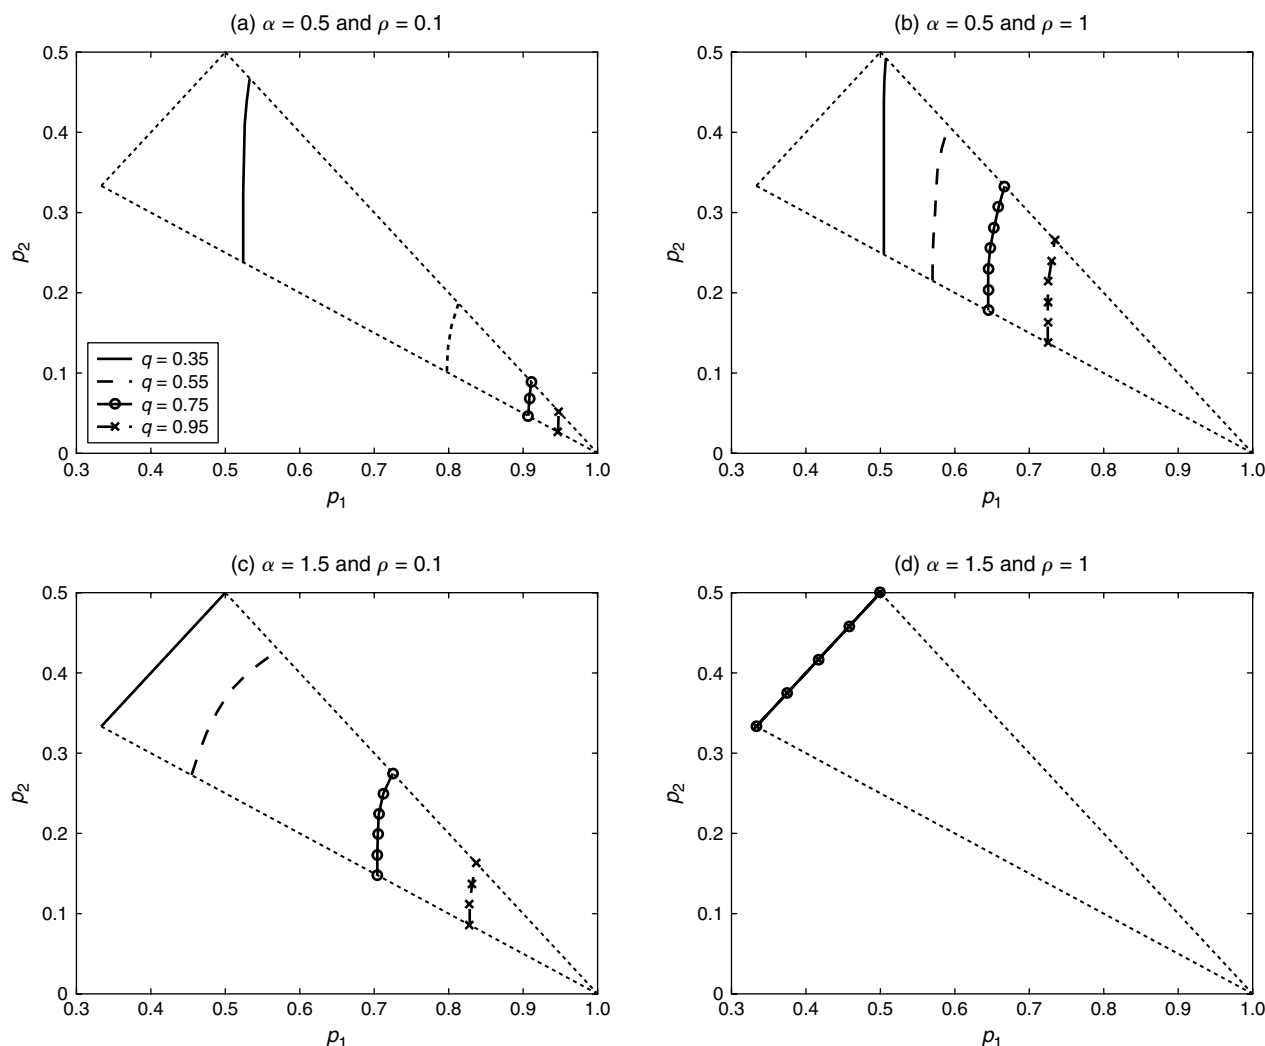
$$\frac{1}{n} \geq \frac{\rho}{1+\rho}(1-\alpha) + \frac{1}{1+\rho} q.$$

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 §6. Subsection 7.1 examines the three-cell case and in §7.2 we analyze the effect of number of cells on the expected cost.

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 $\rho \in \{0.1, 1\}$.



Note. The engage region lies to the southeast of each curve.

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 §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.

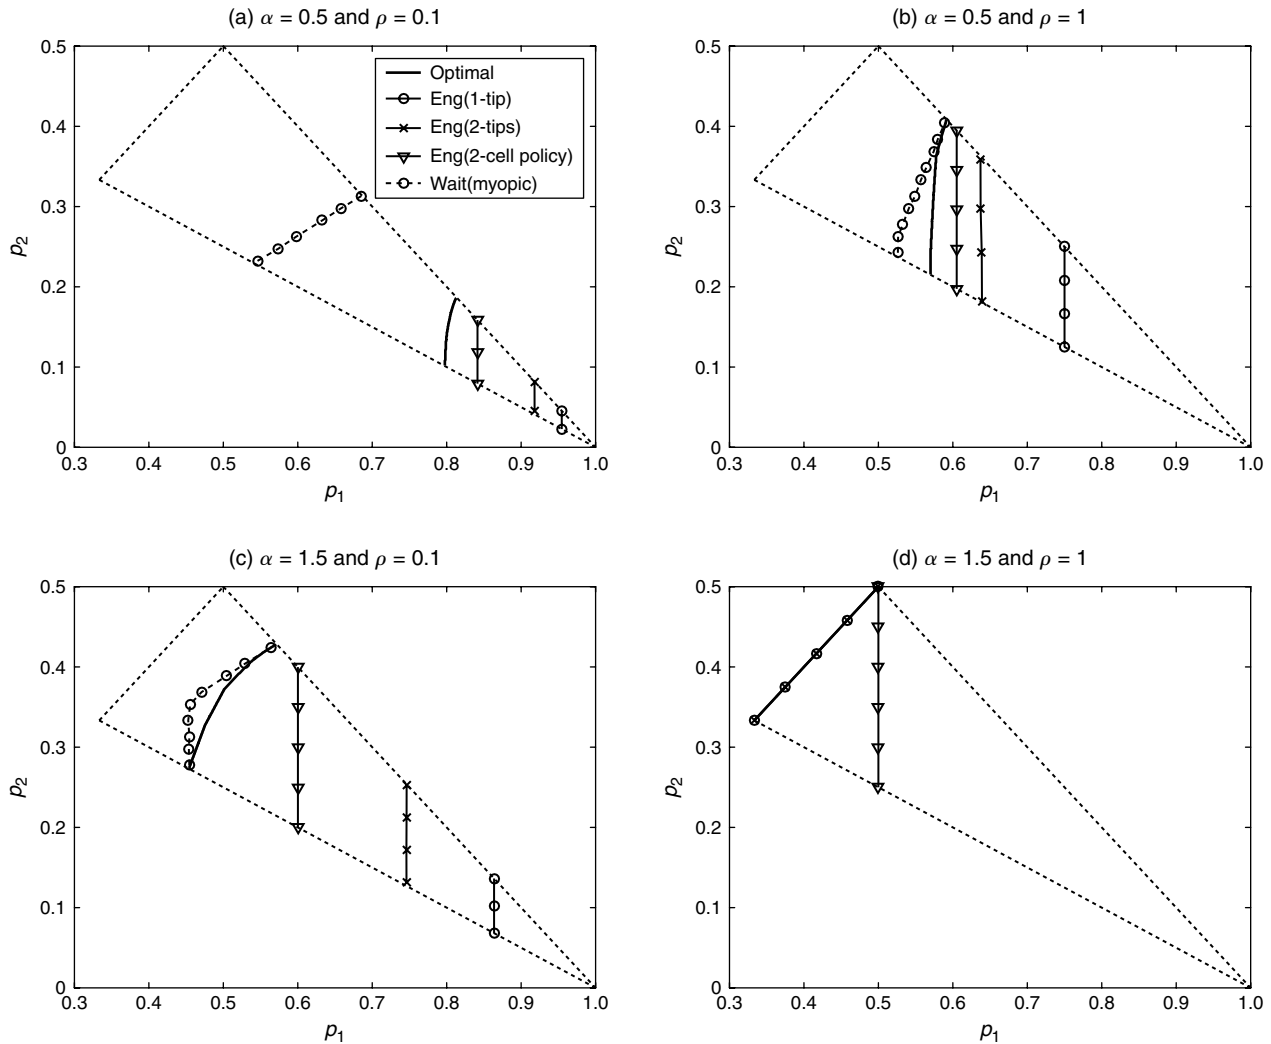
Similarly to the two-cell case, the engage region decreases with the reliability of the informant because the benefit from additional tips increases. Larger values of α or ρ increase the size of the engage region because the cost or likelihood of an attack increases. For larger value of ρ (Figures 5(b) and 5(d)), the boundaries for the various reliability values are closer together than for smaller ρ (Figures 5(a) and 5(c)).

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 5(d) 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 .

In §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 §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 $-o-$ curve.

Figure 6. Engage region for various heuristic policies for $q = 0.55$ and combinations of $\alpha \in \{0.5, 1.5\}$ and $\rho \in \{0.1, 1\}$.



Note. The engage region lies to the southeast of each curve.

- Condition (19), corresponding to perfect detection after two tips, is denoted *eng(2-tips)* and represented by the $- \times -$ curve. As discussed in §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 $- \nabla -$ curve.

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

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

The *wait(myopic)* heuristic performs very well except for small values of α and ρ (Figure 6(a)). 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 6(a). Intuitively, engaging seems like the right decision for this state because 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 6(a).

We also examine how much the cost increases using a heuristic instead of the optimal policy by generating 84,000 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 6(a) 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 §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 7(a))

and (b) high-cost, high-risk (Figure 7(b)). 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:

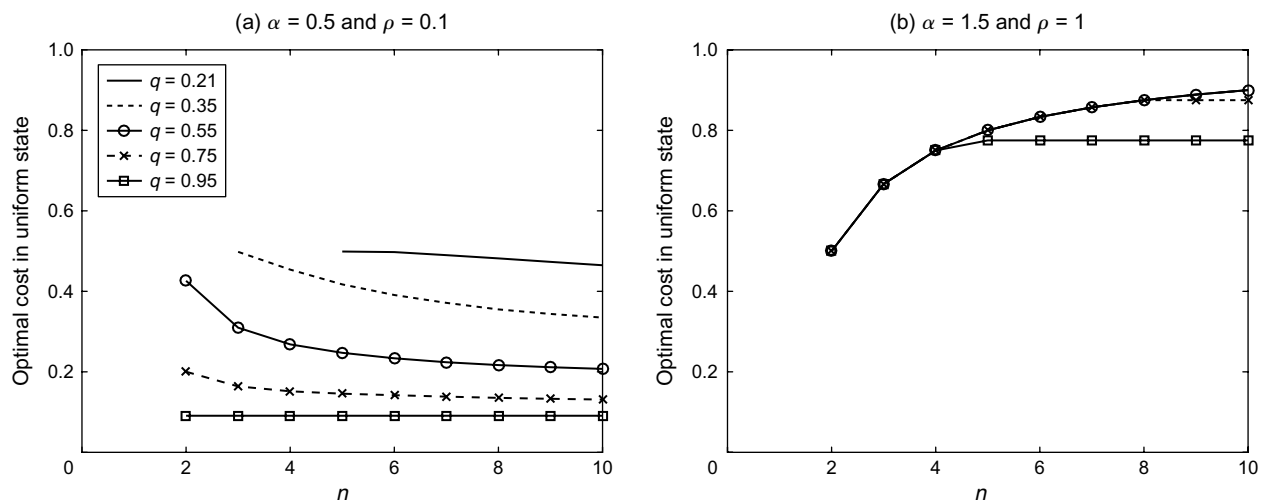
1. Blind engagement scenario: searcher engages a cell uniformly at random incurring cost of $(n - 1)/n$.
2. The searcher obtains one tip and engages the corresponding cell, which incurs cost $(\rho/(1 + \rho))\alpha + (1/(1 + \rho))(1 - q)$.
3. The searcher obtains at least two tips.

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 7(b), the searcher chooses either option 1 (when the curves increase) or option 2 (when the curves flatten out). For small α and ρ (Figure 7(a)), the searcher chooses either option 2 or 3. Even though the cost is nonincreasing with n in Figure 7(a), the cost significantly decreases for only moderate values of q and the curves flatten out quickly.

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 §§3–6 generalize in a natural way. Because of 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

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 .



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 nonexponential 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 to 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 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,

$\mathbf{P}[\text{informant says } i \mid p, \text{ target in } k]$

$$= \begin{cases} \frac{q}{q + (|A(p)| - 1)(1 - q)/(n - 1)} & \text{if } i = k \\ \frac{(1 - q)/(n - 1)}{q + (|A(p)| - 1)(1 - q)/(n - 1)} & \text{if } i \neq k. \end{cases}$$

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

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

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

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

Although 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:

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

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 §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 c_j/p_j , $j = 1, \dots, n$. Let $g(i)$ denote the index of the i th smallest value of c_j/p_j in $A(p)$. Thus, $g(1)$ and $g(|A(p)|)$ are the indices of the cells with the smallest and largest ratios c_j/p_j , respectively. Let $K(p)$ denote the cost of this policy. In the Appendix K we show that

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

The searcher should engage a cell if $K(p) \leq (\rho/(1 + \rho))d$. 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

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

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

9. Summary and Conclusions

In this paper we study a time-critical variant of the whereabouts 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 counterterror and counterinsurgency operations, which rely on human intelligence and intercepted communications. Unlike the original whereabouts 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 action—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 trade-offs 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 §8 and Appendix L to capture alternative scenarios. These include the search continuing after an incorrect engagement, multiple types of informants, and nonexponential 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 the following: (1) the optimal number of tips received by the searcher may not be monotone as a function of the informant reliability (see §4) and (2) in many cases there is little to no reduction in the optimal cost as we increase the number of cells (see §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 §8.1). In this case the searcher could search multiple cells 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.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/opre.2016.1488>.

Acknowledgments

The authors would like to thank their colleagues in the Department of Operations Research 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 to thank the referees and associate editor for invaluable comments that significantly improved the paper. The first and second author are supported by the Office of Naval Research. The third author is supported by the European Union Seventh Framework Programme [FP7-SSH/2007-2013, Grant Agreement 320270-SYRTO].

References

- An B, Brown M, Vorobeychik Y, Tambe M (2013) Security games with surveillance cost and optimal timing of attack execution. *Proc. 2013 Internat. Conf. Autonomous Agents and Multi-Agent Systems* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC), 223–230.
- Bechhofer RE, Elmaghraby S, Morse N (1959) A single-sample multiple-decision procedure for selecting the multinomial event which has the highest probability. *Ann. Math. Statist.* 30(1):102–119.
- Berger J (1985) *Statistical Decision Theory and Bayesian Analysis* (Springer, New York).
- Carlstein E, Müller H-G, Siegmund D (1994) *Change-Point Problems* (Institute of Mathematical Statistics, Hayward, CA).
- Chen P (1988) An integrated formulation for selecting the most probable multinomial cell. *Ann. Institute Statist. Math.* 40(3):615–625.
- Chow YS, Robbins H, Siegmund D (1971) *Great Expectations: The Theory of Optimal Stopping* (Houghton Mifflin, Boston).
- Dayanik S, Yu AJ (2013) Reward-rate maximization in sequential identification under a stochastic deadline. *SIAM J. Control Optim.* 51(4):2922–2948.
- Ferguson TS (2004) *Optimal Stopping and Applications* (Mathematics Department, UCLA), <http://www.math.ucla.edu/~tom/Stopping/Contents.html>.
- Frazier P, Yu AJ (2007) Sequential hypothesis testing under stochastic deadlines. Platt J, Koller D, Singer Y, Roweis S, eds. *Advances in Neural Information Processing Systems 20 (NIPS 2007)* (MIT Press, Cambridge, MA), 465–472.
- Kadane JB (1971) Optimal whereabouts search. *Oper. Res.* 19(4):894–904.
- Kaplan EH (2010) Terror queues. *Oper. Res.* 58(4):773–784.
- Kaplan EH (2012) Estimating the duration of jihadi terror plots in the United States. *Stud. Conflict and Terrorism* 35(12):880–894.
- Kim S-H, Nelson BL (2006) Selecting the best system. Henderson S, Nelson B, eds. *Handbooks in Operations Research and Management Science: Simulation* (Elsevier, Amsterdam), 501–534.
- Komiya T, Iida K, Hohzaki R (2006) An optimal investigation in two stage search with recognition errors. *J. Oper. Res. Soc. Japan* 49(2):130–143.
- Kress M, Lin KY, Szechtman R (2008) Optimal discrete search with imperfect specificity. *Math. Methods Oper. Res.* 68(3):539–549.
- Lange R-J (2012) *Brownian motion and multidimensional decision making*. Unpublished doctoral dissertation, University of Cambridge, UK.
- Ramey JT, Alam K (1979) A sequential procedure for selecting the most probable multinomial event. *Biometrika* 66(1):171–173.

- Shiryayev A (2007) *Optimal Stopping Rules* (Springer, Berlin).
- Stone LD (1975) *Theory of Optimal Search* (Academic Press, New York).
- Szechtman R, Kress M, Lin K, Cfir D (2008) Models of sensor operations for border surveillance. *Naval Res. Logist. (NRL)* 55(1):27–41.
- Vieira H, Sanchez SM, Sanchez PJ, Kienitz K, Belderrain M (2014) A restricted multinomial hybrid selection procedure. *ACM Trans. Modeling Comput. Simulation (TOMACS)* 24(2):10.
- Wald A, Wolfowitz J (1948) Optimum character of the sequential probability ratio test. *Ann. Math. Statist.* 19(3):326–339.
- Washburn A (2014) *Two-Person Zero-Sum Games*, 4th ed. (Springer, New York).
- Wilson KE, Szechtman R, Atkinson MP (2011) A sequential perspective on searching for static targets. *Eur. J. Oper. Res.* 215(1):218–226.

Michael P. Atkinson is an associate professor in the Operations Research Department at the Naval Postgraduate School. His research focuses on applying stochastic models to study homeland security and military applications.

Moshe Kress is a professor in the Department of Operations Research at the Naval Postgraduate School. His general research area is defense operations research with focus on combat and counter-terror modeling and operational logistics.

Rutger-Jan Lange is currently a post-doctoral researcher at VU University Amsterdam. His research interests range from Markov decision processes to time-series econometrics with financial applications.