Supplemental Material: Proofs appear in online appendix.

informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000

Submitted to *Operations Research*

ISSN 0030-364X, EISSN 1526-5463

Hard and Soft Defense Against a Sequence of Aerial Threats

Michael P. Atkinson, Moshe Kress
Operations Research Department, Naval Postgraduate School, mpatkins@nps.edu, mkress@nps.edu

Abstract. The increasing prevalence of missiles and drones (hereafter referred to as *threats*) in attacks by both state and non-state actors highlights the critical need for a robust defense system to counter these threats. We develop a combat model for the engagement between a Blue defender who is subject to repeated attacks by Red threats. The defender employs two types of defenses: hard interceptors, such as anti-ballistic missiles, and soft measures, such as directed-energy weapons and jamming. Employing strategies for these two types of defensive options are evaluated by a two-dimensional measure of effectiveness: expected number of leaking Red threats, and the expected expenditure of hard interceptors. We define efficient frontiers on this two-dimensional space and identify defense strategies that compose these frontiers.

Funding: This research was supported by Naval Postgraduate School, Naval Research Program (PE 0605853N/2098).

Key words: Firing theory, combat model, salvos, missile defense, soft measures, interceptors

History: Received: May 20, 2024; Accepted: February 18, 2025; Published Online: Month DD, YYYY

1. Introduction

Recent military conflicts in Ukraine and Israel (Lendon 2024), and events in the Red Sea (Hjelmgaard and Brook 2024), underscore the importance of air defense against missiles and drones. Indeed, several air-defense missile systems, such as the Arrow (Majumdar 2018) and the Iron-Dome (Armstrong 2014) have been developed and deployed in recent years. Henceforth we call such defensive missiles *interceptors*. For short-range threats (e.g., attack by rockets and/or mortar bombs) the engagement opportunity time-window for interceptors is narrow and a defense system may only be able to engage the threat once. However, for long-range threats, such as drones and ballistic or cruise missiles, the defender may have several opportunities to engage the incoming $threat$ with salvos of interceptors. In that case, the defender can observe the outcome of each engagement before deciding if to engage the threat again. This tactic, which is the focus of this paper, is called shoot-look-shoot (SLS) (Washburn and Kress 2009). Specifically, consider the situation where a defender (called henceforth *Blue*) is subject to an attack by a sequence of homogeneous *Red* threats launched one at a time. Blue defends itself by firing several salvos of interceptors at each threat. After each salvo, Blue performs *battle damage assessment* (BDA) about the outcome of the salvo, which is binary: either the threat is intercepted and *killed* or it is still *alive* and threatening Blue. We call the use of interceptors *hard defense* (HD). While our modeling framework allows for imperfect BDA in HD engagements, we assume neither false-positive errors (threat is declared killed while it is still alive) nor false-negative errors (threat is declared alive while actually killed). The physical characteristics of HD engagements make this assumption reasonable.

HD is expensive, with each interceptor costing between 2 and 10 million USD (Williams 2024). Therefore, HD may be augmented by *soft measures* (SM), comprising actions such as jamming the command channel to the threat, employing decoys that divert the threat from its planned trajectory, and engaging the threat with directed-energy weapons (van Hooft and Boswinke 2021). A SM, which we assume is effectively an unlimited defense resource, is applied during the engagement opportunity time-window along with the HD. Similar to the HD, we assume no false-positive signals for the SM; the threat is not declared killed if it is still alive. However, unlike the HD, false-negative errors may occur for SM. Blue may consider the threat alive even if the threat has already been killed by the SM. HD combined with SM is *combined defense* (CD). In this paper we address both HD and CD.

We model and study questions about how Blue decides what is the maximum number of interceptors to allocate for each threat, how to distribute this number among the possible salvos fired at a threat, and how these parameters are affected by the presence and characteristics of SM. We focus on two primary measures of effectiveness (MOE): expected number of Red *leakers* - threats that penetrate the missile-defense - and the expected number of expended HD interceptors. The defender wishes to minimize both MOEs.

SLS firing models have been studied for over five decades (Eckler and Burr 1972, Soland 1987, Aviv and Kress 1997). A review, as of 2004, and extensions of SLS models are given in Glazebrook and Washburn (2004). Washburn (2002) gives an overview of firing theory and focuses on SLS models in section 4. Our problem is related to the "Bomber problem" (Simons and Yao 1990, Weber 2013) and the "Fighter problem" (Bartroff and Samuel-Cahn 2011). In the Bomber problem, a bomber attempts to survive a sequence of attacks by enemy aircraft. In the fighter problem, a fighter plane fires at a sequence of enemy targets. Other related references for salvo analysis and missile defense are Hughes Jr (1995), Armstrong (2004, 2005), Brown et al. (2005), Davis et al. (2017), Kline et al. (2019), Karasakal et al. (2011), Dutta (2014). The most relevant reference to our paper is by Kalyanam and Clarkson (2021) who prove that, in the case of a single threat, the optimal salvo size is monotonic non-decreasing from one salvo to the next. We extend the single-threat scenario to the more realistic sequential attack and study the combined effect of HD and SM against aerial attacks.

The references in the previous paragraph focus on HD. There is also work on SM in a variety of research domains. One strand examines effective implementation of SM from a physics and geometric perspective. For example, where to place decoys and what direction to aim jammers. See Sun et al. (2021) and Wu et al. (2024) for examples. Certain SMs affect defensive interceptors. For example, chaff or jamming might cause interceptors to veer from their intended course. Blodgett et al. (2001, 2002), Sui et al. (2009) examine tactics to ensure that there is a minimum amount of deconfliction between HD and SM.

Our work examines how the defender can bolster its use of HD with SM. Much of the operations research work that studies SM analyzes the flip scenario: how an attacker can utilize decoys, chaff, and jamming to enhance its missile attack. These penetration aids (Critchlow and Williams 1982) are meant to decrease the ability of Blue's defensive measures to kill the incoming threats. Examples include Gramann et al. (1963), Wilkening (2000), Zhai et al. (2016).

There are several works that, like us, examine defensive uses of SM. Engelbrecht (1997) and Popa et al. (2018) take a simulation approach to analyze how jamming and chaff can be used to defend against incoming missile threats. Moore (2002) utilizes genetic algorithms to examine how an aircraft should execute turns, thrust changes, and chaffs and flares to avoid anti-aircraft missiles. Carr et al. (2016) uses simulation to evaluate the effectiveness of using a decoy unmanned aerial vehicle to protect a high-value asset. Maskery et al. (2007) formulates a game theory model to develop a decentralized missile defense plan. Each ship can deploy HD and/or SM. The main focus of Maskery et al. (2007) is on centralized vs. decentralized planning; the distinction between HD and SM is a secondary consideration. Our approach differs from these papers in that we develop analytic models that more readily generate insight into how SM improves missile defense.

This work makes two main contributions

1. Provides a framework for decision makers to evaluate cost vs. benefit tradeoffs for different firing policies for HD.

2. Incorporates SM into a firing theory model of missile defense. To the best of our knowledge, this has not been done before.

This work is extremely important given current and potentially future world events. States (Iran, Russia, Ukraine, Israel, United States) and non-state actors (Hamas, Hezbollah, Houthis) have the capabilities and desire to fire rockets, missiles, and drones at their enemies. In April 2024, Iran launched 300 drones and missiles (primarily cruise missiles) at Israel in a single day (Lendon 2024). This was followed by an attack of higher-velocity threats in October 2024, consisting of 200 ballistic and hypersonic missiles (George 2024). The Houthis have severely restricted shipping in the Red Sea due to their missile attacks (Hjelmgaard and Brook 2024). The United States spent over 1 billion USD in munitions in 2023–2024 executing missile defense missions (Williams 2024). Given that HD interceptors often cost several orders of magnitude more than their targets (Seligman and Berg 2024), it is imperative to develop more efficient HD firing policies and incorporate SM into an overall CD strategy. Our paper directly address these issues.

The paper is organized as follows: Section 2 provides more details on the operational setting, Section 3 presents the HD model, and Section 4 adds SM to produce the CD model. The Online Appendix contains a list of acronyms and notation as well as all proofs.

2. Operational Setting

A Blue defender is subject to a Red attack by a sequence of W homogeneous threats. Treating threats as homogeneous is reasonable because each type of threat is typically countered by a specific type of HD–for example, the Arrow system targets ballistic missiles, while Iron Dome intercepts rockets.

Blue has, at the beginning of the battle, an initial inventory of M interceptors for HD. Each interceptor can, independently, intercept a threat with probability p. Blue also has a SM capability, which we assume is an infinite resource (e.g., jamming or directed-energy weapon). We defer discussion of the probability SM neutralizes a threat until Section 4

The defense against a certain threat is time-driven by the *salvo engagement opportunities* (SEO), which are the potential SLS cycles. There are K SEOs. At each SEO, either a salvo of one or more interceptors is fired (HD), or the SM is activated, or both modes of defense are exercised. No action is taken if the threat is evidently killed in a previous SEO. That is, the number of *actual* SEOs is bounded by K. A realized SEO is followed by BDA. While BDA for interceptors is perfect, and there is no false-positive errors for SM, there is a false-negative probability that a kill by the SM is missed.

In each SEO, we assume that Blue receives SM BDA *before* firing its HD salvo. If the SM BDA reports a kill, then the salvo is not fired. Furthermore, a correct SM BDA signal may occur in an SEO occurring after the SM achieves a kill. For example, the SM kills the threat in SEO 1, and Blue receives false-negative BDA reports in SEOs 1 and 2 and the correct BDA report in SEO 3. The false-negative probability can change over time. For example, it is reasonable to assume that the false-negative probability decreases as the time from the actual SM kill increases. As time passes, the signal of a successful SM (e.g., threat veering from its expected trajectory) becomes more likely. Finally, we assume that SM can be active after the HD sequence concludes.

Figure 1 presents an event tree illustrating the sequence of actions and outcomes during an engagement of one threat when the number of SEOs is $K = 3$. Blue first executes SM at SEO 1 (node *S1*). If the SM fails to kill the threat, Blue immediately recognizes that the threat is still alive (no SM false positives) and proceeds to fire its first hard salvo (upper *H1* node). Conversely, if SM successfully kills the threat in SEO 1, Blue receives BDA (leftmost *BDA1* node).

If Blue receives accurate BDA confirming the threat's destruction via SM, the engagement with this threat concludes (indicated by the leftmost Blue star). However, if Blue receives a false-negative BDA indicating the threat is still alive, Blue fires its first hard salvo at the already neutralized threat (lower *H1* node). In this scenario, Blue may achieve both an SM-kill (e.g., disabling the threat's command and control system) and an HD-kill (e.g., physically destroying the threat).

If SM kills the threat in SEO 1 with false-negative BDA, and HD fails to destroy the threat in SEO 1, the engagement proceeds to SEO 2 along the bottom row of Figure 1. After a delay, Blue may eventually receive the correct BDA confirming the SM-kill in SEO 1 (the lowest *BDA2*, *BDA3*, *BDA4* nodes).

If neither an SM-kill nor an HD-kill is achieved in SEO 1, the engagement progresses to SEO 2 with the next round of Blue SM (node *S2*). This cycle continues in a similar pattern until the final SEO. After the last SEO, Blue has a final opportunity to execute SM and obtain BDA, as illustrated in the last column of Figure 1.

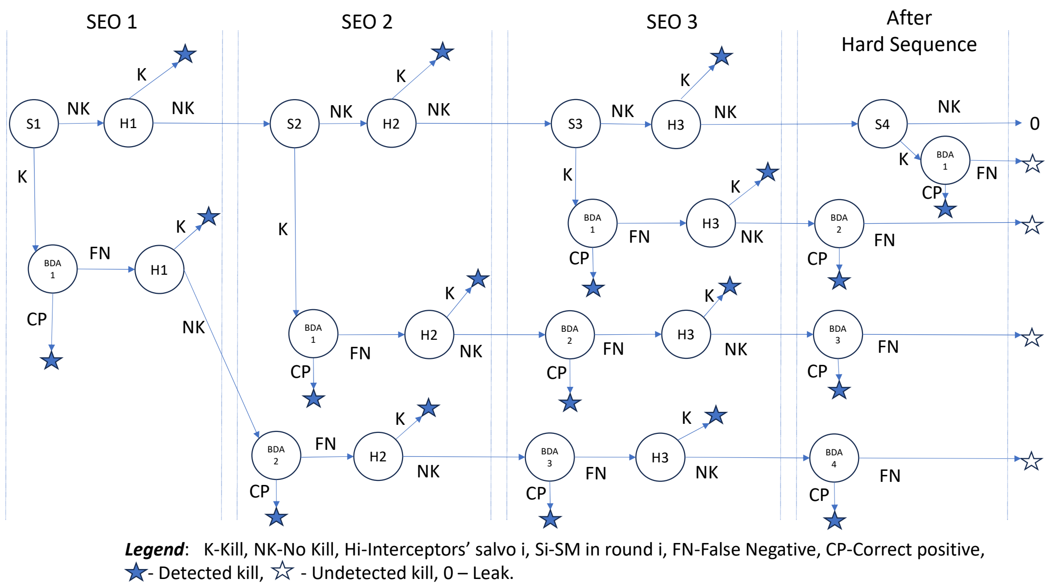

Figure 1 Event tree describing the engagement of one threat for $K = 3$ SEOs. Note that SM can be active both prior to the initiation of the HD sequence (S1 occurs before H1) and after the HD sequence concludes (S4 occurs after H3).

As mentioned above, The *actual* number of SEOs is between 1, if the threat is killed, and correctly identified as such, in the first SEO, and K, if there were no kill indications during the first $K-1$ SEOs. An *engagement tactic* for HD is a firing sequence $\mathbf{s} = (s_1, s_2, \ldots, s_K)$ where s_k is the number of interceptors launched in salvo k. Obviously, the firing sequence against a certain threat terminates at SEO k^* if the "look" following this SEO indicates that the threat has been killed. The firing sequence depends upon the type of threat. Thus, given a barrage of homogeneous threats, (e.g., ballistic missiles), we assume the same firing sequence against every threat.

Recall we consider two MOEs: Expected number of leakers and expected number of interceptors expended. We say that a firing sequence $\mathbf{s} = (s_1, s_2, \ldots, s_K)$ is *efficient* for a fixed K if there is no firing sequence that results in smaller expected number of leakers AND a smaller expected number of interceptors expended. All efficient firing sequences generate an *efficient frontier*. We will use efficient frontiers to evaluate firing sequences. We start by modeling the HD case in Section 3 as an introduction to the more involved CD case in Section 4.

3. Hard Defense (HD)

We first summarize the notation introduced in Section 2

- W: number of threats launched by Red. W is a random variable with finite support : $W \in [W_{\min}, W_{\max}]$.
- M: Blue's initial inventory of interceptors.

- p: single-shot probability of kill (SSPK) of one Blue interceptor against one Red threat
- K: number of SEOs
- $\mathbf{s} = (s_1, s_2, \ldots, s_K)$: Blue's planned firing sequence of interceptors-fired-per-threat. Every threat faces the same firing sequence.

The total number of interceptors fired when a firing sequence executes in full (in all K SEOs) is an important quantity, which we denote as N:

$$N \equiv \sum_{k=1}^{K} s_k. \tag{1}$$

The inputs to the model are W, M, p, and K. The output is a list of efficient sequences $\mathbf{s}$ (and corresponding N) that generate small values of the MOEs.

We next define $X, X \leq N$, as the number of interceptors fired against a certain threat. This random variable has probability mass function $x(i) \equiv \mathbf{P}[X = i]$. $x(i)$ is positive for only K values, $\sum_{\ell=1}^{k} s_\ell$ for $k = 1, 2, \ldots, K$, corresponding to the K possible SEOs. There are two scenarios to consider when defining $x(i)$.

1. Blue fires $k < K$ salvos. This occurs if Blue fails on the first $k-1$ SEOs (with probability $(1-p)^{\sum_{\ell=1}^{k-1} s_\ell}$) and succeeds on SEO k (with probability $(1-(1-p)^{s_k})$). In this case, we have

$$x(i) = (1-(1-p)^{s_k})\,(1-p)^{\sum_{\ell=1}^{k-1} s_\ell}, \tag{2}$$

for $i = \sum_{\ell=1}^{k} s_\ell$, corresponding to some $k = 1, 2, \ldots, K-1$.

2. Blue fires all K salvos (and N interceptors). This occurs when Blue fails on the first $K-1$ SEOs, regardless of the outcome of SEO K. In this case :

$$x(i) = (1-p)^{\sum_{\ell=1}^{K-1} s_\ell}, \qquad \text{if } i = N = \sum_{\ell=1}^{K} s_\ell. \tag{3}$$

The single-threat leak probability and expected number of interceptors-fired-per-threat play a main role in our analysis:

$$\mathbf{P}[\text{Leak}] = (1-p)^N \tag{4}$$

$$\begin{aligned}\mathbf{E}[X] &= \mathbf{E}[\text{Interceptors-per-threat}] \\ &= \sum_{i=1}^{N} i x(i) = \sum_{k=1}^{K} s_k (1-p)^{\sum_{\ell=1}^{k-1} s_\ell}.\end{aligned} \tag{5}$$

We next derive expressions for our two MOEs. The first MOE is the expected number of leakers, which is a measure for the total damage Red inflicts on Blue with its offensive barrage. We compute the expected number of leakers via a recursive formulation. $L(w, m)$ is the expected number of future leakers when Blue has m interceptors remaining in inventory and is currently about to engage threat number w. Index $w = 1$ corresponds to the first threat. Consequently, our desired MOE is $L(1, M)$. To ensure the recursion converges, we assume that W has a finite support: $W \in [W_{\min}, W_{\max}]$.

First consider the situation where $m < N$. In this situation, Blue cannot execute its entire firing sequence $\mathbf{s}$ against the current threat. In the extreme case, where Blue has expended all of its interceptors ($m = 0$), Blue cannot engage any additional threats. We use the tilde notation ($\tilde{\mathbf{s}} = (\tilde{s}_1, \tilde{s}_2, \ldots, \tilde{s}_K)$, $\tilde{N}$, $\tilde{X}$, $\tilde{x}(i)$) to denote the actual firing sequence variables used against the current threat in the $m < N$ situation. While there are several reasonable ways to generate $\tilde{\mathbf{s}}$, we assume, for simplicity and consistency with $\mathbf{s}$, a "truncation from the end" policy. That is, Blue attempts to use $\mathbf{s}$ until it runs out of interceptors:

$$\tilde{s}_k = \min\left(s_k, m - \sum_{\ell=1}^{k-1} \tilde{s}_\ell\right). \tag{6}$$

Once we have $\tilde{\mathbf{s}}$, we compute $\tilde{N}$, $\tilde{x}(i)$, $\mathbf{E}[\tilde{X}]$ using (1)–(5), with s_k replaced by $\tilde{s}_k$. We stress that $\tilde{\mathbf{s}} = \mathbf{s}$ except in the special case when Blue is nearly out of interceptors: $m < N$.

We now present the recursive formula for the expected number of leakers. The recursion consists of two parts: the likelihood that the current threat leaks plus the expected additional leakers from future threats. The number of threats, W, is a random variable, and Blue does not know its realization in advance. Thus, the term for future leakers must account for the possibility that Blue may face no further threats.

$$
\begin{aligned}
L(w,m) = {} & (1-p)^{\tilde{N}} \\
& + \frac{\mathbf{P}[W \geq w+1]}{\mathbf{P}[W \geq w]} \Bigg((1-p)^{\sum_{\ell=1}^{K-1} \tilde{s}_\ell} \\
& \qquad \times L\big(w+1, m-\tilde{N}\big) \\
& \quad + \sum_{k=1}^{K-1} (1-p)^{\sum_{\ell=1}^{k-1} \tilde{s}_\ell} \big(1-(1-p)^{\tilde{s}_k}\big) \\
& \qquad \times L\Big(w+1, m - \sum_{\ell=1}^{k} \tilde{s}_\ell\Big)\Bigg)
\end{aligned}
\tag{7}
$$

The first line in (7) is the leak probability of the current threat, and lines 2 through 5 of (7) compute the expected number of future leakers recursively. We multiply lines 2 through 5 by the conditional probability there will be additional threats. Lines 2 and 3 of (7) occur when Blue fails on the first $K-1$ SEOs and executes the entire firing sequence. Lines 4 and 5 of (7) correspond to the situation where the threat is killed on salvo $k < K$. If $w < W_{min}$, then Blue faces future threats with certainty: $\mathbf{P}[W \geq w+1] = \mathbf{P}[W \geq w] = 1$ and the conditional probability factor in line 2 of (7) is 1. If $w = W_{max}$, $\mathbf{P}[W \geq w+1] = 0$, Blue faces no future threats, and the recursion in (7) terminates with $L(W_{max}, m) = (1-p)^{\tilde{N}}$

Similarly we define a recursion for the number of interceptors fired: $F(w,m)$ is the expected number of additional interceptors fired when Blue has m interceptors remaining in inventory and is currently about to engage threat number w.

$$
\begin{aligned}
F(w,m) = \mathbf{E}[\tilde{X}]& \\
+ \frac{\mathbf{P}[W \geq w+1]}{\mathbf{P}[W \geq w]}\Bigg(&(1-p)^{\sum_{\ell=1}^{K-1}\tilde{s}_\ell} \\
&\times F\big(w+1, m-\tilde{N}\big) \\
&+ \sum_{k=1}^{K-1}(1-p)^{\sum_{\ell=1}^{k-1}\tilde{s}_\ell}\big(1-(1-p)^{\tilde{s}_k}\big) \\
&\times F\Big(w+1, m-\sum_{\ell=1}^{k}\tilde{s}_\ell\Big)\Bigg)
\end{aligned} \tag{8}
$$

The first line in (8) is the expected number of interceptors fired at the current threat. Lines 2 through 5 of (8) compute the recursive portion in the same fashion as (7)

While there is little analysis we can directly perform on $L(w,m)$ and $F(w,m)$, we can easily compute the recursions in (7)–(8) numerically to obtain values for the two MOEs.

3.1. HD Approximation

To generate insight, we approximate our MOEs by considering two scales of Red attack: a *small*-W case and a *large*-W case. In the small-W case, Blue still has remaining interceptors at the end of the engagement. In the large-W case, Blue depletes its interceptor inventory before Red's assault finishes.

If we assume Blue has an infinite supply of interceptors ($M = \infty$) and fires X_w interceptors against the w-th threat, then $\{X_w, w \geq 1\}$ is a sequence of independent and identically distributed (IID) random variables. Consequently, Blue fires $\mathbf{E}[W]\mathbf{E}[X]$ interceptors in expectation against all the threats in the $M = \infty$ setting, as W and X_w are independent. Returning to the finite M case, we compare $\mathbf{E}[W]\mathbf{E}[X]$ to M to determine when to use the small-W vs. large-W approximation. When $\mathbf{E}[W]\mathbf{E}[X] < M$, we use the small-$W$ approximation, which we denote with subscript AS. Otherwise we use the large-W approximation and subscript AL.

In the small-W case, when $\mathbf{E}[W] < \frac{M}{\mathbf{E}[X]}$

$$\begin{aligned}
\mathbf{E}[\text{leakers}] &\approx \mathbf{E}[\text{leakers}_{AS}] \\
&\equiv \mathbf{E}[W] \times (1-p)^N && (9) \\
\mathbf{E}[\text{interceptors fired}] &\approx \mathbf{E}[\text{interceptors fired}_{AS}] \\
&\equiv \mathbf{E}[W]\mathbf{E}[X]. && (10)
\end{aligned}$$

We formally derive these expressions in the proof for Theorem 1 below. Informally, if W is small and deterministic, the number of leakers follows a binomial distribution with W trials and success probability $(1-p)^N$. Note, that the expected number of leakers in (9) only depends upon the maximum number of interceptors fired at a single threat, N, and not on the exact nature of the firing sequence. The number of interceptors fired in (10) does depend upon the specific structure of the firing sequence. Consequently, for a fixed N, Blue chooses a firing sequence to minimize $\mathbf{E}[X]$.

For large-W, $\mathbf{E}[W] > \frac{M}{\mathbf{E}[X]}$

$$\begin{aligned}
\mathbf{E}[\text{leakers}] &\approx \mathbf{E}[\text{leakers}_{AL}] \\
&\equiv \mathbf{E}[W] - \mathbf{E}[\text{Kills}] && (11) \\
\mathbf{E}[\text{interceptors fired}] &\approx \mathbf{E}[\text{interceptors fired}_{AL}] \\
&\equiv M, && (12)
\end{aligned}$$

where $\mathbf{E}[\text{Kills}]$ is the expected number of threats that Blue intercepts. By definition, in the large-W case Blue fires all available interceptors. Hence, Blue chooses a firing sequence that maximizes the number of intercepted threats. Analyzing $\mathbf{E}[\text{Kills}]$ is complicated because the number of interceptors fired against the last threat engaged before Blue runs out of interceptors has a different distribution than the number fired against earlier threats. Since we do not need to derive a general expression for $\mathbf{E}[\text{Kills}]$ for any of our analysis, we defer further discussion of computing $\mathbf{E}[\text{Kills}]$ to Online Appendix C.2.

We base the above approximations on the extreme cases $W \ll M$ and $W \gg M$. In these extreme cases, the approximations are exact. Theorems 1 and 2 provide the conditions.

THEOREM 1. *If* $W_{max}N \leq M$ *then*

$$\begin{aligned} \mathbf{E}[leakers] =& \mathbf{E}[leakers_{AS}] \\ \equiv& \mathbf{E}[W] \times (1-p)^N \end{aligned} \tag{13}$$

$$\begin{aligned} \mathbf{E}[interceptors\ fired] =& \mathbf{E}[interceptors\ fired_{AS}] \\ \equiv& \mathbf{E}[W]\mathbf{E}[X]. \end{aligned} \tag{14}$$

$W_{max}N$ is an upper-bound on the number of interceptors fired, so the condition in Theorem 1 ensures Blue cannot run out of interceptors.

Theorem 2 covers the other extreme: when $W \gg M$. Specifically Theorem 2 is the case where the number of interceptors M is evidently smaller than the number of threats W. The AL approximation is also exact in this situation because Blue runs out of interceptors with certainty. As mentioned above, we discuss $\mathbf{E}[\text{Kills}]$ in Online Appendix C.2 and provide there a simple approximation in the large-W case.

THEOREM 2. *If* $W_{min} > M$ *then*

$$\begin{aligned} \mathbf{E}[leakers] =& \mathbf{E}[leakers_{AL}] \\ \equiv& \mathbf{E}[W] - \mathbf{E}[Kills] \end{aligned} \tag{15}$$

$$\begin{aligned} \mathbf{E}[interceptors\ fired] =& \mathbf{E}[interceptors\ fired_{AL}] \\ \equiv& M. \end{aligned} \tag{16}$$

The proofs for Theorems 1 and 2 appear in Online Appendix C.1.

Note that the distinction between the small-W and large-W case ($\mathbf{E}[W]$ vs. $\frac{M}{\mathbf{E}[X]}$) depends upon the firing sequence through $\mathbf{E}[X]$. Therefore, for a fixed $\mathbf{E}[W]$, one firing sequence might lie in the small-W regime but another sequence might lie in the large-W regime.

Theorems 1–2 provide theoretical justification for using the approximations presented in this section to further analyze the system. In Online Appendix C.3 we illustrate numerically that these approximations perform very well in practice.

3.2. Evaluating HD Efficiency

Using the AS and AL approximations in 3.1, we examine which firing sequences Blue should consider to minimize the two MOEs.

In our analysis, we fix the number of SEOs at K and evaluate various firing sequences across different values of N. Instead of identifying a single optimal firing sequence and corresponding optimal N, we reveal a set of efficient sequences, each associated with a distinct value of N. Recall a sequence is efficient for fixed K if no other firing sequence generates both a smaller expected number of leakers and a smaller expected number of interceptors fired.

3.2.1. Large-W For large-W Blue expends all M interceptors, so efficiency corresponds to minimizing the expected number of leakers. From (11), this is equivalent to Blue maximizing the number of threats killed. Theorem 3 specifies the structure of an efficient firing sequence.

THEOREM 3. *In the large-W case, Blue maximizes the expected number of threats killed by executing a sequence that fires no more than one interceptor in every salvo. That is* $s_k \in \{0,1\}, \forall k = 1,2,\ldots K$.

We refer to the sequences in Theorem 3 as *one-interceptor-per-salvo* sequences. Theorem 3 follows by using an accounting argument. For any firing sequence, an upper bound on the number of kills is the number of successful interceptors. A one-interceptor-per-salvo sequence achieves this upper bound: each successful interceptor corresponds to a kill.

There are a number of one-interceptor-per-salvo sequences. For example, for $K = 3$, sequences $\mathbf{s} = (0,0,1)$, $\mathbf{s} = (0,1,1)$, and $\mathbf{s} = (1,1,1)$, corresponding to $N = 1,2,3$, respectively, are all equally efficient. Furthermore, if $K > N$ it does not matter which SEOs Blue chooses to utilize: $(0,1,1)$, $(1,0,1)$ and $(1,1,0)$ are equivalent.

We can compute $\mathbf{E}[\text{Kills}]$ explicitly for these one-interceptor-per-salvo sequences:

THEOREM 4. *When Blue executes a one-interceptor-per-salvo sequence such that* $s_k \in \{0,1\}, \forall k = 1,2,\ldots K$*, then*

$$\mathbf{E}[Kills] = Mp. \tag{17}$$

Theorem 4 follows by defining the outcome of each interceptor as an IID sequence of Bernoulli random variables with success probability p. Each successful Bernoulli has a one-to-one correspondence with a Red kill. Hence the number of kills follows a Binomial distribution with M

trials and success probability p. This approach does not work with other firing sequences because multiple Bernoulli successes can correspond to one kill.

3.2.2. Small-W Equations (9)–(10) manifest a tradeoff for Blue who wants to fire many interceptors (large N) to drive down the leak probability, but also wishes to limit the number of interceptors fired (small $\mathbf{E}[X]$). A small $\mathbf{E}[X]$ is also beneficial because it makes the preferred small-W case more likely than the large-W case.

There is no analog to Theorem 3 in the small-W case as the chosen firing sequence will depend on Blue's acceptable risk for leakers. However, there is a dominance relation among sequences that identifies efficient sequences. The maximum number of interceptors-fired-per-threat, N, determines the expected number of leakers (see (9)). If two firing sequences have the same N (same risk for leakers), then dominance is dictated by the expected number of interceptors fired at each threat (see (10)). We highlight this notion of dominance in the following remark to define *efficiency*, as we utilize it throughout this section

REMARK 1. If firing sequences $\mathbf{s^A}$ and $\mathbf{s^B}$ both have the same number of total interceptors fired ($N_A = N_B$), then $\mathbf{s^A}$ is said to be more *efficient* than $\mathbf{s^B}$ if and only if $\mathbf{E}[X_A] < \mathbf{E}[X_B]$

For a given number of SEOs K, we generate the efficient frontier of the firing sequences as we vary the leak probability, which is, for a fixed p, equivalent to varying N. For a fixed N, a firing sequence that generates the smallest $\mathbf{E}[X]$ is efficient and lies on the efficient frontier. Blue only needs to consider firing sequences that lie on the efficient frontier, as all other sequences are dominated by an efficient firing sequence. Next we present several theoretical results about the properties of efficient firing sequences. We numerically illustrate these efficient frontiers in Section 3.3.

For a fixed K and N, the efficient firing sequence may not be unique. See Online Appendix B.1 for an example. However, if we perturb the the scenario input data (e.g., p) slightly, we can break ties and force a unique solution. Consequently, we treat the solution as unique and present the analysis accordingly. See Online Appendix B.1 for further discussion on uniqueness.

The first result specifies that Blue prefers to spread its interceptors out over as many salvos as possible

THEOREM 5. *If $K \leq N$, then any firing sequence that does not utilize all available SEOs (i.e., $s_k = 0$ for some $k = 1, 2, \ldots K$) is inefficient.*

The proof for Theorem 5 appears in Online Appendix G. Theorem 5 implies Blue prefers K to be as large as possible. Theorem 5 also implies the following result

COROLLARY 1. *Assume the total number of interceptors in a firing sequence is fixed at N. The firing sequence that fires exactly one interceptor in each salvo for $K = N$, $\mathbf{s}^* = (1, 1, 1, \ldots, 1)$, is more efficient than all other firing sequences with N interceptors.*

Note that the dominance of the one-interceptor-per-salvo sequence for the small-W case in Corollary 1 resembles the optimality of one-interceptor-per-salvo sequences for the large-W case in Theorem 3. If K is small, the sequence $\mathbf{s}^*$ in Corollary 1 might not generate a small enough leak probability to satisfy Blue. However, the similarity between Corollary 1 and Theorem 3 suggests that $\mathbf{s}^*$ in Corollary 1 will often be a robust choice that performs well regardless of the value of W.

The previous results stress the importance of utilizing as many SEOs as possible. The next several results fix the number of SEOs at K and examine the structure of efficient firing sequences. Intuitively Blue wants to conserve its inventory by firing smaller salvos earlier and larger salvos later. The following theorem formalizes this monotonicity property reported earlier in Kalyanam and Clarkson (2021).

THEOREM 6. *If $\mathbf{s}$ is the efficient sequence for fixed N and K, then*

$$s_k \leq s_{k+1} \quad \forall 1 \leq k \leq K - 1. \tag{18}$$

The proof for Theorem 6 appears in Online Appendix H. The following corollary follows immediately from Theorem 5 and Theorem 6.

COROLLARY 2. *If $N = K + 1$ then the efficient firing sequence is $\mathbf{s}^* = (1, 1, 1, \ldots, 2)$*

Corollary 2 suggests a heuristic: fire one interceptor in the first $K - 1$ SEOs, and then fire the remaining interceptors in the final SEO. The intuition behind this "extreme" sequence is that Blue conserves interceptors if it can avoid firing the last salvo. The next theorem specifies that this sequence is efficient when p is large enough.

THEOREM 7. *For fixed N and K, $N > K$, the efficient sequence is $(1, 1, \ldots, N - K + 1)$ if and only if*

$$p \geq \frac{N - K - 1}{N - K}. \tag{19}$$

The proof for Theorem 7 appears in Online Appendix I.

3.3. Numerical Analysis of the HD Case

In this Section, we provide some scenario-dependent operational insights regarding efficient firing sequences. Recall that the objective is to minimize both MOEs: expected number of leakers and expected number of expended interceptors. The inputs to the model are p, K, M, and W. There are several observations that determine the scope of the operational analysis and corresponding ranges of valid input.

- Anti-missile interceptors generally have high success rates p, which is typically higher than 80% (Landau and Bermant 2014, Eyal Pecht and Weingold 2013). Thus, for the analysis that follows we consider just two values for our SSPK (p): a *pessimistic* kill probability, 0.7, and an *optimistic* kill probability, 0.9.
- The number of SOEs (K) for short- and mid-range threats is relatively small; the time-window for SLS sequences is quite limited. Thus we assume that, within the limited window of engagement, the number of SEOs ranges between $K = 1$ (no BDA opportunity) and $K = 3$ (two BDA opportunities).
- We set the initial inventory to $M = 100$ throughout, which is consistent with the typical inventory of interceptors on board United States Navy battle ships (Stöhs 2021).
- We focus on the small-W setting for two reasons. From an analysis perspective, the large-W is not very interesting. Blue depletes its inventory and should use a one-interceptor-per-salvo policy throughout. From an operational standpoint, the small-W case is more realistic. Red often will fire fewer than 10 threats in a time window of interest (USCENTCOM 2024, Jewish Virtual Library 2024, Sewell and Mroue 2024, Azhari et al. 2024).
- N is not a direct model input. As discussed in Section 3.2.2, for a fixed K, there is a unique efficient sequence for each N. In this section, we vary N from 1 to 5 to generate the efficient frontiers. There is no reason to consider higher values of N because $p = 0.7$ and $N = 5$ correspond to a 0.998 single-threat kill probability, which is likely more than sufficient. In practice, N would be an input selected to meet a specified interception probability threshold and reflects a commander's risk tolerance for leakers.

In Section 3.3.1, we present the efficient frontiers introduced in Section 3.2.2. Then in Section 3.3.2, we illustrate how Red could use our framework to determine the launch range of their threats. Note all figures in the paper use the recursions in (7) and (8) to compute the MOEs exactly. Furthermore, the efficient frontiers are determined numerically rather than derived from the theoretical results in Section 3.2.2. Nonetheless, the insights from the approximation results in Section 3.2.2 apply, because we focus on the small-W setting.

3.3.1. Efficient Frontier Figure 2 plots the efficient frontier in the Expected Number of Interceptors Fired vs. Expected Number of Leakers space. The two figures are for $p = 0.7$ (Figure 2a) and $p = 0.9$ (Figure 2b). We vary the number of SEOs between $K = 1$ (solid line), $K = 2$ (dashed line), and $K = 3$ (dotted line). We set W to $\sim U[10,20]$ and fix $M = 100$. Note the figures use log-scale for the x-axis.

Recall from Section 3.2.2 that we construct the efficient frontier by finding the firing sequence that minimizes $\mathbf{E}[X]$ for each fixed N. The results in Section 3.2.2 determine the efficient frontier. For example, let us focus on $K = 3$ (dotted line). The sequences (0,1,1) (c) and (1,1,1) (f) follow from Theorem 5, (1,1,2) (j) follows from Corollary 2, and (1,1,3) (o) follows from Theorem 7 (because $p = 0.7 > 0.5$).

Arguably, Blue prefers firing sequences closer to the origin, which represent fewer expected leakers and fewer expected expenditure of interceptors. The $N = 1$ sequence (a) is on the far right and bottom in each figure, representing the highest expected number of leakers, and lowest expected number of interceptors. In the case of $N = 5$, the efficient firing sequences are (1,1,3) (o) for $K = 3$, (1,4) (l) for $K = 2$, and (5) (l) for $K = 1$ (note (5) expends 75 interceptors, so it does not appear on the figure). The sequences corresponding to (i), (m), and (n) are inefficient for $p \in \{0.7, 0.9\}$. Throughout the paper, we use an asterisk (*) to denote any sequence that does not appear on any efficient frontier in a given figure.

Obviously Blue prefers larger p and larger K. There is a significant difference between $K = 1$ and $K = 2$ and a much smaller difference between $K = 2$ and $K = 3$, especially for $p = 0.9$. For $p = 0.9$ it is likely that Blue kills the threat before firing a third salvo. This suggests that Blue should invest in command and control technology (e.g., radar) to obtain one BDA look, but further investments to obtain additional looks might not be cost-effective.

For a fixed K, Blue chooses a firing sequence on the corresponding efficient frontier based on the trade-off between the damaging effect of the threat and the cost of the interceptors. In cases where even one leak causes significant damage, Blue might choose $N = 4$ or 5 when $p = 0.7$. If $p = 0.9$ and $K = 2$ or $K = 3$, the efficient frontiers are nearly flat, especially for $K = 3$. This suggests that Blue should choose $N = 5$ ((o) or (l) in Figure 2b). With such a high SSPK, Blue rarely fires the later salvos.

Figure 2 illustrates the small-W case. In Online Appendix D, we show how the frontiers evolve as $\mathbf{E}[W]$ increases, transitioning from the small-W case to the large-W case.

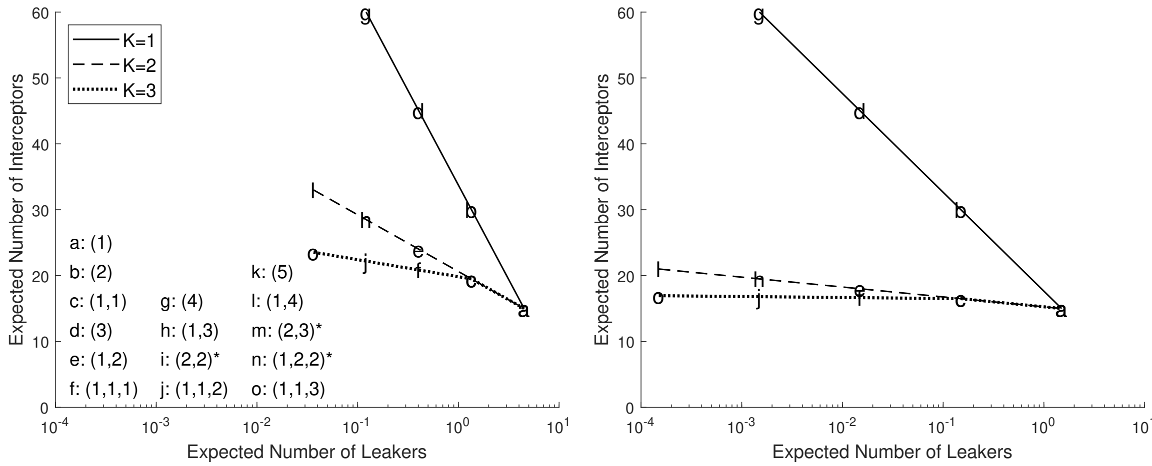

(a) $p = 0.7$ (b) $p = 0.9$

Figure 2 Efficient Frontier in Expected Number of Interceptors vs Expected Number of Leakers space for various values of K and p. x-axis is on a log scale. In the legend, sequences are marked with an asterisk (*) if they are inefficient for all combinations of K and p.

3.3.2. Red Decision: Threats W vs Salvo Opportunities K In this section we study a hypothetical Red decision: at what range should Red engage Blue. This problem is manifested in the trade-off between the number of threats (W) Red can launch and the number of SEOs (K) Blue has available. If Red chooses to launch closer to the Blue target, Blue will have fewer SEOs (smaller K). However, the closer Red launches to Blue, the easier it is for Blue to detect the launch location and utilize counter measures, which may reduce the number of threats Red can launch (smaller W). Note we are not modeling the mechanism by which Blue reduces W (e.g., a Blue drone strike against a Red launch site), we just assume in this section that W depends upon K.

In this section, Red is the decision maker and considers three possible ranges for launching its attack: short-range, medium-range, and long-range. Each range corresponds to a different value of Blue SEOs, K. When Red fires at short-range, Blue has $K = 1$ SEOs, whereas at long range Blue has $K = 3$ SEOs. To simplify the presentation, Red selects its launch range by choosing K.

Once Red selects K, Red fires $W_K \sim U[\mu_K - 5, \mu_K + 5]$ threats, where $\mathbf{E}[W_K] = \mu_K$. Here, the number of threats W_K explicitly depends upon the range via K, with μ_K increasing in K. This increase reflects that Red can avoid detection longer at greater ranges and thus launch a larger number of threats. μ_K is an input to the model, and we denote $\mu = (\mu_1, \mu_2, \mu_3)$ as the vector containing the mean number of threats by range.

Figure 3 plots the efficient frontier for $K = 1$ (solid line), $K = 2$ (dashed line), and $K = 3$ (dotted line) and for $p \in \{0.7, 0.9\}$. We consider two different mean threat vectors μ: $\mu = (12, 15, 18)$ (top row of Figure 3) and $(6, 12, 18)$ (bottom row). Each line corresponds to a different number of SEOs, K, and a different threat distribution, W_K.

Figure 3 is similar to Figure 2. The key difference is that the distribution of W is fixed in Figure 2 for all values of K, but the threat distribution varies with K in Figure 3. For example, in the top row of Figure 3, the $K = 1$ curves correspond to $W_1 \sim U[7, 17]$ threats, and the $K = 3$ curves correspond to $W_3 \sim U[13, 23]$ threats.

Red prefers an outcome that leads to more leakers (more effective attack) and more interceptors launched (depleting Blue's arsenal). Red is unlikely to choose the middle option of $K = 2$ as the $K = 3$ option dominates (i.e., is to the northeast) or nearly dominates the $K = 2$ frontier. Only in Figure 3c does Red have a clearly dominant option. In this case where $p = 0.7$ and $\mu = (6, 12, 18)$, Red should choose the $K = 3$ salvo option and launch an attack with $\mu_3 = 18$ threats. In this case, Blue has a lower SSPK and moving to the $K = 3$ option significantly increases the number of threats Red can fire. For the $\mu = (12, 15, 18)$ case, the $K = 1$ is a fairly robust option as it lies to the northeast of the other two frontiers over most of the domain. In this case, Red prefers the higher attrition to Blue from using 1-salvo sequences, even though Red launches slightly fewer threats. Finally in figure 3d, Red has no good options. With a SSPK of $p = 0.9$ and a limited number of threats, Blue can fairly successfully counter any Red plan.

There are other types of similar analysis our model can support. For example, Blue could examine a trade-off between K and p: should Blue invest in command and control to improve BDA and increase the number of SEOs K or should Blue invest in improved accuracy and lethality and thus increase the SSPK p.

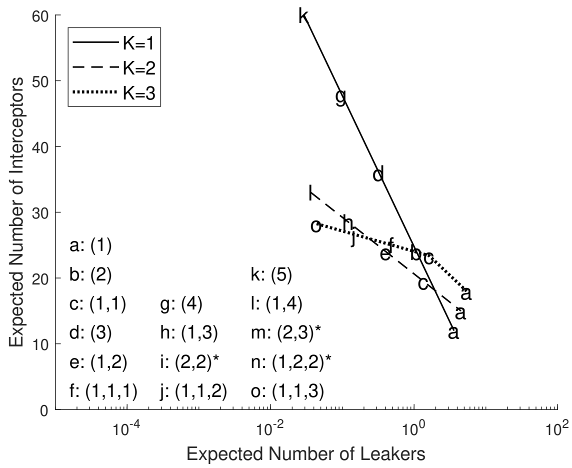

(a) $p = 0.7$, and $\mu = (12, 15, 18)$

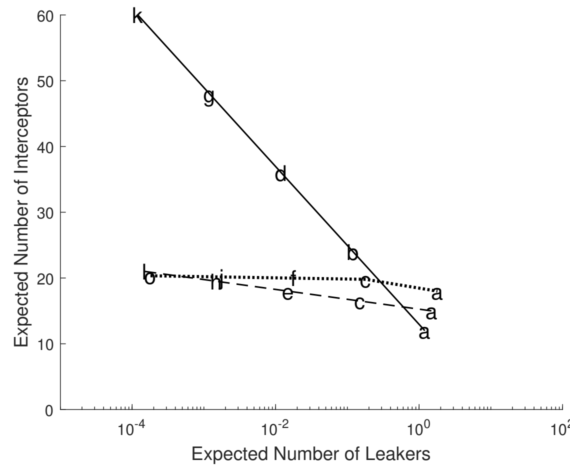

(b) $p = 0.9$, and $\mu = (12, 15, 18)$

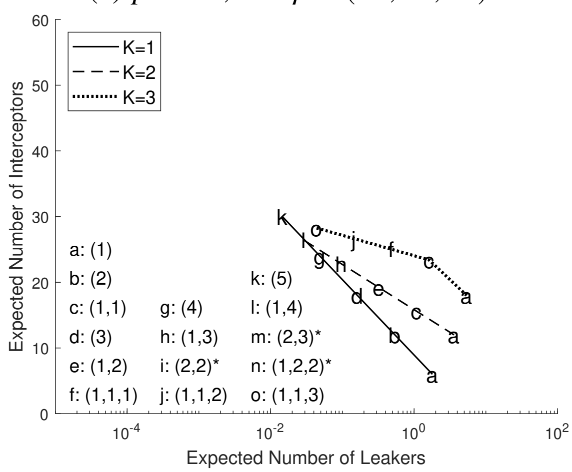

(c) $p = 0.7$, and $\mu = (6, 12, 18)$

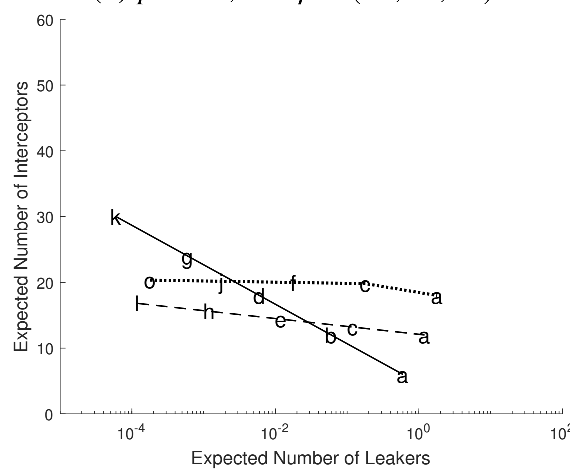

(d) $p = 0.9$, and $\mu = (6, 12, 18)$

Figure 3 Expected Number of Interceptors vs Expected Number of Leakers when Blue has K SEOs against W_K threats. x-axis is on a log scale. In the legend, sequences are marked with an asterisk (*) if they are inefficient for all combinations of K and p.

4. The Combined Defense (CD): Incorporating SM

HD interceptors are very expensive (Williams 2024), and therefore Blue might also utilize SM in conjunction with the interceptors. Possible SMs are jamming, decoys, chaff, and directed-energy weapons (van Hooft and Boswinke 2021, Keller 2021). In this section we formulate a CD model by incorporating SM into the HD scenario from Section 3. As we describe the CD model, the reader may find it useful to refer back to Figure 1.

We assume that Blue has one type of SM, which can be active for any part (including the entire duration) of a threat engagement. The SM can kill the threat at the beginning of each SEO and also after the last SEO. Hence, if there are K SEOs, there are $K+1$ *rounds* where SM can be active. See Figure 1. The SM and the interceptors operate independently. Round k of a SM refers to the period just prior to the HD salvo in SEO k. Hence, round 1 occurs before the first salvo, and round $K+1$ occurs after Blue fires all the HD interceptors in the firing sequence. In each round, the SM can kill the threat. For round $k, k = 1 \dots K$ we assume Blue receives the BDA report just before firing the salvo in SEO k and therefore only fires the salvo if the BDA reports the threat is still alive; in Figure 1, node *Sk* comes before *Hk* in SEO k.

While the HD kills (by interceptors) have perfect BDA, the BDA for SM is imperfect and subject to false-negative error. We assume no false-positive error, because declaring a threat as killed in such scenarios requires strong evidence.

Let θ_k denote the conditional probability that the SM fails to kill the threat in round k given that the threat survived the first $k-1$ rounds.

If Blue kills a threat with SM, there may be a delay until Blue determines the correct BDA. For example, if SM kills the threat in round 1, Blue may not be able to correctly identify the kill until round 3, when, say, the threat evidently veers off course. In Figure 1, this example corresponds to the star below the lowest *BDA3* node. Formally, given an SM kill at round k, the BDA false-negative probability at round $k+j$ is denoted by γ_j for $j = 0, \dots K+1-k$. The false-negative probability γ_j is a conditional probability given that the BDA in previous rounds generated false-negative reports. Arguably, γ_j is monotone non-increasing in j.

Ignoring for a moment the HD kills (by interceptors), let ψ_k denote the probability that the SM's BDA reports that the threat is still alive at the end of round k, just prior to firing the interceptors in SEO k. Such a report could be obtained in two scenarios: (1) the threat is still alive after round k with probability ψ_k^{alive}, or (2) the threat was killed by SM in round j, $j = 1, \dots, k$, but at round k it is falsely reported as being alive. We denote the probability of the second scenario as $\psi_{k,j}^{fn}$. We have:

$$\psi_k = \psi_k^{alive} + \sum_{j=1}^{k} \psi_{k,j}^{fn}, \tag{20}$$

where

$$\psi_k^{alive} = \prod_{i=1}^{k} \theta_i, \tag{21}$$

and

$$\psi_{k,j}^{fn} = \prod_{i=j}^{k} \gamma_{i-j} \times (1-\theta_j) \prod_{i=1}^{j-1} \theta_i. \tag{22}$$

Equation (21) follows directly from the definition of θ_k. The expression for $\psi_{k,j}^{fn}$ in (22) has two terms: (1) the probability the threat is killed in round j: $(1-\theta_j)\prod_{i=1}^{j-1}\theta_i$, and (2) the probability all BDA reports from round j to round k have been false negatives: $\prod_{i=j}^{k}\gamma_{i-j}$.

We summarize the notation for SM below. The notation for the HD interceptors is given in the beginning of Section 3.

- $K+1$: number of rounds of SM
- θ_k: probability threat survives SM at round k; $\theta = (\theta_1, \ldots, \theta_{K+1})$
- γ_j: probability SM BDA reports a false negative j periods after a SM kill occurs; $\gamma = (\gamma_0, \ldots, \gamma_K)$
- ψ_k: probability SM BDA reports "threat is alive" after round k. $\psi = (\psi_1, \ldots, \psi_{K+1})$
- ψ_k^{alive}: probability threat is alive after round k
- $\psi_{k,j}^{fn}$: probability that the threat was killed by SM in round $j (\leq k)$ and was falsely reported to be alive in round k.

The first three parameters in the list above are design inputs while the last three are derived quantities.

We assume that Blue will deploy its SM in every round since there is no associated cost. However, operational constraints may limit SM usage, such as restrictions on range. For rounds in which SM is inactive due to such constraints, we still consider SM as deployed, but we set $\theta_k = 1$.

To modify the recursion equations (7)–(8) to incorporate SM, we first update equations (4)–(5):

$$\mathbf{P}[\text{Leak}] = (1-p)^N \prod_{i=1}^{K+1} \theta_i \tag{23}$$

$$\begin{aligned}\mathbf{E}[X] &= \mathbf{E}[\text{Interceptors-per-threat}] \\ &= \sum_{k=1}^{K} \psi_k s_k (1-p)^{\sum_{\ell=1}^{k-1} s_\ell}.\end{aligned} \tag{24}$$

We next specify when Blue stops firing HD interceptors. Blue terminates its HD firing sequence after SEO k if all three of the following conditions hold:

1. Blue has not scored a HD kill in SEOs $1, 2, \ldots, k-1$ (probability $(1-p)^{\sum_{\ell=1}^{k-1} s_\ell}$)
2. SM BDA reports the threat is "alive" after round k (probability ψ_k)
3. One of the two following conditions hold:
   (a) Blue kills the threat in SEO k with an interceptor (probability $1-(1-p)^{s_k}$)
   (b) Blue fails to score a HD kill in SEO k (probability $(1-p)^{s_k}$) **and** SM BDA reports "kill" in round $k+1$

Conditions 2 and 3b imply that one of the scenarios we must examine is when SM BDA reports "kill" in round $k+1$ after reporting "alive" in round k. We present the probability of this event in the following equation:

$$\begin{aligned}&\mathbf{P}\Big[\text{ SM BDA reports "kill" in round } k+1 \\ &\qquad \Big| \text{ SM BDA reports "alive" in round } k\Big] \\ &= \frac{\psi_k^{alive}}{\psi_k}(1-\theta_{k+1})(1-\gamma_0) \\ &\quad + \sum_{j=1}^{k} \frac{\psi_{k,j}^{fn}}{\psi_k}(1-\gamma_{k-j+1}).\end{aligned} \tag{25}$$

The first term in (25) is when a SM kill occurs in round $k+1$ and the immediate BDA is correct. The summation term in (25) is when a SM kill happens in round $j < k+1$ and round $k+1$ is the first round when the BDA is correct.

We now put all the pieces together and present the recursive equations for the expected number of leakers and interceptors fired. As in the HD model of Section 3, we use the tilde notation $\tilde{\mathbf{s}}$ to account for truncated firing sequences when current inventory m is low. See the discussion around equation (6) for more details about the truncated sequences. We use the following recursion to compute the expected number of leakers.

$$
\begin{aligned}
L(w,m) &= (1-p)^{\tilde{N}} \prod_{i=1}^{K+1} \theta_i \\
&\quad + \frac{\mathbf{P}[W \geq w+1]}{\mathbf{P}[W \geq w]} \Bigg(\psi_K (1-p)^{\sum_{\ell=1}^{K-1} \tilde{s}_\ell} \\
&\qquad \times L\big(w+1, m-\tilde{N}\big) \\
&\qquad + \sum_{k=0}^{K-1} \psi_k (1-p)^{\sum_{\ell=1}^{k-1} \tilde{s}_\ell} \Bigg(\big(1-(1-p)^{\tilde{s}_k}\big) \\
&\qquad\quad + (1-p)^{\tilde{s}_k} \Bigg(\frac{\psi_k^{alive}}{\psi_k} (1-\theta_{k+1})(1-\gamma_0) \\
&\qquad\quad + \sum_{j=1}^{k} \frac{\psi_{k,j}^{fn}}{\psi_k} (1-\gamma_{k-j+1}) \Bigg) \Bigg) \\
&\qquad \times L\Bigg(w+1, m - \sum_{\ell=1}^{k} \tilde{s}_\ell \Bigg) \Bigg)
\end{aligned} \tag{26}
$$

The first line in (26) is the probability the current threat leaks (see equation (23)). The second and third lines of (26) correspond to the special case when Blue fires the entire firing sequence (use $k = K$ in Conditions 1 and 2 in the list below (24)).The fourth line of (26) is when Blue terminates its firing sequence at salvo $k < K$ because of a HD kill (Condition 3a in the above list). Finally lines 5 and 6 of (26) occur when Blue stops firing because SM BDA reports “kill” (see equation (25)).

We next present a similar recursion for the expected number of interceptors fired.

$$\begin{aligned}
F(w,m) &= \mathbf{E}[\tilde{X}] \\
&+ \frac{\mathbf{P}[W \geq w+1]}{\mathbf{P}[W \geq w]} \Bigg(\psi_K (1-p)^{\sum_{\ell=1}^{K-1} \tilde{s}_\ell} \\
&\quad \times F\big(w+1, m-\tilde{N}\big) \\
&\quad + \sum_{k=0}^{K-1} \psi_k (1-p)^{\sum_{\ell=1}^{k-1} \tilde{s}_\ell} \Bigg(\big(1-(1-p)^{\tilde{s}_k}\big) \\
&\quad\quad + (1-p)^{\tilde{s}_k} \Bigg(\frac{\psi_k^{alive}}{\psi_k} (1-\theta_{k+1})(1-\gamma_0) \\
&\quad\quad + \sum_{j=1}^{k} \frac{\psi_{k,j}^{fn}}{\psi_k} (1-\gamma_{k-j+1}) \Bigg) \Bigg) \\
&\quad \times F\Big(w+1, m - \sum_{\ell=1}^{k} \tilde{s}_\ell \Big) \Bigg)
\end{aligned} \tag{27}$$

Solving these recursions numerically is straightforward.

4.1. Approximation

As in Section 3.1, we generate approximations using the same AS and AL subscripts for the small-W and large-W cases, respectively. The underlying logic and expressions for the approximations in this section are very similar to those in Section 3.1.

In the small-W case, when the expected number of threats is relatively small, that is $\mathbf{E}[W] < \frac{M}{\mathbf{E}[X]}$, we have:

$$\begin{aligned}
\mathbf{E}[\text{leakers}] &\approx \mathbf{E}[\text{leakers}_{AS}] \\
&\equiv \mathbf{E}[W] \times (1-p)^N \prod_{i=1}^{K+1} \theta_i
\end{aligned} \tag{28}$$

$$\begin{aligned}
\mathbf{E}[\text{interceptors fired}] &\approx \mathbf{E}[\text{interceptors fired}_{AS}] \\
&\equiv \mathbf{E}[W]\mathbf{E}[X].
\end{aligned} \tag{29}$$

Equation (29) is identical to (10). Equation (28) is similar to (9), but (28) must also account for the likelihood the SM kills the threat.

For the large-W case, when $\mathbf{E}[W] > \frac{M}{\mathbf{E}[X]}$, we have:

$$\begin{aligned} \mathbf{E}[\text{leakers}] \approx & \mathbf{E}[\text{leakers}_{AL}] \\ \equiv & \mathbf{E}[W] - \mathbf{E}[\text{total kills}] \\ & = (\mathbf{E}[W] - \mathbf{E}[\text{HD Kills}]) \times \prod_{i=1}^{K+1} \theta_i \end{aligned} \tag{30}$$

$$\begin{aligned} \mathbf{E}[\text{interceptors fired}] \approx & \mathbf{E}[\text{interceptors fired}_{AL}] \\ \equiv & M. \end{aligned} \tag{31}$$

$\mathbf{E}[\text{total kills}]$ is combined kills from both interceptors and SM. $\mathbf{E}[\text{HD Kills}]$ is the expected number of threats killed by Blue interceptors. Equations (30)–(31) are very similar to (11)–(12). (31) is identical to (12) and equations (30) and (11) are both defined as $\mathbf{E}[W] - \mathbf{E}[\text{total kills}]$. However, there is a crucial difference between the large-W case in the HD model of Section 3.1 and the CD model here regarding how we compute $\mathbf{E}[\text{total kills}]$. In the HD case of (11) all threats not killed by interceptors leak. However, in the CD model in (30), Blue can still kill a threat even if the interceptors fail. Thus we multiply the "HD leakers" in (30) by the SM survival probability $\prod_{i=1}^{K+1} \theta_i$ to generate the final expression for the expected number of leakers when Blue also employs SM in the CD case.

In Online Appendix C.3.2 we show these approximations perform well numerically.

4.2. Evaluating CD Efficiency

As in Section 3.2 we use the AS and AL approximations from Section 4.1 to determine effective firing sequences. Blue benefits from SM by conserving interceptors while maintaining an acceptable single-threat kill probability. In this section we examine the properties of Blue's preferred HD firing sequences.

There are essentially two ways to incorporate SM in CD: parallel and sequential. In the parallel case, the operation of SM is parallel to and independent of the launching of the interceptors. That is, the SM parameters θ, γ and ψ are independent of the sequence $\mathbf{s}$ of salvos. Examples of a parallel SM are decoys and directed-energy weapons. In the sequential setting, the SM can only be executed prior to initiating the HD salvo sequence (or after the HD sequence completes). For example, jamming the threat's command channel may be performed in such a sequential fashion to avoid jamming Blue's own interceptors (e.g., (Blodgett et al. 2001, 2002, Sui et al. 2009)). The CD model handles both parallel and sequential settings, however the results presented in this section apply to the parallel case. We address the sequential setting in Online Appendix F.

4.2.1. Large-W Case When Red's attack comprises many threats, SM does not help Blue to conserve its interceptors, and the entire inventory depletes before Blue engages all threats. Consequently, Blue's aim is to minimize the expected number of leakers. Examining equation (30), we see the same implication as from the HD model in (11): Blue wants to maximize the expected number of total threats killed. In the HD model in Section 3.2.1, there is a one-to-one correspondence between kills and successful interceptors. That is not the case in the CD model because of SM. To fully leverage SM and to avoid "wasted" HD interceptions that coincide with SM kills, Blue desires all successful interceptions to occur in the final SEO. The following theorem summarizes this logic

THEOREM 8. *In the large-W case, Blue maximizes the number of threats killed by utilizing the firing sequence* $\mathbf{s} = (0, 0, 0, 0, \ldots, 0, 1)$.

In the HD model in Section 3.2.1, Blue is indifferent among the various one-interceptor-per-salvo sequences (Theorem 3). However, in the CD model Blue prefers the one-interceptor-in-final-SEO sequence of Theorem 8 to fully take advantage of SM.

While in this large-W case, SM cannot conserve inventory, SM still plays a role in determining how large $\mathbf{E}[W]$ needs to be so that the large-W case occurs. This happens when $\mathbf{E}[W] > \frac{M}{\mathbf{E}[X]}$, where we note that the expression for $\mathbf{E}[X]$ in (24) depends upon SM via the ψ_k parameters. For example, if Blue utilizes the firing sequence in Theorem 8, then $\mathbf{E}[X] = \psi_K$. If SM is very effective, either through kill probabilities or accurate BDA, then the ψ_k values will be small, decreasing $\mathbf{E}[X]$ and increasing the threshold for when the large-W case occurs.

4.2.2. Small-W Case In the small-W case, SM allows Blue to potentially conserve interceptors by employing interceptors later in the SEO's sequence. We first note that since we assume here that SM is an unlimited and no-cost resource that does not interfere with the employment of interceptors, it is a dominant strategy for Blue to continuously employ SM throughout the engagement sequence, starting at SEO 1.

Since Blue's use of SM is unrestricted, Blue prefers SM with higher lethality (lower θ_k) and reduced false-negative probabilities (lower γ_k), particularly in early rounds where these attributes are most valuable. We present a formalization of these points in Online Appendix E, though we omit them here due to their intuitive nature.

We now examine the properties of efficient sequences (see Remark 1) and how SM affects them. The monotonicity result for efficient sequences for the HD model in Theorem 6 holds also for the

CD case. There is also a monotonicity result for N, the total number of interceptors in the firing sequence. As N increases, the number of interceptors fired in each SEO of an efficient sequence can only increase. Kalyanam and Clarkson (2021) present a similar result in the HD setting; here we extend it to the CD case.

THEOREM 9. *For a fixed K, If $\mathbf{s}(N)$ is the efficient firing sequence for N interceptors, then*

$$s_k(N) \leq s_k(N+1) \quad \text{for } 1 \leq k \leq K. \tag{32}$$

The proof for Theorem 9 appears in Online Appendix J. Theorem 9 provides the basis for a simple algorithm to compute the efficient firing sequence. First set $\mathbf{s}(1) = (0,0,0,\ldots,0,1)$ by monotonicity. We then construct $\mathbf{s}(N)$ by examining K modifications of $\mathbf{s}(N-1)$. Modification k adds one interceptor to SEO k. The modification that generates the smallest $\mathbf{E}[X]$ is $\mathbf{s}(N)$. This interceptor allocation algorithm follows a similar approach to the Maximum Marginal Return (MMR) strategy commonly used in Weapon Target Assignment models (denBroeder Jr et al. 1959). In MMR, weapons are assigned sequentially to targets, with each new assignment selected to yield the maximum possible improvement in the objective function. Analogously, the interceptor allocation algorithm suggested by Theorem 9 applies this logic to assign weapons (interceptors) to SEOs rather than targets.

The next two results show that the firing sequence becomes "more" monotonic as SM is introduced. That is, an efficient sequence with SM pushes more interceptors to later SEOs.

THEOREM 10. *For fixed K and N, define $\mathbf{s}^*$ as the efficient firing sequence for HD (i.e., $\psi_k = 1$ for all k). Define $\mathbf{s}^*(\psi)$ as the efficient sequence for CD with SM parameter vector ψ. If*

$$s_j^*(\psi) > s_j^* \tag{33}$$

for some j, then

$$s_i^*(\psi) \geq s_i^* \quad \forall i > j. \tag{34}$$

The proof for Theorem 10 appears in Online Appendix K.

Theorem 10 examines the relation between efficient sequences for HD to CD. Theorem 11 compares two different CD scenarios by examining how the efficient firing sequence changes if the SM improves after a certain SEO. Specifically, if SM B is equivalent to SM A up to round H but better ($\psi(B)$ is less than $\psi(A)$) after round H, then Blue shifts more interceptors to SEOs after round H under SM B compared to SM A.

THEOREM 11. *For fixed number of SEOs K and number of interceptors N, define* $\mathbf{s}^*(\psi)$ *as the efficient firing sequence as a function of* ψ*. Let A and B be two SMs with parameter vectors* $\psi(A)$ *and* $\psi(B)$*, respectively.* $\psi(B)$ *is a modification of* $\psi(A)$ *such that* $\psi_k(B) = \psi_k(A)$ *if* $k \le H$*, and* $\psi_k(B) \le \psi_k(A)$ *if* $H < k \le K$*. Then*

$$s_k^*(\psi(B)) \le s_k^*(\psi(A)) \quad \forall k \le H. \tag{35}$$

The proof for Theorem 11 appears in Online Appendix L. It might seem intuitive that Blue would maintain the same firing sequence up to SEO H, given that SM A and SM B are equivalent up to that point. However, Blue anticipates the improved performance of SM B in subsequent rounds and strategically reallocates more interceptors to later SEOs. This adjustment seeks to conserve interceptors by leveraging the improved SM's capabilities to achieve a kill in the later rounds.

Recall that in the HD case Theorem 5 states that Blue utilizes all of its K SEOs. This is no longer necessarily true in the CD case. The following result generalizes Theorem 5 and provides a condition for when Blue continues to utilize all SEOs even with SM.

THEOREM 12. *If* $N \ge K$ *and*

$$\frac{\psi_K}{\psi_1} > \frac{1}{1 + p(1-p)^{K-2}}, \tag{36}$$

then the efficient Blue firing sequence utilizes all SEOs (i.e., $s_k > 0$ *for all k)*

The proof for Theorem 12 appears in Online Appendix G. A relatively large value of $\frac{\psi_K}{\psi_1}$, which is never greater than 1 (see Lemma 1 in Online Appendix B.3), implies that the improvement in the effectiveness of the SM from round to round during the threat engagement is small. Hence, Blue is more likely to utilize all SEOs for launching interceptors. Condition (36) specifies the threshold for this to happen. Because $\frac{\psi_K}{\psi_1} = 1$ in the HD setting, condition (36) always holds in the HD case.

Recall from the HD case that if $N \geq K$, the "extreme" firing sequence is the one where only one interceptor is launched in each of the first $K-1$ SEOs, and all the remaining $(N-K+1)$ interceptors are launched in the last $(K-th)$ SEO. Theorem 7 provides a threshold for when this sequence is an efficient firing sequence. As discussed earlier, in the CD case Blue does does not necessarily utilize all SEOs for launching interceptors. Thus the corresponding "extreme" firing sequence in the CD case is when Blue fires all interceptors in the last SEO. Theorem 13 is the analog to Theorem 7 and provides the condition for when this extreme sequence is efficient in the CD case.

THEOREM 13. *An efficient Blue sequence only utilizes the last SEO (i.e.,* $\mathbf{s} = (0,0,0,\ldots,0,N)$ *) if and only if*

$$p < \frac{1}{N-1}\frac{\psi_{K-1} - \psi_K}{\psi_K}. \tag{37}$$

The proof for Theorem 13 appears in Online Appendix M. The key factor in Theorem 13 is the relative drop from ψ_{K-1} to ψ_K. A large decrease from ψ_{K-1} to ψ_K makes it more likely Blue uses the fire-all-in-last-SEO sequence.

4.3. Numerical Analysis of the CD Case

This section illustrates the effect of the SM on the efficient frontier. In Online Appendix F, we consider a scenario where the SM must be executed prior to launching interceptors; this introduces a dependence between the fire sequence $\mathbf{s}$ and the SM's survival parameter θ. For the figures in this section, we use the recursions in (26) and (27) to calculate the MOEs.

The following assumptions are similar to those in Section 3.3

- $p \in \{0.7, 0.9\}$
- $N \in \{1,2,3,4,5\}$
- $M = 100$
- $W \sim U[10,20]$: small-W setting
- $K = 3$
- The SM is only lethal in rounds 1–3, not in the final round after the last SEO (i.e., $\theta_4 = 1$). For example, the SM is not feasible at short range.

- The false-negative probability in the round a SM kill occurs is γ_0. The false-negative probability j rounds after the SM kills occurs is $\gamma_j = \gamma_0^{j+1}$. Many SMs (jamming, directed-energy, decoys) affect the command and control system of the threat, causing the threat to deviate from its intended course. It may take time for Blue to recognize such a deviation and determine that the threat has been neutralized. Hence the false-negative probability decreases with the number of rounds from the kill.

The combination of kill probability and BDA capability represent the *effectiveness* of the SM. Theorems 12–13 show that, depending on SM effectiveness, Blue may opt to skip early SEOs – not launch any interceptors – in the hope that the SM kills the threat. Thus saving the limited resource of interceptors. In this section, we explore how effective SM needs to be for Blue to consider altering its firing sequence to utilize fewer SEOs. Figure 4 plots how Blue's efficient frontier changes as the SM parameters θ and γ change.

Recall that so far we assume that Blue can execute HD and SM in parallel. A directed-energy weapon (e.g., laser) is an example of such parallel execution. In Online Appendix F, we examine a situation where SM, such as jamming, may interfere with the launching of interceptors and therefore the two measures must be employed sequentially.

In Figure 4 we assume that the SM survival probability is the same in each of the first three rounds: $\theta_1 = \theta_2 = \theta_3$. The x-axis in Figure 4 is the overall SM kill probability across all three rounds: $1 - \theta_1^3$. An overall SM kill probability of 0.8 on the x-axis of Figure 4 corresponds to one-round survival probability of $\theta_1 = \theta_2 = \theta_3 = (1 - 0.8)^{1/3} = 0.585$ and a one-round kill probability of 1-0.585 = 0.415. The y-axis in Figure 4 is the instantaneous false-negative probability: γ_0.

We compute the efficient frontier for various values of SM false-negative probability γ_0 and SM kill probability $1 - \theta_1^3$. Figure 4 partitions the (kill probability) x (false-negative probability) space according to the efficient frontier associated with each parameter combination. In the northwest corner of the figure (high false-negative error and low kill probability), the efficient frontier is the same as in the HD case because the SM is essentially useless. As we move to the southeast, and thus improve the performance of the SM, the efficient frontier uses fewer SEOs and places more interceptors in the last salvo.

We next explain the labels for each region in Figure 4. Each region comprises a set of firing sequences that lie on the efficient frontier. For example (1,1,1) is a firing sequence with $N = 3$ that launches one interceptor in each SEO.

- A: $\{$(1); (1,1); (1,1,1); (1,1,2); (1,1,3) $\}$. This is the efficient frontier for HD in the small-W case by Theorem 7.
- B: $\{$(1); (1,1); *(1,2)*; (1,1,2); (1,1,3)$\}$. When $N = 3$, the efficient firing sequence (1,1,1) switches to (1,2).
- C: $\{$(1); (1,1); (1,2); *(1,3)*; (1,1,3)$\}$. When $N = 4$, the efficient firing sequence (1,1,2) switches to (1,3).
- D: $\{$(1); (1,1); (1,2); (1,3); *(1,4)*$\}$. When $N = 5$, the efficient firing sequence (1,1,3) switches to (1,4). The SM is so effective that Blue forgoes the first SEO.
- E: $\{$(1); *(2)*; (1,2); (1,3); (1,4) $\}$. Blue now starts skipping the second SEO as the efficient sequence for $N = 2$ switches from (1,1) to (2).
- F: as we move further to the southeast, Blue continues to opt out of the second SEO for larger values of N. In the extreme case, in the very southeast corner when SM is extremely effective, Blue only utilizes the last SEO: $\{$(1); (2); (3); (4); (5)$\}$

Figure 4 illustrates how Blue's efficient frontier compresses as the SM performance improves (i.e., we move to the southeast in the figure). When the SM is not effective (region A), Blue maintains its HD firing sequence. As SM improves, Blue considers skipping the first SEO in the hope the SM can kill the threat without Blue firing any interceptors. Blue's efficient frontier first swaps (1,2) for (1,1,1) (region B), then (1,3) for (1,1,2) (region C), and finally (1,4) for (1,1,3) (region D). This shift in Region D from (1,1,3) to (1,4) is potentially very costly; if the first interceptor misses, Blue must fire 4 additional interceptors in the final salvo (as opposed to one additional interceptor if Blue used (1,1,3)). Consequently, Blue must be very confident in the performance of the SM to prefer (1,4) to (1,1,3). Eventually the SM is so effective that Blue switches to the 1-salvo strategies in regions E and F: (2), (3), (4), and eventually (5). In the lower right corner, Blue is willing to skip two SEOs because SM is likely to achieve a kill on its own.

The boundaries for Regions A, B, and C are all farther to the northwest for $p = 0.9$ (Figure 4b) than for $p = 0.7$ (Figure 4a). Blue can be more aggressive with skipping the first salvo with $p = 0.9$ because Blue knows it will likely kill the threat with one HD interceptor in SEO-2 if SM has not already killed the threat. Note that Region D is much larger for $p = 0.9$ compared to $p = 0.7$. As we move from Region D to Region E, the efficient frontier starts including 1-salvo sequences (e.g., (3) over (1,2)). With $p = 0.9$, 1-salvo strategies are usually overkill. Consequently, Blue will not consider 1-salvo sequences unless the SM is so effective that Blue will rarely have to actually fire the HD interceptors.

Finally, the thresholds between regions in Figure 4 are nearly vertical for moderate false-negative probabilities (i.e., below 0.5). This suggests that as long as Blue has reasonable BDA capabilities, the SM kill probability primarily determines when Blue skips early SEOs.

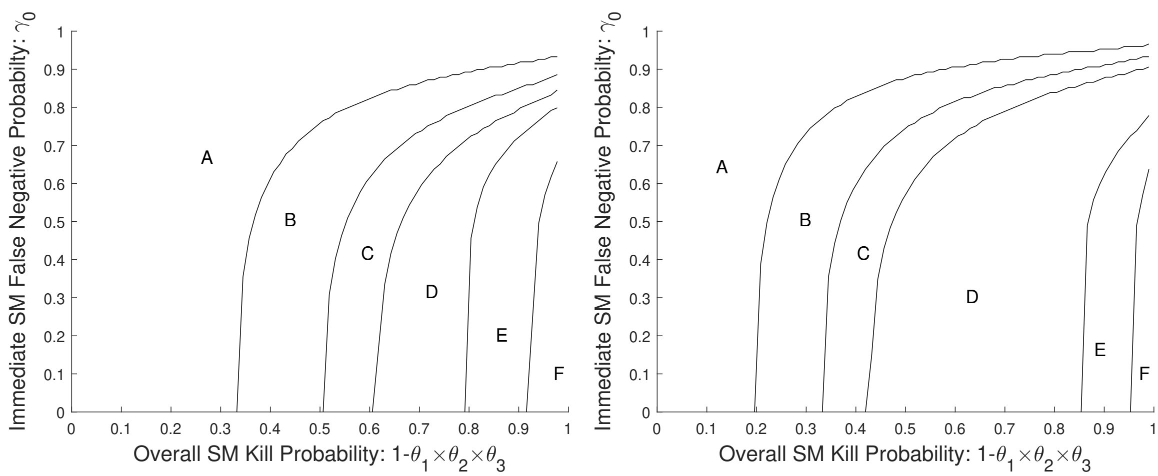

(a) $p = 0.7$ (b) $p = 0.9$

Figure 4 Efficient Frontier Regions for $K = 3$ and $N \in \{1,2,3,4,5\}$ for different SM false-negative and SM kill probability combinations. In each region the efficient frontier is the same. See the text for the efficient frontier corresponding to each region

5. Conclusion

We begin by discussing the main contributions of our work, followed by an examination of its limitations and directions for future research.

5.1. Contributions

Current and potentially future world events underscore the crucial importance of missile defense against long-range aerial threats. In particular, the high cost of HD necessitates a careful management of interceptor usage and the employment of SM. In this paper, we address this need by developing effective and efficient tactics for such missile defense missions. To the best of our knowledge, we have, for the first time, formulated a firing theory model to study the optimal joint employment of interceptors and general SM in missile defense. By defining two relevant MOEs – one representing cost (interceptors expended), and one representing damage (leakers) – we construct efficient frontiers for defense strategies. For certain realistic combat situations we specify optimal firing sequences for the interceptors, and for other situations, the efficient frontiers can guide mission planners about the trade-off between the expenditure of interceptors and the acceptable level of damage, measured by the number of leakers. In the following we summarize the main operational results and managerial insights:

- For the large-W case when Red launches an overwhelming number of threats that will likely deplete Blue's interceptor inventory, Blue should counter with a one-interceptor-per-salvo sequence for (HD) (Theorem 3) and use the $(0, 0, \dots, 0, 1)$ sequence for CD (Theorem 8).

For the small-W case:

- If the HD SSPK p is sufficiently large for fixed values of K and N, then Blue's efficient sequence is to launch one interceptor in each of the first $K - 1$ SEOs and allocate the remaining $(N - K + 1)$ interceptors to the final (Kth) SEO (Theorem 7).
- Blue should invest in command and control technology (e.g., radar) to obtain a single BDA look (i.e., $K = 2$); however, additional investments to obtain more looks may not be cost-effective (Section 3.3.1).
- The presence of SM shifts more interceptors to later SEOs (Theorem 10, 11).
- Blue utilizes all SEOs for HD when SM has limited effectiveness (Theorem 12).
- If the HD SSPK p falls below a threshold determined by SM lethality and BDA capability, the efficient firing sequence concentrates all interceptors in the final (Kth) SEO. (Theorem 13).
- Figure 4 illustrates the combinations of SM lethality and BDA capability that enable Blue to conserve interceptors by skipping SEOs.

5.2. Model Limitations

As a first attempt at combining hard and soft missile defense, the proposed model has certain limitations that suggest avenues for future research. For instance, the model currently only accounts for one type of BDA error: a false negative signal of a SM kill. While this is the most likely error, other types of BDA errors could arise, adding complexity to the modeling process. For example, to incorporate false positives, we might assume that after a delay, Blue becomes aware that the threat has not been destroyed and resumes engagement. Blue may not simply return to its original firing sequence but could adopt a more aggressive sequence to make up for the SEOs lost while the threat was mistakenly deemed killed.

We assume that Blue's SSPK p remains constant across SEOs, though in reality interceptor performance often varies with range. Additionally, Blue may deploy an assortment of types of HD interceptors, such as short- and long-range options. While our modeling framework–specifically the recursive formulas in equations (7), (8), (26), and (27)–could generalize to capture these complexities, this would complicate the bookkeeping, and we would likely be unable to derive any theoretical results on efficient frontier properties. Efficient frontiers could still be generated for a variable-SSPK scenario, but would require considering all possible sequences of N interceptors rather than just monotonic sequences. In the case of multiple interceptor types, we would need to track the expected number of interceptors fired by type. The efficient frontier could then incorporate an aggregate metric of interceptor expenditure cost, weighted across types, in place of the expected number of interceptors fired.

Finally, we assume that incoming threats are launched and engaged one at a time. In practice, however, Red may attack with *raids*, which are multiple threats launched simultaneously. If Blue has sufficient interceptor capacity to engage each threat in a raid with Blue's full firing sequence, then our results and analysis remain valid. However, if Blue faces restrictions on the number of interceptors it can deploy at once, Blue may need to adapt its strategy by developing a family of firing policies tailored to varying raid sizes.

5.3. Future Work

In addition to the research directions suggested in the previous subsection, we outline several areas for future exploration. The most realistic and complex scenario involves Red launching a series of raids, each comprising a mixture of different threat types such as cruise missiles, ballistic missiles, and drones. In response, Blue would deploy various interceptor types, potentially fired from multiple launch locations, each with distinct inventories of interceptors. The firing sequence would depend upon the threat type. Analyzing this scenario will likely require a simulation approach, although elements of weapon target assignment may also play a significant role.

When Blue's SSPK varies with range, Blue may face an interesting timing decision. Blue could delay firing the current salvo to increase the SSPK for that salvo, but this delay might reduce the number of SEOs available for future engagements.

Finally, our work highlights several interesting inventory management challenges. For an extended battle scenario, Blue will need to determine its inventory replenishment policy. If multiple interceptor launch locations are involved, Blue faces the choice between centralized and distributed inventory strategies. Distributing interceptors across launch sites enhances responsiveness by ensuring that each site has inventory readily available for immediate use when needed. However, a centralized inventory might be more efficient overall, especially if launch sites experience different usage rates; moving interceptors between launch locations could be costly and time-consuming. An additional complexity is that inventory sites themselves may be vulnerable to attrition.

6. Code and Data Disclosure

The code to support the numerical experiments in this paper can be found at URL.

7. Author Biographies

Dr. Michael Atkinson is a Professor of Operations Research at the Naval Postgraduate School. He joined NPS in 2009 after earning his Ph.D. in Computational and Mathematical Engineering from Stanford University. His research applies stochastic models to military applications such as combat modeling, logistics, search and detection, intelligence collection, and missile defense. He teaches courses in statistics, machine learning, computational methods, stochastic models, search theory, decision theory, and game theory.

Dr. Moshe Kress is a Professor of Operations Research at the Naval Postgraduate School, where he teaches and conducts research in combat modeling and related areas. His current research interests are operations research application in military and defense intelligence, counter-insurgency modeling, homeland security problems, and military logistics. He published four books and over 80 papers in refereed journals. Dr. Kress has been twice awarded the Koopman Prize for military operations research (2005 and 2009) and the 2009 MOR Journal Award. Prior to joining NPS, Dr. Kress was a senior analyst at the Center for Military Analyses in Israel.

References

Armstrong MJ (2004) Effects of lethality in naval combat models. *Naval Research Logistics* 51(1):28–43.

Armstrong MJ (2005) A stochastic salvo model for naval surface combat. *Operations Research* 53(5):830–841.

Armstrong MJ (2014) Modeling short-range ballistic missile defense and Israel's Iron Dome system. *Operations Research* 62(5):1028–1039.

Aviv Y, Kress M (1997) Evaluating the effectiveness of shoot-look-shoot tactics in the presence of incomplete damage information. *Military Operations Research* 79–89.

Azhari T, al Badrani J, Rasheed A, Ali I (2024) Five rockets fired from Iraq towards U.S. military base in Syria. *Reuters* URL `https://www.reuters.com/world/five-rockets-fired-iraq-towards-us-military-base-syria-security-sources-say-2024-04-2`, accessed 01 May 2024.

Bartroff J, Samuel-Cahn E (2011) The fighter problem: optimal allocation of a discrete commodity. *Advances in Applied Probability* 43(1):121–130.

Blodgett D, Paquet S, Plamondon P, Chaib-draa B, Kropf P (2001) Coordinating plans for agents performing AAW hardkill and softkill for frigates. *Proceedings of The 2001 AAAI Fall Symposium Series*, 1–8.

Blodgett D, Plamondon P, Chaib-draa B, Kropf P, Bossé E (2002) A method to optimize ship maneuvers for the coordination of hardkill and softkill weapons within a frigate. *7th International Command and Control Research and Technology Symposium (7thICCRTS)*.

Brown G, Carlyle M, Diehl D, Kline J, Wood K (2005) A two-sided optimization for theater ballistic missile defense. *Operations research* 53(5):745–763.

Carr RW, Torf J, Cobb R (2016) Cooperative control for missile evasion. *AIAA Guidance, Navigation, and Control Conference*, 1862.

Critchlow CL, Williams RC (1982) A simulation model for analyzing reentry vehicle/antiballistic missile engagements. Master's Thesis, Air Force Institute of Technology, Wright-Patterson AFB, OH.

Davis MT, Robbins MJ, Lunday BJ (2017) Approximate dynamic programming for missile defense interceptor fire control. *European Journal of Operational Research* 259(3):873–886.

denBroeder Jr G, Ellison R, Emerling L (1959) On optimum target assignments. *Operations Research* 7(3):322–326.

Dutta D (2014) Probabilistic analysis of anti-ship missile defence effectiveness. *Defence Science Journal* 64(2):123.

Eckler AR, Burr SA (1972) *Mathematical models of target coverage and missile allocation* (Military Operations Research Society Alexandria, VA).

Engelbrecht GN (1997) A dynamic simulation model to optimize soft kill tactics employed by a generic fast attack craft against a generic surface-to-surface, fire-and-forget missile. Master's of Science, University of South Afrrica.

Eyal Pecht AT, Weingold N (2013) On the choice of multi-task R&D defense projects: A case study of the Israeli missile defense system. *Defence and Peace Economics* 24(5):429–448, URL `http://dx.doi.org/10.1080/10242694.2012.717205`.

George S (2024) Iran mustered fiercer attack than in the spring, warns of even worse. *Washington Post* URL `https://www.washingtonpost.com/world/2024/10/02/iran-israel-ballistic-missile-attack/`, accessed 18 October 2024.

Glazebrook K, Washburn A (2004) Shoot-look-shoot: A review and extension. *Operations Research* 52(3):454–463.

Gramann RH, Butler BM, Manoff MK (1963) A combinatorial model for determining penetration probability of warheads accompanied by decoys. Technical Report Weapons Systems Division Technical Paper RAC-TP-93, Research Analysis Corporation.

Hjelmgaard K, Brook T (2024) U.S. intercepts and shoots down Houthi missile aimed at Navy destroyer in Red Sea. *USA Today* URL `https://www.usatoday.com/story/news/world/2024/01/15/yemen-houthis-airstrikes-red-sea/72229760007/`, accessed 29 April 2024.

Hughes Jr WP (1995) A salvo model of warships in missile combat used to evaluate their staying power. *Naval Research Logistics* 42(2):267–289.

Jewish Virtual Library (2024) Rocket & mortar attacks against Israel by date. Https://www.jewishvirtuallibrary.org/palestinian-rocket-and-mortar-attacks-against-israel, Accessed 01 May 2024.

Kalyanam K, Clarkson J (2021) Sequential attack salvo size is monotonic nondecreasing in both time and inventory level. *Naval Research Logistics* 68(4):485–495.

Karasakal O, Özdemirel NE, Kandiller L (2011) Anti-ship missile defense for a naval task group. *Naval Research Logistics* 58(3):304–321.

Keller J (2021) Navy asks Northrop Grumman to build electronic warfare (EW) with soft kill capability for surface warships. *Military+Aerospace Electronics* URL `https://www.militaryaerospace.com/sensors/article/14203208/surface-warships-electronic-warfare-soft-kill`, accessed 2 February 2024.

Kline A, Ahner D, Hill R (2019) The weapon-target assignment problem. *Computers & Operations Research* 105:226–236.

Landau EB, Bermant A (2014) Iron Dome protection: missile defense in Israel's security concept. *The lessons of operation protective edge* 37–42.

Lendon B (2024) How Israel and allied defenses intercepted more than 300 Iranian missiles and drones. *CNN* URL `https://www.cnn.com/2024/04/14/middleeast/israel-air-missile-defense-iran-attack-intl-hnk-ml/index.html`, accessed 29 April 2024.

Majumdar S (2018) Arrow atbm: Lethal evolution. *Vayu Aerospace and Defence Review* 3:75–77.

Maskery M, Krishnamurthy V, O'Regan C (2007) Decentralized algorithms for netcentric force protection against antiship missiles. *IEEE Transactions on Aerospace and Electronic Systems* 43(4):1351–1372.

Moore FW (2002) A methodology for missile countermeasures optimization under uncertainty. *Evolutionary Computation* 10(2):129–149.

Popa CH, Stone SP, Aw EH, Teo CPJ (2018) Distributed maritime operations and unmanned systems tactical employment. Systems Engineering Analysis (SEA27), Naval Postgraduate School, Monterey, CA.

Seligman L, Berg M (2024) A $2m missile vs. a $2,000 drone: Pentagon worried over cost of Houthi attacks. *Politico* URL `https://www.politico.com/news/2023/12/19/missile-drone-pentagon-houthi-attacks-iran-00132480`, accessed 29 April 2024.

Sewell A, Mroue B (2024) Israel's military and Hezbollah exchange fire along tense Lebanon-Israel border. *Time* URL `https://time.com/6331589/israel-hezbollah-fire-strikes-lebanon-border/`, accessed 01 May 2024.

Simons G, Yao YC (1990) Some results on the bomber problem. *Advances in Applied Probability* 22(2):412–432.

Soland RM (1987) Optimal terminal defense tactics when several sequential engagements are possible. *Operations Research* 35(4):537–542.

Stöhs J (2021) *How High?: The Future of European Naval Power and the High-end Challenge* (Djof Publishing and Centre for Military Studies).

Sui J, Xing C, Wu L (2009) The conflict resolution of hardkill and softkill weapon in fleet cooperative air defense. *2009 Chinese Control and Decision Conference*, 5024–5028 (IEEE).

Sun W, Zhang Y, Wang Y (2021) Research on an effective jamming technique for infrared guided missiles. *2021 21st International Conference on Control, Automation and Systems (ICCAS)*, 2235–2239 (IEEE).

USCENTCOM (2024) Iranian-backed Houthi terrorists conduct multiple anti-ship ballistic missile attacks in the Southern Red Sea and Gulf of Aden. Https://www.centcom.mil/MEDIA/PRESS-RELEASES/Press-Release-View/Article/3668425/iranian-backed-houthi-terrorists-conduct-multiple-anti-ship-ballistic-missile-a/, Accessed 01 May 2024.

van Hooft P, Boswinke L (2021) Surviving the deadly skies: Integrated air and missile defence 2021-2035. The Hague Centre for Strategic Studies.

Washburn AR (2002) Notes on firing theory. Department of Operations Research, Naval Postgraduate School, Monterey, CA. https://faculty.nps.edu/awashburn/Files/Notes/FiringTheory.pdf.

Washburn AR, Kress M (2009) *Combat modeling*, volume 139 (Springer).

Weber R (2013) ABCs of the bomber problem and its relatives. *Annals of Operations Research* 208:187–208.

Wilkening DA (2000) A simple model for calculating ballistic missile defense effectiveness. *Science & Global Security* 8(2):183–215.

Williams L (2024) Mideast missile duels have cost us navy nearly $1b, secretary says. *Defense One* URL `https://www.defenseone.com/threats/2024/04/mideast-missile-duels-have-cost-us-navy-nearly-1b-secretary-says/395791/`, accessed 29 April 2024.

Wu Z, Luo Y, Hu S (2024) Optimization of jamming formation of USV offboard active decoy clusters based on an improved pso algorithm. *Defence Technology* 32:529–540.

Zhai C, He F, Hong Y, Wang L, Yao Y (2016) Coverage-based interception algorithm of multiple interceptors against the target involving decoys. *Journal of Guidance, Control, and Dynamics* 39(7):1647–1653.