

Naval Gunfire Support Under Uncertainty

T. L. Keevers^a

^a*Joint and Operation Analysis Division, Defence Science and Technology Group, Sydney, Australia*
Email: Thomas.Keevers@dst.defence.gov.au

Abstract: Warfare is inherently conducted under risk and uncertainty. Decisions need to be made quickly, often on the basis of incomplete and dynamically changing information. The ability to accurately recalibrate beliefs and modify tactics in response to new information is essential to effective warfighting.

A Bayesian framework is used to examine the optimal combat tactics and resource allocation in a naval gunfire support scenario, specifically, the optimal timing to withdraw from combat. Enemy targets reside in an area of operation and a friendly platform is available that can fire upon the targets. Elimination of the targets would improve the regional stability and supply chain effectiveness. However, there is an opportunity cost associated with allocating the platform to that particular region, creating a trade-off between the potential benefits of eliminating the adversary and the opportunity cost of the platform.

There is an additional complication in that the platform's ability to eliminate the targets is poorly characterized. This could arise from the novelty of the platform itself, a change in the operational environment, or changes in the adversary's tactics. Immediate feedback about whether the gunfire support hit or missed is available after each shot, so the platform's capabilities become better characterized throughout combat, leading to unpredictable and dynamic changes in the optimal tactics. The cases of a single target and an ensemble of homogeneous targets are modelled. It is found that ensemble combat motivates the platform to continue attacking longer than in the single target case as information gain becomes more valuable.

We compare the optimal Bayesian framework against other strategies such as unconditional attack, unconditional withdrawal, use of the posterior mode, and using the mean value of the prior distribution without updating the posterior distribution. The penalty of these approaches depends on the precise situations. We show that these strategies can fail catastrophically and identify the situations in which this is likely to occur.

This work provides a theoretical basis for incorporating uncertainty and information gain into contemporary combat models. It identifies ways in which classical approaches involving maximum likelihood estimates and basic sensitivity analysis can lead analysts astray. We anticipate this work will be important given the proliferation of new domains for combat (such as space and cyberspace) and difficulty in providing timely and low cost assessments of emerging strategies and weapon systems.

Keywords: *Naval gunfire support, Bayesian statistics, optimal stopping time, decision making*

1 INTRODUCTION

Uncertainty and risk are undesirable, yet inevitable, aspects of warfare. They can lead to misallocation of resources, slower and poorer quality decisions, and more volatile outcomes. Combat models were developed to ameliorate the effects of uncertainty. Early combat models, like the Lanchester model, tended to be deterministic and simple in the sense that they could be solved analytically [Przemieniecki, 2000]. The development of faster computers and more robust numerical algorithms motivated agent-based simulations with stochastic events and time-varying costs. The focus of these models is usually optimizing tactics given information provided prior to the engagement, as opposed to during combat [Glazebrook and Washburn, 2004].

We explore changes in uncertainty and optimal combat tactics over time using a Bayesian framework. Bayesian reasoning has been used to estimate the parameters for combat models, such as the Lanchester model [Wiper et al., 2000; Pettit et al., 2003], but not to model changes in combat expectations (like the probability of hit) over time. We investigate the role of information gain using a toy problem in which one or more targets reside within an area of operation. These targets each impose a one-off cost of r if they survive, which represents the cost of regional stability and an impact on supply chains. Fortunately, a friendly platform is available to attack the enemy targets by firing one or more shots at a cost of one per shot. This is a prototypical naval gunfire support scenario (see Fig. 1).

The shots follow an independent and identically distributed Bernoulli distribution. Each shot either hits or misses with a fixed probability, and the outcome of each shot is immediately known. For this work, we will assume that only a single hit is required to kill a target. The probability of hit is unknown. Initial estimates for this parameter, the ‘prior distribution’, can be derived from historical data or subject matter experts.

Our methodology fits within the conventional framework of Markov Decisions Processes (see [Glazebrook and Washburn, 2004]). The differentiating factor is our treatment of uncertainty associated with the free parameters, and not just state variables such as whether the target is alive or dead. As we will see, this produces remarkable qualitative and quantitative changes in the optimal behaviour.

Following each hit or miss, the prior distribution should be rationally updated using Bayes’ rule,

$$P(\theta|x) = \frac{P(x|\theta)P(\theta)}{P(x)}, \quad (1)$$

where θ is the probability of hit, x indicates a hit or miss, $P(\theta)$ is our prior beliefs about the probability of hit, $P(x|\theta)$ is a likelihood function that describes how likely we believe a hit is given our prior distribution, $P(\theta|x)$ is the ‘posterior distribution’ that describes our new beliefs, and $P(x)$ is the probability of observing x . Bayes’ rule allows us to reweigh our beliefs given new evidence.

For mathematical convenience, we restrict the prior distribution for the probability of hit to a Beta distribution,

$$p_\theta(\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1}, \quad (2)$$

where α and β are shape parameters ($\alpha, \beta > 0$), θ is the probability of kill and p_θ is the prior distribution of θ .

Usually we would expect the posterior distribution to converge to the true probability over time. However, when there are only a small number of adversaries, only a limited number of successful hits can ever be recorded, so the decision to continue attacking or to withdraw always must be made under uncertainty. After n misses, the posterior distribution is given by a Beta($\alpha, \beta+n$) distribution, following Bayes’ theorem.

One strategy is to update the prior distribution using Bayes rule and maximize the expectation value of a decision over all future states (with appropriate weighting). This is optimal by construction. To evaluate the practical importance of rationally updating the prior beliefs, this approach is benchmarked against several other strategies, including using the mode of the posterior distribution, unconditional withdrawal, unconditional attack, and using the mean value of the prior distribution without updating it.

2 SINGLE TARGET

The simplest case we can consider is a single target and a sharply defined probability of hit (θ). There is no information updating for this scenario and the marginal value (m) per shot is fixed at

$$m = r\theta - 1. \quad (3)$$




Figure 1. A diagrammatic representation of naval gunfire support and the associated tradeoff between the cost of continuing the engagement and allowing the adversary to survive.

The rational course of action is to withdraw when the marginal value is negative or attack when the marginal value is positive. The expected value of the encounter ($E_\theta[\text{Encounter}]$) is the overall gain or loss of value associated with all of the attacker's choices. For a single target (summing over the sequence of decisions until either all targets are eliminated or the platform withdraws) this is

$$E_\theta[\text{Encounter}] = \begin{cases} -\frac{1}{\theta}, & \text{if } r \geq \frac{1}{\theta}, \\ -r, & \text{otherwise.} \end{cases} \quad (4)$$

There is a sharp decision boundary at $\theta = r^{-1}$ that separates withdrawal and attack. The two expectation values correspond to the expectation value of a geometric series and the cost of allowing the target to survive. When the probability of kill is uncertain, the decision boundary is a function of the number of misses and the values above become limiting cases. The marginal value after n consecutive failures is calculated as

$$m = r \frac{\alpha}{\alpha + \beta + n} - 1. \quad (5)$$

Noting that θ has been replaced by $\langle \theta \rangle$ by integrating over the Beta distribution. This process retains all of the relevant information in the prior distribution because of the updating process, thereby differing from classical point estimates.

The rational strategy is to shoot while the marginal value is positive, producing the decision boundary

$$r = \frac{1}{\alpha} n + \frac{\alpha + \beta}{\alpha}. \quad (6)$$

A minimum reward of $\frac{\alpha + \beta}{\alpha}$ is required for an initial engagement and an increase in the minimum reward of α^{-1} is required for an additional shot. We call this strategy 'Bayesian mean'. The initial cost is indicative of the initial expected value of kill and the gradient corresponds to how firmly the initial beliefs are held. A sharp probability distribution yields no gradient. The decision boundaries for several strategies are shown in Fig. 2a. The blue line (Bayesian mean) in Fig. 2a provides a visual description of uncertainty and can be used for expert elicitation, and provides a method for model validation.

One challenge is estimating the sensitivity of the decision boundaries to perturbations in the prior density. One method is to differentiate Eq. 6 to calculate the change in the boundary position as a function of α and β :

$$\frac{\partial r}{\partial \alpha} = -\frac{n + \beta}{\alpha^2}, \quad \frac{\partial r}{\partial \beta} = \frac{1}{\alpha}. \quad (7)$$

These terms represent how perturbations in the prior distribution affect the minimum reward to attack. A positive number means that a larger incentive is required, and vice versa for a negative value. An increase in α lowers the required reward, indicating the attacker is more confident, and obeys an inverse square law. In contrast, β always discourages engagement with the effect inversely correlated with α . The decision boundary is more sensitive to α than β because only a single kill can occur, while misses are potentially unlimited.

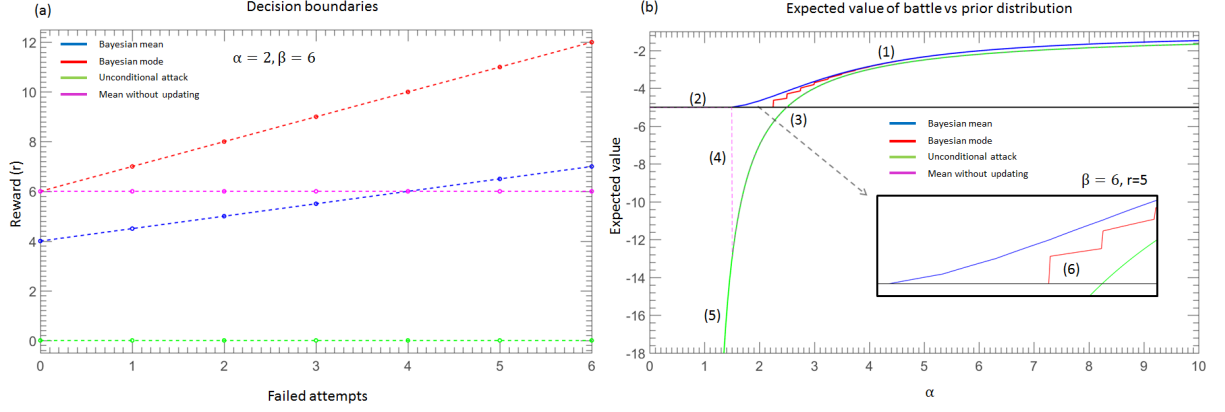


Figure 2. (a) The decision boundaries of several strategies. A reward above the dashed line indicates another attack will be made and below indicates withdrawal. (b) The expected value of a battle against a single adversary when $\beta=6$ and $r=5$ as a function of α . The inset is a zoomed in section of the plot showing the expected value of the Bayesian mode strategy is not smooth.

The probability that the first n shots miss and the subsequent shot hits is given by

$$P(n \text{ misses, then a hit}) = \underbrace{\frac{\beta}{\alpha + \beta}}_{P(\text{Miss}|0 \text{ hits, 0 misses})} \underbrace{\frac{\beta + 1}{\alpha + \beta + 1}}_{P(\text{Miss}|0 \text{ hits, 1 miss})} \cdots \underbrace{\frac{\beta + n}{\alpha + \beta + n - 1}}_{P(\text{Hit}|0 \text{ hits, } n-1 \text{ misses})} \underbrace{\frac{\alpha}{\alpha + \beta + n}}_{P(\text{Hit}|0 \text{ hits, } n \text{ misses})}. \quad (8)$$

or

$$P(n \text{ misses, then a hit}) = \frac{\alpha}{\beta + n} \prod_{N=0}^{n-1} \frac{\beta + N}{\alpha + \beta + N}. \quad (9)$$

The expected value of the encounter ($E_\theta[\text{Encounter}]$) can be calculated by summing over each possible future and weighting the probability of each future from the posterior distribution at each point in time. Repeating the process of sequentially updating the posterior distribution, the expectation value for the entire encounter is

$$E_{\alpha, \beta}[\text{Encounter}] = -r + \sum_{n=0}^{\lfloor \alpha(r-1) - \beta \rfloor} \underbrace{\left(r \frac{\alpha}{\alpha + \beta + n} - 1 \right)}_{\text{Expected reward for } n\text{-th shot}} \underbrace{\left(\frac{\Gamma(\alpha + \beta) \Gamma(\beta + n)}{\Gamma(\alpha + \beta + n) \Gamma(\beta)} \right)}_{\text{Probability of } n\text{-th shot occurring}}. \quad (10)$$

3 PHASE TRANSITIONS

So far we have treated the prior distribution consisting of shape parameters α and β as fixed, and examined the influence of varying the reward, r . An interesting and complementary picture is developed when considering r as a fixed quantity, while varying α and β . This is done by rearranging Eq. 5,

$$\beta = (r - 1)\alpha - n, \quad (11)$$

which produces a family of lines with intercepts $\frac{n}{r-1}$ and gradient $r-1$ in α and β -space, shown in Fig. 3a. These lines are decision boundaries, corresponding to zero marginal value following n misses. The Bayesian solution produces many possible decisions, while the classical model (Fig. 3b) produces only two.

The behaviour we see in Fig. 3 is analogous to phase transitions encountered in condensed matter physics. Traditionally, we might vary temperature and pressure, and observe transitions between the solid, liquid, and gas phases of a material. Here, we vary α and β , and see transitions in the number of shots that occur. This geometric viewpoint provides a clear notion of prior sensitivity, and provides a transparent process for comparing decisions. It also demonstrates the issues with validating classical decision making with sensitivity analysis. The right-side of Fig 3b shows a constant, but suboptimal, decision for a range of prior distributions.

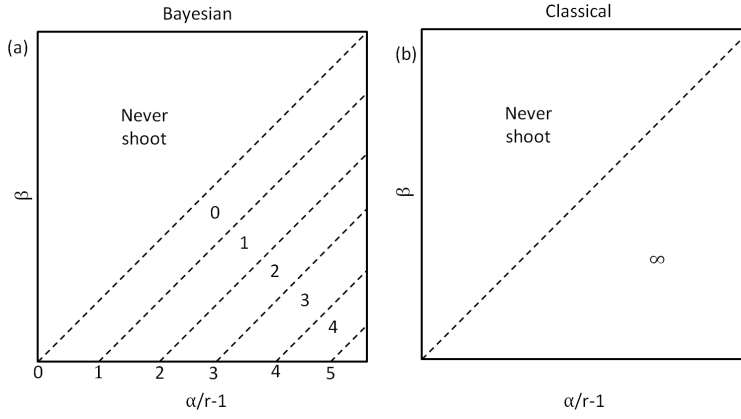


Figure 3. The number of misses that the attacker will observe for a fixed reward r as a function of α and β before withdrawing. The Bayesian solution is shown in (a) and the classical solution is shown in (b).

Table 1. The expected value for combat against a single target using various strategies.

Strategy	Expected value	Domain
Bayesian mean	$-r + \sum_{n=0}^{\lfloor \alpha(r-1) - \beta \rfloor} (r \frac{\alpha}{\alpha+\beta+n} - 1) (\frac{\Gamma(\alpha+\beta)\Gamma(\beta+n)}{\Gamma(\alpha+\beta+n)\Gamma(\beta)})$	
Bayesian mode	Undefined	$1 \leq \alpha, \beta$
	$-r + \sum_{n=0}^{\lfloor (\alpha-1)r - \alpha - \beta + 2 \rfloor} (r \frac{\alpha}{\alpha+\beta+n} - 1) (\frac{\Gamma(\alpha+\beta)\Gamma(\beta+n)}{\Gamma(\alpha+\beta+n)\Gamma(\beta)})$	$\alpha, \beta < 1$
Mean without updating	$-r$	$\frac{\alpha}{\alpha+\beta} r \leq 1$
	$-1 - \frac{\beta}{\alpha-1}$	$1 < \frac{\alpha}{\alpha+\beta} r$
Unconditional attack	$-1 - \frac{\beta}{\alpha-1}$	$1 \leq \alpha$
	Integral doesn't converge	$\alpha < 1$
Unconditional withdrawal	$-r$	

4 COMPARISON WITH OTHER STRATEGIES

The Bayesian approach, while optimal, is cumbersome to use. In this section, we examine whether simpler strategies can provide comparable results. We find they tend to perform well when the prior distribution is highly concentrated, but fail close to decision boundaries and for diffuse priors. The expectation values for each strategy is summarized in Table 1.

Unconditional withdrawal ensures the target survives, locking in a loss of $-r$. There is zero variance in the outcome and it is the optimal strategy when $r \frac{\alpha}{\alpha+\beta} \leq 1$ (the black line in region (2) of Fig. 2b).

Unconditional attack The other extreme is to continue attacking the target until he or she is eliminated. The expected number of shots follows a geometric distribution. The expected number of shots (s) is

$$E_{\alpha,\beta}[s] = E_{\alpha,\beta}[1/\theta] = \frac{\Gamma(\alpha+\beta)\Gamma(\alpha-1)}{\Gamma(\alpha)\Gamma(\alpha+\beta-1)} = 1 + \frac{\beta}{\alpha-1} \text{ (for } 1 < \alpha \text{).} \quad (12)$$

It is shown by the green line in Fig. 2b. The strategy works well when the probability of kill is high (region (1)). However, it can potentially lead to unbounded losses. The expected value diverges when $\alpha \leq 1$. Unconditional attack is the worst possible strategy when $r \frac{\alpha}{\alpha+\beta} \leq 1$ (region (5) in Fig. 2b).

Mean without updating A compromise between unconditional withdrawal and unconditional attack is to use the mean of the prior distribution, but not update it in response to new information. This helps to moderate the extremes of unconditional attack and unconditional withdrawal. This strategy can fail when the prior distribution is diffuse and the marginal value is slightly positive, since attacking will not switch off after several failures. The mean without updating matches the green curve in Fig. 2b until point (4), when it matches the black line (unconditional withdrawal). Between points (3) and (4)




Figure 4. The backpropagation algorithm for ensemble combat and the final decision tree that is produced. States with a positive marginal value are shown in green, states with a negative marginal value are shown in red. The platform will attack on green states and withdraw on red states.

the strategy performs worse than unconditional withdrawal, which is the critical region for dynamically adjusting tactics. This critical region is bounded by

$$\frac{1}{r} \leq \frac{\alpha}{\alpha + \beta} \text{ and } \frac{\beta}{\alpha - 1} \leq r - 1. \quad (13)$$

Bayesian mode Using the mode of the posterior distribution with Bayesian updating is close to the optimal strategy and, unsurprisingly, performs well. The mode can be more optimistic or pessimistic than the mean, depending on the specific prior distribution. Bayesian updating using the mode is shown in red in Fig. 2b, with the inset (6) emphasizing the presence of steps in the expected value. This is caused by the edge of decision boundaries between withdrawal and attack. The mode cannot be used when $\alpha \leq 1$ or $\beta \leq 1$, since the prior distribution then diverges at $\theta \rightarrow 0$ or $\theta \rightarrow 1$.

5 ENSEMBLE COMBAT

The decision about when to withdraw forces becomes more complicated when an ensemble of homogenous targets is present. Information gained by engagement with earlier targets informs the optimal tactics against later targets. This motivates increased aggression against early targets.

We will assume that the probability of hit is identical for each target. The case of independent probability distributions is trivial: The analysis for a single target is simply applied to each target individually. A more involved hierarchical model would be more realistic, but is beyond the scope of this work. As before, the platform has the option of continuing to attack, at a cost, or withdrawing and suffering a penalty from the target's continued operation. Target switching is not relevant since they're indistinguishable.

The decision tree is generated recursively by defining a grid structure that describes all potential belief states that the platform can reach (see Fig. 4), noting that the number of hits is limited by the number of targets, and that the number of misses is limited by the decision boundaries. The maximum number of misses that can occur corresponds to all but one target destroyed, and the state in which one more miss will cause the marginal value to become negative. This is used to define the size of the grid (a) and the marginal value is calculated for this point. The marginal value is propagated back to all the states from which it can be reached (b), so that each point represents the expected value until either all targets are destroyed or the platform withdraws.

When multiple targets are present, this backwards propagation introduces a non-negative value associated with information gain. For states with a negative value (the red circle in (c)), a value of zero is propagated backwards, since the platform will withdraw if it reaches the red state. Once this process is completed, a complete decision tree is generated (d). The circle in the top-left hand corner in Fig. 4 represents the platform's initial state. Pseudocode for this backpropagation algorithm is given in Algorithm 1.

6 CONCLUSIONS

In this paper, we considered the role of risk and uncertainty in the context of naval gunfire support. This was demonstrated through a toy problem in which a single platform can attack one or more targets, but with the

Algorithm 1 Algorithm for determining the optimal decision tree for ensembles of targets.

```

1: procedure BACKPROPAGATION
2:    $\alpha, \beta \leftarrow$  Prior parameters
3:    $r \leftarrow$  Reward per eliminated adversary
4:    $t \leftarrow$  Number of targets destroyed,  $T \leftarrow$  Number of targets
5:    $n \leftarrow$  Number of misses,  $N \leftarrow$  Maximum number of misses
6:   Calculate maximum number of misses with a positive marginal value:  $N \leq (r - 1)(\alpha + t - 1)$ 
7:    $V \leftarrow$  matrix of size  $(T, N+1)$  (zero indexing)
8:    $V(T-1, N) = r \frac{\alpha+T}{\alpha+T+\beta+N} - 1 \leftarrow$  Marginal value of last entry
9:   for diagonal entries from entries with the largest sum of indices to the smallest sum of indices:
10:    if  $V(t+1, n)$  defined and positive then
11:       $V(t, n) \leftarrow V(t, n) + \frac{\alpha+t}{\alpha+t+\beta+n} V(t+1, n)$ 
12:    if  $V(t, n+1)$  defined and positive then
13:       $V(t, n) \leftarrow V(t, n) + \frac{\beta+n}{\alpha+t+\beta+n} V(t, n+1)$ 
14:    Shoot on positive marginal values, otherwise withdraw

```

caveat that the platforms probability of hit is unknown. Over time the platform is able to collect data about its performance, and can dynamically adjust its tactics in response to this new information. We compared the optimal Bayesian strategy against several alternative strategies. The performance of these alternative strategies were found to sensitively depend on the particular scenario. Unconditional withdrawal and unconditional attack were especially vulnerable to producing sub-optimal outcomes. The benefit of a fully Bayesian approach tended to be strongest when there was high initial uncertainty, and the prior was close to the decision boundary.

In work not reported here, we examined a low-information scenario in which the targets survival is only known after the vessel has withdrawn from the area of operation, and a stochastic case in which the targets elimination or survival is only sometimes reported after a shot. We found that using a full Bayesian treatment also introduces new types of behaviors in these situations. Notably that a power law, rather than an exponential, decrease in the marginal value was found for the low information case. A paradox also occurs in that a miss may increase the marginal value in the stochastic feedback case.

Despite the simplicity of the toy model, it demonstrates the presence of sharp decision boundaries and the value of information. Our model shows that static estimates for the probability of kill combined with sensitivity analysis can produce misleading results. Contemporary combat models may not be fit for purpose.

This work provides a theoretical foundation for incorporating uncertainty into combat models. Tactics can be dynamically modified and the value of information can be rigorously quantified within our model. We have only considered the probability of hit. However, the importance of updating parameters can be extended to other areas, such as the potential reward for eliminating an enemy target. This is especially important given the proliferation of new domains for combat (space and cyberspace) and the difficulty in providing timely and low cost assessments of new strategies and weapon systems. We expect these approaches to be incorporated into other combat models, strategic forecasting, and force design.

ACKNOWLEDGMENTS

We thank Marc West, Paul Chircop, Maria Athanassenas, and Nicholas Armstrong for helpful comments.

REFERENCES

- Glazebrook, K. and A. Washburn (2004). Shoot-look-shoot: A review and extension. *Operations Research* 52(3), 454–463.
- Pettit, L., M. P. Wiper, and K. Young (2003). Bayesian inference for some Lanchester combat laws. *European Journal of Operational Research* 148(1), 152–165.
- Przemieniecki, J. S. (2000). *Mathematical methods in defense analyses*. Aiaa.
- Wiper, M., L. Pettit, and K. Young (2000). Bayesian inference for a Lanchester type combat model. *Naval Research Logistics (NRL)* 47(7), 541–558.