

A Verification Study of the Stochastic Salvo Combat Model¹

Michael J Armstrong

Department of Finance, Operations, & Information Systems

Faculty of Business, Brock University

St Catharines, ON, L2S 3A1, Canada

Phone 905-688-5550 x 5323

Fax 905-378-5723

michael.armstrong@brocku.ca

Abstract. When the stochastic version of the salvo combat model was designed, several assumptions and approximations were made to keep its mathematical structure relatively simple. This paper examines the impact of those simplifications by comparing the outputs of the stochastic model to those from a Monte Carlo simulation across 486 scenarios. The model generally performed very well, even where the battle size was relatively small or the damage inflicted by each missile was not normally distributed. The model's accuracy did decrease where missiles were positively correlated instead of independent.

Keywords: *missile combat; military operations research; simulation; stochastic models.*

1. Introduction

The salvo combat model was developed by Hughes (1995) to study naval battles in which the exchange of fire consists of discrete salvos of missiles. The model represents, in a simple aggregated form, the interaction between offensive firepower (e.g., Harpoon anti-ship cruise missiles) and defensive firepower (e.g., Standard surface-to-air missiles). It therefore describes naval missile combat in much the same way that the Lanchester model describes naval gunfire combat. While there have been relatively few such battles to date (Hughes, 2000), the potential for them has not disappeared. Indeed, technological developments are expanding the possibilities for future salvo-style battles to include anti-ship ballistic missiles (USNI, 2009) and super-cavitating torpedoes (Hambling, 2009).

A limitation of the original salvo model was that its deterministic nature made no provision for the variability or risk inherent in combat. But as other researchers have noted, it is important to account for variability in warfare models

¹ Published as: Armstrong MJ, 2011, A verification study of the stochastic salvo combat model, *Annals of Operations Research* 186 #1 (Jun), 23-38, doi 10.1007/s10479-011-0889-0. The final publication is available at Springer via <http://dx.doi.org/10.1007/s10479-011-0889-0>

(Jaiswal, 1987). Consequently, Armstrong (2005) developed a stochastic version that was simple enough to be implemented in spreadsheet software, as in (Armstrong, 2007) and (Armstrong & Powell, 2005).

One of the stated goals for the stochastic version was to retain most of the salvo model's relative simplicity, and so it incorporated a number of assumptions and approximations in its mathematical structure. This was not unreasonable; as Hughes (1997: p 38) noted in regards to combat models in general, "The unpredictability of a future environment is so wide ranging that one suffers no great additional penalty in accuracy by keeping the model uncomplicated and transparent." Nonetheless, these simplifications did raise questions about the accuracy of the stochastic model. Do the approximations significantly distort the model's outputs, or merely result in minor differences at the margins? And how robust is the model with respect to violations of its assumptions?

This paper addresses these questions by comparing the stochastic model's behavior with its design intent. It therefore aims at the model's *verification*: i.e., how well its "... implementation and its associated data accurately represent the developer's conceptual description and specifications" (MSCO, 2006). That is, does the model work as intended? This is an important issue to address, as analysts would understandably be reluctant to use a model if the quality of its outputs was poor or unknown.

Before proceeding further, it is worth noting several issues that this study does not address. Firstly, it does not test how well the model reflects actual naval combat, a concept known as *validation*: "...the degree to which a model and its associated data provide an accurate representation of the real world from the perspective of the intended uses of the model" (MSCO, 2006). Validation is a challenging task for computer models, and is especially difficult for combat

simulations. Ideally, it would require repeated comparisons of the model's predictions against actual naval battles; but this is rarely feasible. The studies of aircraft carrier battles in (Armstrong & Powell, 2005) and (Hughes, 2000), however, have provided at least a degree of *face validation*: "... whether a model or simulation seems reasonable to people who are knowledgeable about the system under study, based on performance" (MSCO, 2006).

The work herein also does not evaluate any simplifications that were originally made in the deterministic model and then carried forward into the stochastic model. For example, both models are aggregated treatments of combat, and so both ignore the allocation of specific attackers against specific enemy targets, as in (Ghose et al, 2002). The potential impact of these factors is significant, but beyond the scope of the present work.

Finally, this study does not attempt to prove that the stochastic model is "perfect", as it was deliberately designed with several major simplifications. Thus there is little point in testing whether the model's outputs are identical to those of exact calculations, as they clearly can not be. Rather, "... it is more useful to ask whether the differences between the system and the model are significant enough to affect any conclusions derived from the model" (Law & Kelton, 2000).

The rest of the paper proceeds as follows. The next section provides a brief review of the salvo combat model; readers wanting more detail should consult (Hughes, 1995) and (Armstrong, 2005). That is followed by a discussion of the main approximations and assumptions in the stochastic model: normal distributions to approximate binomial random variables, sufficiently large salvo sizes, normally distributed damage, and independence of missile successes. This is followed by a description of the numerical study that evaluated these factors. It compared the outputs of the stochastic model to those of a Monte Carlo

simulation across a set of 486 scenarios. Both were implemented in Excel spreadsheet software.

The numerical data showed that the model performed very well in the kind of scenarios for which it was originally intended. Surprisingly, the model continued to do well even when the number of missiles per salvo was quite small, and/or when the damage inflicted by each missile was not normally distributed. The model's accuracy did decrease, however, when missiles were positively correlated instead of being independent. Analytical work showed that this was because the model underestimated the variation in the number of successful missiles. A modified version of the model was able to partially compensate for this discrepancy. The paper finishes with a discussion of the study's results, along with their implications and limitations. It concludes that the accuracy and consistency of the model are respectable enough to justify its use, at least where correlation between missiles is minimal.

2. Overview of the Salvo Combat Model

Suppose that 2 opposing naval forces, Red and Blue, are meeting in battle. Red has A identical warships, each of which simultaneously fires α accurate offensive missiles at the Blue force. Blue has B identical ships, each with defensive power z , which is the number of incoming enemy missiles successfully intercepted per salvo by its active defenses. Each Red missile that is not intercepted causes damage v to a Blue ship. The number of Blue ships lost can be expressed as follows.

$$Loss_B = \sum^{Net_{AB}} v, \quad \text{where} \quad Net_{AB} = Offense_A - Defense_B = \sum^A \alpha - \sum^B z \quad (1 \& 2)$$

In the deterministic model (Hughes, 1995) all of these parameters were constants. By contrast, the stochastic model (Armstrong, 2005) allowed α , z , and

v to be random variables. For example, consider α as a random variable that represents n_α missiles fired by Red ship, where each missile has independent probability p_α of being accurate. Then the total number of accurate Red missiles $Offense_A$ will have a binomial distribution with mean and variance as follows.

$$E[Offense_A] = An_\alpha p_\alpha \quad \text{and} \quad \text{Var}[Offense_A] = An_\alpha p_\alpha (1 - p_\alpha) \quad (3 \ \& \ 4)$$

Likewise, if each of n_z attempted interceptions per Blue warship has an independent probability of success p_z , then Blue's total successful defensive fire $Defense_B$ will follow a binomial distribution, with $E[Defense_B] = Bn_z p_z$ and $\text{Var}[Defense_B] = Bn_z p_z (1 - p_z)$. The net quantity of Red missiles that survive to hit Blue will be $Net_{AB} = Offense_A - Defense_B$. The damage v caused by each hit is likewise an independent random variable with mean μ_v and variance σ_v^2 .

Given this situation, Armstrong (2005) derived the probability distribution of Blue's losses. The main assumptions and approximations that it employed are described in the next section.

3. Primary Assumptions and Approximations

3.1. Binomials Replaced by Normal Distributions

As mentioned above, if the probability of success for each missile is independent from every other missile, then the number of accurate offensive missiles and the number of successful defensive interceptions will follow binomial distributions. The stochastic salvo model uses a known approximation (Moore & McCabe, 1993: p 382) to replace each binomial distribution with a normal distribution. However, the model goes beyond past practice by using the difference of 2 normal variables to approximate the difference of 2 binomials. The performance of this multi-step approximation can be tested by examining how well the model works when its assumptions are all met.

3.2. Sufficiently Large Salvo Sizes

The normal approximation to a binomial generally fits better as the number of random events increases, and so the salvo model was intended for “... battles involving moderately small to moderately large forces on each side” (Armstrong, 2005). But how small is “moderately small”? For example, (Moore & McCabe, 1993: p 383) recommends that normal approximations only be used where the sample size n and probability of success p are large enough that $np \geq 10$ and $n(1-p) \geq 10$. For the probabilities used in this paper, this implies that there should be at least 20 to 60 missiles per salvo. How quickly does the model’s performance diminish as the number decreases below this range?

3.3. Damage Distributions

The stochastic model calculates the total damage caused by a salvo by summing the random damage v caused by each successful missile over the random number Net_{AB} of such missiles. Since Net_{AB} follows a normal distribution and the damage per missile is also assumed to follow a normal distribution, the model treats the compound distribution of the total damage as if it were normal as well. However, it is not clear whether the normal distribution is an appropriate representation of damage per missile. Thus the next factor to be investigated is the role of the probability distribution for missile damage. Does its shape materially affect the stochastic model’s outputs?

3.4. Missile Independence

The stochastic model assumes that each missile’s probability of success is statistically independent. However, there are factors in combat that can affect many missiles simultaneously. For example, an attacker might inadvertently aim its salvo of missiles along the wrong target bearing. Likewise, a defender might use counter measures such as chaff to confuse the guidance systems of an entire

salvo of incoming missiles. These factors would induce some positive association or correlation, and so the number of accurate missiles would no longer be binomially distributed.

This can be modeled by splitting the probability of success, such as p_α for offensive missiles, into two components. First, let probability q_α represent random influences that are specific to individual missile success and are statistically independent. Meanwhile, let probability r_α represent influences that are shared across all missiles in the salvo and are perfectly correlated. The product $p_\alpha = q_\alpha r_\alpha$ gives the combined probability of success for each missile, while the correlation between any two missiles is given by the following formula (all derivations are contained in the appendix).

$$\text{Cor}[\text{missile}_i, \text{missile}_j] = (q_\alpha - r_\alpha q_\alpha) / (1 - r_\alpha q_\alpha) \quad (5)$$

The number of successful missiles Offense_A would then have the following mean and variance.

$$E[\text{Offense}_A] = n_\alpha q_\alpha r_\alpha \quad \text{and} \quad \text{Var}[\text{Offense}_A] = n_\alpha q_\alpha r_\alpha (1 - q_\alpha + n_\alpha q_\alpha - n_\alpha q_\alpha r_\alpha) \quad (6 \ \& \ 7)$$

To illustrate, consider a salvo of 2 offensive missiles. If $r_\alpha = 1$ and $q_\alpha = 0.5$, then the probability of each missile hitting is $1 \times 0.5 = 0.5$, and their correlation is $(0.5 - 1 \times 0.5) / (1 - 1 \times 0.5) = 0$, so they are independent. The total number of hits therefore follows a binomial distribution with mean $2 \times 1 \times 0.5 = 1$ and variance $2 \times 1 \times 0.5 (1 - 0.5 - 2 \times 0.5 (1 - 1)) = 0.5$. The probability of 0 successful missile hits is 0.25, the probability of 1 hit is 0.50, and the probability of 2 hits is 0.25.

Next, suppose that $r_\alpha = 0.7143$ and $q_\alpha = 0.7$. The probability of each missile hitting is still $0.5 = 0.7143 \times 0.7$, but missile successes now have correlation 0.4. In this case, the average number of hits is still 1, but their distribution is no longer binomial and the variance increases to 0.7. The probability of 0 hits is 0.35, the probability of 1 hit is 0.30, and the probability of 2 hits is 0.35.

The extreme case occurs with $r_\alpha = 0.5$ and $q_\alpha = 1$, as then the correlation is 1. The hit probability is $0.5 \times 1 = 0.5$ as before, but both missiles either hit or miss together. The probability of 0 hits is 0.50, the probability of 1 hit is 0.00, and the probability of 2 hits is 0.50. The mean is still 1, but the variance becomes 1.

Given that such correlation can exist, it is worth investigating how well the model tolerates it.

4. Design of the Study

To address these research questions, this study compared the behavior of the stochastic salvo combat model to that of a Monte Carlo computer simulation. This numerical work used the same inputs (scenarios) with both approaches and then examined the extent to which their outputs (battle outcomes) differed. The stochastic salvo combat model (hereafter, “the model”) was implemented in Excel spreadsheet software, as in (Armstrong & Powell, 2005). The Monte Carlo simulation (“the simulation”) was also run in Excel, using the Crystal Ball add-in. The simulation involved the same combat mechanics as the model, but without its simplifying approximations. Both methods shared the same set of scenarios, most of which deliberately stretched the model beyond its original context.

For example, since the fit of a normal to a binomial improves with larger samples, the test scenarios were kept small. The largest one involved 18 offensive missiles fired in a salvo against a defender with 18 ships that can attempt 18 interceptions in total. This “18-missile-on-18-missile” scenario represented a relatively strong defense. This was accompanied by an 18-on-9 scenario (i.e., 18 offensive missiles fired at 18 defending ships that can attempt 9 interceptions in total) to represent a relatively strong offense. The next pair of scenarios was three times smaller, 6-on-6 and 6-on-3 in terms of the number of missiles per salvo.

The last two scenarios were a further three times smaller, just 2-on-2 and 2-on-1.

In each case, the number of defending ships was set equal to the number of incoming offensive missiles; in effect, each missile had its own target.

For each of these salvo combinations the probability p_α of offensive missile success was varied across 3 values: 50%, 67%, and 83%. These figures were previously used in (Armstrong, 2007), where the 67% approximated wartime missile accuracy as reported in (Hughes, 2000: p 276). The same 3 values were separately used for the probability p_z of defensive interception.

The simulation tried 3 differently shaped probability distributions for the damage per missile. The first was a normal distribution with a mean of 0.33 ships and a standard deviation of 0.11 ships, which was the base case in (Armstrong, 2007). The second was a uniform distribution over the range 0.1395 to 0.5205, which gave the same mean and standard deviation but with a different shape. The third was a triangular distribution, with a lower limit at 0.1744, a peak at 0.1744, and an upper limit at 0.6411. Although it was heavily skewed, this distribution again had the same mean and standard deviation.

The simulation also allowed for positively correlated missiles. It is theoretically possible to have correlation between two offensive missiles, or between two defensive missiles, or even between an offensive and a defensive missile. But to keep the numerical study to a manageable size, only offensive missile correlation was included, as it arguably was of greatest practical interest. The simulation tried 3 levels of correlation: 0.0 (independence), 0.2, and 0.4. By contrast, the model implicitly assumed independence in all cases.

Altogether these choices yielded 486 combinations of input parameters. The simulation ran each one of these combinations for 50000 trials (battles). Each run calculated the mean loss inflicted, measured in ships; the standard deviation of

the loss, in ships; the 95th percentile of the loss distribution, in ships; and the probability of suffering zero loss. These measures all dealt with what Armstrong (2005) referred to as the model's "actual" loss; i.e., after adjusting the "nominal" loss normal distribution to account for the fact that losses can neither fall below 0 nor can exceed the number of defending ships B . The mean and standard deviation evidently are key summary statistics of the loss distribution. The other two metrics provide information about the upper and lower tails of the distribution, and therefore quantify more extreme (rare) outcomes.

In terms of salvo combat *lethality* (Armstrong, 2004), these scenarios were of low lethality (e.g., the 18-on-18 case with $p_a = 50\%$ and $p_z = 83\%$) and moderate lethality (e.g., 18-on-9, $p_a = 83\%$, $p_z = 50\%$). The study avoided high lethality scenarios, as their tendency towards excessive losses or "overkill" would have overshadowed any differences between model and simulation.

5. Numerical Results

5.1. An Illustrative Example

To better understand the numerical results, consider first a single example scenario. Suppose the attacker has 6 ships that in total fire 6 statistically independent offensive missiles, each with a probability of 0.67 of hitting a target unless intercepted. The defender has 6 ships that in total attempt to shoot down 3 of these missiles, with each attempt having a probability of 0.67 for success. The loss caused by each offensive missile follows a normal distribution with a mean of 0.33 ships and a standard deviation of 0.11 ships.

For this 6-on-3 example, the loss per battle across the simulation's 50000 trials had a mean of 0.679 ships and a standard deviation of 0.467 ships. In the upper tail of the distribution, the 95th percentile was 1.499 ships; i.e., the defender

lost 1.499 ships or less in 95% (47500) of the trials. In the lower tail, the probability of zero loss was 0.141; i.e., 14.1% (7065) of the trials resulted in no losses. Figure 1 displays the range of simulation results in histogram form.

For the same scenario, the model estimated a mean loss of 0.678 ships, just 0.001 ships (0.1%) lower than the simulation. The standard deviation was 0.463, 0.004 ships (0.8%) lower than the simulation; the 95th percentile lay at 1.473, 0.026 ships (1.7%) lower; and the probability of zero loss was 0.156, or 0.015 (10.1%) higher. This latter difference was a side effect of the continuity correction (Moore & McCabe, 1993: p 384) that the model uses; it allocates probability mass for near-zero losses onto the zero point. Figure 2 shows the probability distribution produced by the model; it is similar but not identical to that of the simulation, with a small gap just above the zero loss.

With 50000 trials per simulation run, there should be relatively little difference from one run to another. To check this, the simulation was run 9 times more with the same inputs. Within this sample, the mean varied from 0.675 to 0.683, a range of 1.1%. The other ranges were also small: the standard deviation varied from 0.465 to 0.468 (0.6%), the 95th percentile from 1.488 to 1.502 (0.9%), and the probability of zero loss from 0.138 to 0.143 (3.4%).

5.2. The Overall Data Set

Now consider the full data set of 486 scenarios. One way to assess the fit between the model and the simulation is via scatter plots, as in Figure 3. Each dot represents a mean loss calculated by the model (vertical axis) contrasted with the corresponding value from the simulation (horizontal axis) for the same scenario. The diagram also includes a diagonal reference line to indicate where the dots hypothetically would fall if a perfect fit existed between model and simulation. In this plot, most of the points do fall close to the line, though not perfectly along it.

On the plot for the probability of zero loss (Figure 4) the data points display more scatter than before, though they are still centered around the diagonal. By contrast, the plots for the 95th percentile (Figure 5) and the standard deviation (Figure 6) show points drifting to the lower right, indicating that the model often underestimated the simulation for these metrics.

To quantify these visual impressions, linear regressions were calculated across the pairs of measurements with the regression constants fixed to zero. The top half of Table 1 displays the results of these regressions. Slope coefficients and R^2 values close to 1.000 indicate that the model outputs are similar to the simulation outputs. There is one column for each of the mean loss, the standard deviation of loss, the 95th percentile of loss, and the probability of zero loss.

For example, regression of the mean losses from the model against those of the simulation yielded $R^2 = 0.995$, showing that the model gave extremely consistent estimates of the simulation values; it explained 99.5% of the variation in them. The slope coefficient of 0.970 showed that the model tended to underestimate the simulation by about 3%. The figures for the probability of zero loss were likewise very good. The values for the standard deviation and 95th percentile were less impressive, however, as the model tended to underestimate the simulation.

The bottom half of Table 1 also quantifies the model-to-simulation fit. For this, the difference between each result from the model and the equivalent result from the simulation was calculated, and then the average and variance of these differences across all 486 scenarios was taken. Averages close to zero indicate that the model accurately estimates the simulation, whereas positive averages indicate a tendency to overestimate and negative averages a tendency to underestimate. Similarly, variances close to zero suggest that the model

consistently estimates the simulation, whereas larger variances suggest wider swings from scenario to scenario.

These figures confirmed that the fit was notably better for the mean loss and probability of zero loss than for the standard deviation and 95th percentile. For example, the model underestimated the standard deviation of loss by 0.09 ships on average, with a variance of 0.022. The following subsections explore the sources of these differences between the model and the simulation.

5.3. Consistency: Variance of the Differences

As mentioned above, if the variance of the differences between the model and the simulation is small, this implies that the model is relatively consistent from scenario to scenario. To investigate this, Levene's test (in Minitab statistics software) was used to see whether the variances differed systematically in the 486 scenarios along several dimensions: the shape of the damage distribution, size of salvo, offensive-defensive balance, and correlation of missiles.

The size of the salvo (18, 6, or 3 offensive missiles) did have a statistically significant effect on variability for all 4 metrics (p -value < 0.05): larger salvos naturally tended to show more variation from scenario to scenario than did smaller salvos. Likewise, the correlation between offensive missiles (0.0, 0.2, or 0.4) also made a statistically significant difference for all metrics: higher correlations gave more variability. The relative balance between offensive and defensive missiles (offense outnumbered defense, or offense equals defense) had a statistically significant effect on the means and the standard deviations, but not the 95th percentile or the probability of zero loss. By contrast, the damage distribution's shape (normal, uniform, or triangular) did not have a significant impact on variability for any of those metrics.

The comparisons above evaluated the differences between model and simulation measured in absolute terms. For example, if the standard deviation of loss from the simulation was 0.467 while in the model it was 0.463, then the absolute difference was 0.004 ships. But it would also be reasonable to evaluate this in terms of the relative difference, as in $0.004/0.467 = 0.85\%$. To investigate this alternative viewpoint, Levene's test was also applied to the relative (percentage) differences. Fortunately, the same factors remained significant.

5.4. Accuracy: Average of the Differences

Next consider the average difference between the model and the simulation. The average indicates whether the model's outputs tend to overestimate or underestimate the simulation's outputs overall. To see how these averages varied over the data set, a 4-factor Analysis Of Variance (ANOVA) test was run for each performance metric.

The tests revealed that the damage distribution shape had no statistically significant effect (p -value > 0.05) on average differences for any of the 4 metrics. The other 3 factors did have statistically significant impacts (p -value < 0.05) on all 4 metrics, except that correlation did not affect the probability of zero loss. The fit for all 4 ANOVA models was quite good, with $R^2 = 65.4\%$ for the mean loss, $R^2 = 94.2\%$ for the standard deviation of loss, $R^2 = 77.5\%$ for the 95th percentile of loss, and $R^2 = 51.5\%$ for the probability of zero loss.

These results must be viewed with caution, however, as the preceding Levene's tests had shown that the variances were quite different across the subsamples. This violates ANOVA's assumption of approximately equal variances and makes its results less reliable (Montgomery, 1991: Ch 4). Each multi-factor ANOVA test of means was therefore followed by a set of single

factor Kruskal-Wallis tests of medians. Fortunately, Kruskal-Wallis confirmed a difference in medians wherever ANOVA had indicated a difference in means.

Overall, the model tended to provide the closest average estimates of the simulation when salvos were small and/or missiles were uncorrelated. In scenarios with large correlated salvos, the model underestimated the simulation, especially for the standard deviation and the 95th percentile of loss.

6. Exploring Correlation Further

Given these results, the impact of missile correlation on the model's behavior is worth exploring further. As a base case, first consider the 162 scenarios with statistically independent missiles. Table 2 shows that the fit between model and simulation was very close in this case for all 4 measures. A scatter plot of the standard deviations (Figure 7) displays this clearly.

Table 3 by contrast covers the remaining 324 scenarios where offensive missiles were positively correlated. In this context the model underestimates the simulation, especially for the standard deviation and the 95th percentile. Figure 8 displays this graphically with its points drifting off to the right.

These numerical tendencies can be confirmed with some analytical work. As mentioned earlier, correlation can be modeled by breaking down the overall probability of missile success into two components. If this substitution is made in the model, with Equations 3 & 4 replaced by 6 & 7 respectively, then the model's outputs are affected as follows (see appendix for proofs).

Proposition 1. Consider a modification of the stochastic model in which the probability of offensive missile success p_α is replaced by an individual independent component q_α and a common correlated component r_α , as in $p_\alpha = q_\alpha r_\alpha$.

(Or equivalently for defensive missiles, $p_z=q_zr_z$). Then as missile correlation increases above 0, the model's behavior will change as follows:

- (a) The variance of the number of non-intercepted offensive missiles $\text{Var}[Net_{AB}]$ will increase, while the mean number $E[Net_{AB}]$ will remain the same;
- (b) The standard deviation of loss and the 95th percentile of loss will increase;
- (c) The probability of zero loss will increase if $E[Net_{AB}] > 0$, but it will decrease if $E[Net_{AB}] < 0$; and,
- (d) The mean loss will increase if $E[Net_{AB}] \leq 0$, it will decrease if $E[Net_{AB}] \geq B/\mu_v$, and it could either increase or decrease otherwise.

In addition to these effects, however, positive correlation also alters the shape of the probability distribution for the number of successful missiles. The distribution follows a binomial form when correlation is 0, but becomes a Bernoulli distribution if correlation equals 1, and takes some intermediate shape when correlation is between 0 and 1. This shape change is not accounted for by Equations 6 & 7 or the analytical results of Proposition 1.

Its effects were, however, revealed with some further numerical work. First, the spreadsheet for the model was modified by inserting the new parameters q_α & r_α and the corresponding Equations 6 & 7. Then this modified model was tested on the 324 scenarios with correlated offensive missiles. The q_α and r_α values that replaced p_α for scenarios with correlation 0.4 were $0.700 \times 0.7143 = 0.50$, $0.802 \times 0.8354 = 0.67$, and $0.898 \times 0.9243 = 0.83$, respectively. The values that gave correlation of 0.2 were $0.600 \times 0.8333 = 0.50$, $0.736 \times 0.9103 = 0.67$, and $0.864 \times 0.9606 = 0.83$, respectively. The results are summarized in Table 4.

Whereas the original model (Table 3) had tended to underestimate most metrics, the modified model (Table 4) tended to overestimate them, though to a lesser extent. For example, the regression slope for standard deviation of loss had

been 0.706 (i.e., about 29% too low) for the original model, but it became 1.174 (i.e., about 17% too high) with the modified model. This reversal can also be seen in the scatter plot for the standard deviation of loss (Figure 9), where the dots now drift above the diagonal rather than below it. Thus the distribution's changing shape partially negated the effect of its increasing variation, as the change in the simulation outputs was not as large as that in the model outputs.

7. Discussion & Conclusion

The stochastic salvo combat model of Armstrong (2005) performed surprisingly well overall in this numerical study. To begin with, its outputs closely matched those of the Monte Carlo simulation in the context for which the model was envisioned; i.e., with reasonably large salvo sizes, normally distributed damage, and independent missiles. Furthermore, it continued to do well in many other situations. The shape of the damage distribution had no significant impact on the model's outputs. This robustness means that analysts only need to specify the mean and standard deviation of the damage per missile, rather than its entire empirical form. The salvo size also had minimal impact on model performance, even with just 2 missiles per salvo.

Positive correlation among missiles did affect the model's behavior, however, especially with larger salvos. The model tended to underestimate results relative to the simulation, especially for the standard deviation and 95th percentile of loss; the estimate for the mean remained reasonably good. Correlation affected outcomes by increasing the variation and by changing the shape of the loss distribution function. The model modification that was tested herein incorporated the former effect but not the latter, and so the modified model tended to

overcompensate for correlation. Thus future research is needed to develop a more accurate way to handle positive correlation.

Based upon these results, it seems reasonable to conclude that the stochastic model will perform well as long as the correlation between missiles is negligible. Its accuracy and consistency are comparable to those of Monte Carlo simulation, and free of the latter's sampling issues. Researchers and analysts can therefore feel confident in using the model for their own studies of naval warfare.

As with any numerical study, this project's major limitation is that its results are necessarily a function of the parameter space explored. The other main caveat is that the study focused on verification; that is, on how well the model's behavior matched its design. It did not consider the more difficult issue of validation; that is, how realistically the model depicts actual missile combat. Validation is certainly an important goal, but is beyond the scope of this paper.

Acknowledgements

The author would like to thank the anonymous reviewers for their suggestions for improvements in this paper.

References

- Armstrong MJ. (2007) Effective attacks in the salvo combat model: salvo sizes and quantities of targets. *Naval Research Logistics* 54 #1, 66-77.
- Armstrong MJ (2005). A stochastic salvo model for naval surface combat. *Operations Research* 53 #5, 830-841.
- Armstrong MJ (2004). Effects of lethality on naval combat models. *Naval Research Logistics* 51 #1, 28-43.
- Armstrong MJ & Powell MB (2005). A salvo combat analysis of the Battle of the Coral Sea. *Military Operations Research* 10 #4, 27-38.
- Ghose D, Speyer JL, & Shamma JS (2002). A game theoretical multiple resource interaction approach to resource allocation in an air campaign. *Annals of Operations Research* 109 # 1, 15-40.
- Hambling D (2009). Mine vs. cannon vs. torpedo in high-speed underwater arms race. *Danger Room* blog, <http://www.wired.com/dangerroom/2009/06/mine-vs-cannon-vs-torpedo-in-high-speed-underwater-arms-race/>.
- Hughes WP (2000). *Fleet Tactics and Coastal Combat*. 2nd ed. Naval Institute Press, Annapolis.
- Hughes WP (1997). Overview. In: Hughes WP, editor. *Modeling for Decision Making*. 3rd ed. Military Operations Research Society, Alexandria.
- Hughes WP (1995). A salvo model of warships in missile combat used to evaluate their staying power. *Naval Research Logistics* 42, 267-289.
- Law AM & Kelton WD (2000). *Simulation Modeling and Analysis*. 3rd ed. McGraw Hill, New York.

- Jaiswal NK (1987). Probabilistic analysis of combat models. *Annals of Operations Research* 9 #1, 561-573
- Montgomery DC (1991). *Design and Analysis of Experiments*. 3rd ed. John Wiley, New York.
- Moore DS & McCabe GP (1993). *Introduction to the Practice of Statistics*. 2nd ed. WH Freeman & Company, New York.
- MSCO (2006). *VV&A Recommended Practices Guide*. Modeling & Simulation Coordination Office, Department of Defense, Washington. <http://vva.mscso.mil/>
- USNI (2009). News release: Chinese Develop Special "Kill Weapon" to Destroy U.S. Aircraft Carriers. 31 Mar. United States Naval Institute. <https://www.usni.org/forthemedia/ChineseKillWeapon.asp>.

Appendix

To avoid repetition, most equations below are written in a general form and describe properties of both the offense and the defense. For parameters p , q and r , insert subscript α when considering offensive missiles or z for defensive missiles, respectively. The parameter n represents the total number of missiles fired by one side; it therefore stands for An_α on offense and Bn_z on defensive.

Derivations for Equations 5, 6, & 7.

Let X_i be independent & identically distributed binary random variables that equal 1 with probability q and 0 with probability $(1 - q)$. Let Y be a binary random variable, independent of the X_i , that equals 1 with probability r and 0 with probability $(1 - r)$. Define $W_i = YX_i$ to represent individual missile successes, and $Z = YX_1 + YX_2 + \dots + YX_n$ to represent the total number of successes. These have the following properties.

$$E[W_i] = E[Y]E[X_i] = rq$$

$$\text{Var}[W_i] = E[Y^2X_i^2] - E[YX_i]E[YX_i] = rq(1-rq)$$

$$\text{Cov}[W_i, W_j] = E[Y^2X_iX_j] - E[YX_i]E[YX_j] = rq^2 - r^2q^2$$

$$\text{Cor}[W_i, W_j] = \text{Cov}[YX_i, YX_j] / \text{Var}[YX_i] = (rq^2 - r^2q^2) / (rq - r^2q^2) = (q - rq) / (1 - rq)$$

$$E[Z] = nE[W_i] = nrq \quad (\text{A1})$$

$$\begin{aligned} \text{Var}[Z] &= \text{Var}[Z] = E[Z^2] - E[Z]E[Z] = E[Y^2](nE[X_iX_i] + n(n-1)E[X_iX_i]) - (nrq)(nrq) \\ &= r(nq + n^2q^2 - nq^2) - n^2r^2q^2 = nrq(1 - q + nq - nrq) \quad (\text{A2}) \end{aligned}$$

Derivations for Proposition 1.

For part (a), first show that variation in missile success (σ_α for offense and σ_z for defense) increases as correlation increases. To do this, take the variance for correlated missiles (Equation A2) and subtract the variance for independent missiles, i.e. for a binomial distribution.

$$\begin{aligned}\Delta\text{Var} &= \text{Var}_{\text{Correlated}} - \text{Var}_{\text{Binomial}} = \text{Var}[Z] - np(1-p) \\ &= (nrq - nrq^2 + n^2rq^2 - n^2r^2q^2) - (nrq - nr^2q^2) = -nrq^2 + n^2rq^2 - n^2r^2q^2 + nr^2q^2 \\ &= nrq(-q + nq - nrq + rq) = nrq(nq - nrq - q + rq) = nrq(1-r)(nq - q) \quad (\text{A2})\end{aligned}$$

This gives the increase in variation due to correlation. It equals 0 when $r = 1$ (independence), and is positive when $r < 1$ (positive correlation). To verify how this changes as correlation increases, decrease r while holding $p = rq$ constant. Substitute $p = rq$ and $q = pr^{-1}$ into Equation A2.

$$\Delta\text{Var} = np(npr^{-1} - np - pr^{-1} + p) = np^2(nr^{-1} - n - r^{-1} + 1) \quad (\text{A3})$$

Now take the derivative with respect to r .

$$d/dr [\Delta\text{Var}] = np^2(-npr^{-2} - r^{-2}) < 0 \quad (\text{A4})$$

Since this is negative, it shows that the variance in the number of missile successes (σ_α or σ_z , as appropriate) is increasing as r decreases, and thus as correlation increases. This in turn means that $\text{Var}[Net_{AB}] = \sigma_N^2 = A\sigma_\alpha^2 + B\sigma_z^2$ overall is increasing in correlation.

That $E[Net_{AB}]$ does not change with correlation, i.e. when decreasing r while holding $p = rq$ constant, is a simple consequence of Equation A1.

For part (b), start with the expression for the variance of the nominal loss given in (Armstrong, 2005).

$$\sigma_{\text{Loss}^*}^2 = \mu_N\sigma_v^2 + \sigma_N^2\mu_v^2 - 2\sigma_v^2\mu_N G_N(0) + 2\sigma_v^2\sigma_N^2 g_N(0) \quad (\text{A5})$$

Take the derivative with respect to σ_N .

$$\begin{aligned}d/d\sigma_N [\sigma_{\text{Loss}^*}^2] &= 0 + 2\sigma_N\mu_v^2 - 2\sigma_v^2\mu_N^2\sigma_N^{-1}g_N(0) + 4\sigma_v^2\sigma_N g_N(0) + 2\sigma_v^2\sigma_N^{-1}(\mu_N^2 - \sigma_N^2)g_N(0) \\ &= 2\sigma_N\mu_v^2 + 2\sigma_v^2\sigma_N g_N(0) = 2\sigma_N(\mu_v^2 + \sigma_v^2 g_N(0)) > 0 \quad (\text{A6})\end{aligned}$$

This confirms that σ_{Loss^*} increases as σ_N increases. The increase in variance causes the normal distribution to spread out. Thus all percentiles less than the mean will be decreasing in correlation, while all those greater than the mean (including the 95th percentile) will be increasing.

For part (c), first consider the situation where $E[Net_{AB}] > 0$. Since all the percentiles of the distribution (i.e., their locations) below that mean are decreasing in correlation, it follows that the probability mass at or below 0 will increase; thus the probability of zero loss will increase.

Next consider the $E[Net_{AB}] < 0$ case. Since the percentiles of the distribution above this mean are increasing in correlation, the probability mass above 0 will increase; thus the probability of zero loss will decrease.

For part (d) first consider the situation where $E[Net_{AB}] \geq B/\mu_v$, or equivalently $\mu_v E[Net_{AB}] \geq B$; i.e., the average nominal loss is greater than the number of defending ships. The percentiles of the distribution below this mean are decreasing in correlation, so the probability mass below B will increase, and this will contribute to a reduction in the mean actual loss. Meanwhile, the percentiles above this mean are increasing in correlation; but since they already lie above the maximum possible loss B , their movement will not affect the calculated loss. Thus the net effect is a decrease in the mean actual loss.

A similar argument in reverse applies to the situation where $E[Net_{AB}] \leq 0$, i.e. the average nominal loss is negative. Increased correlation causes more probability mass to shift above 0, where it contributes to an increase in the mean actual loss. Shifts of probability mass lower below the 0 point have no effect, so the net effect is an increase in the mean actual loss.

Table 1. Fit for all 486 scenarios.

	Mean	Std Dev	95th %ile	P[Loss=0]
Regression R^2	0.995	0.958	0.978	0.989
Regression slope	0.970	0.765	0.872	1.019
Average of differences	-0.039	-0.090	-0.196	0.000
Variance of differences	0.005	0.022	0.098	0.003

Table 2. Fit for 162 scenarios with independent offensive missiles.

	Mean	Std Dev	95th %ile	P[Loss=0]
Regression R^2	1.000	0.999	1.000	0.999
Regression slope	1.000	0.987	0.990	0.997
Average of differences	0.000	-0.005	-0.018	0.004
Variance of differences	0.000	0.000	0.001	0.000

Table 3. Fit for 324 scenarios with correlated offensive missiles.

	Mean	Std Dev	95th %ile	P[Loss=0]
Regression R^2	0.993	0.969	0.989	0.992
Regression slope	0.956	0.706	0.830	1.031
Average of differences	-0.058	-0.133	-0.285	-0.002
Variance of differences	0.006	0.028	0.123	0.004

Table 4. Fit of modified model for 324 scenarios with correlated offensive missiles.

	Mean	Std Dev	95th %ile	P[Loss=0]
Regression R^2	0.999	0.993	0.997	0.998
Regression slope	1.012	1.174	1.153	1.008
Average of differences	0.020	0.089	0.244	0.013
Variance of differences	0.001	0.009	0.081	0.001

Figure 1. Distribution of losses for one scenario in the Monte Carlo simulation.

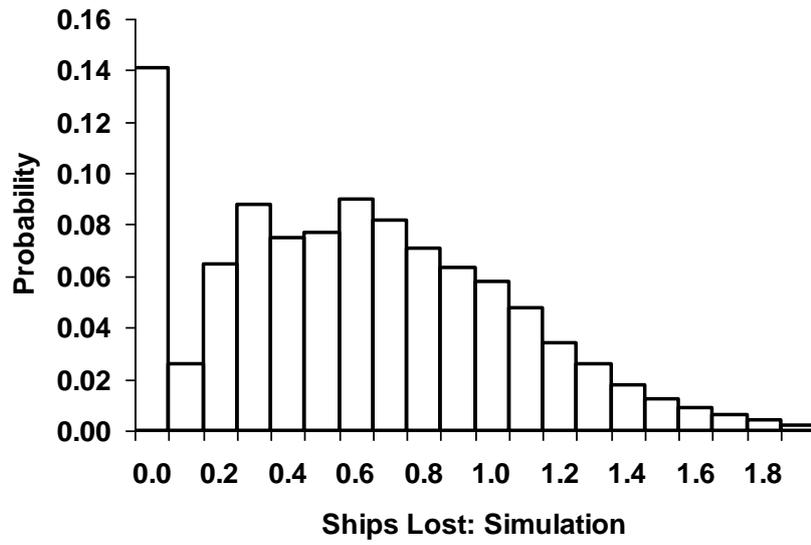


Figure 2. Distribution of losses for one scenario in the stochastic salvo combat model.

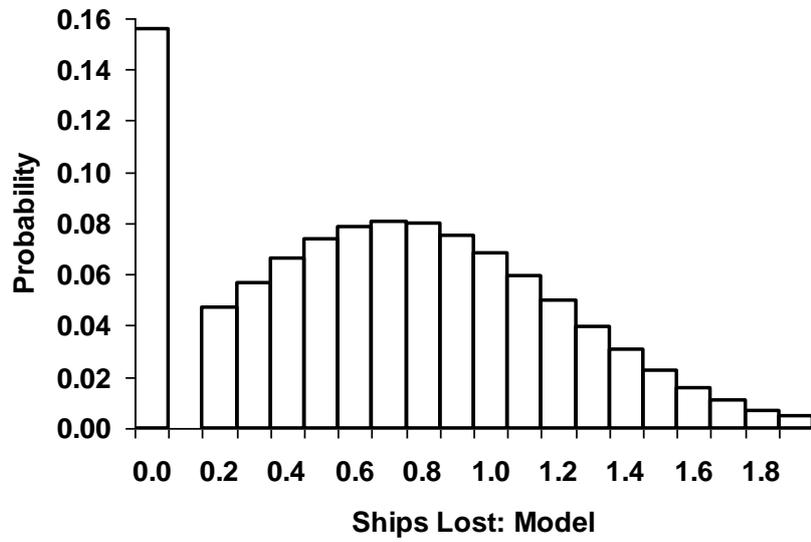


Figure 3. Scatter plot of mean loss for all 486 scenarios.

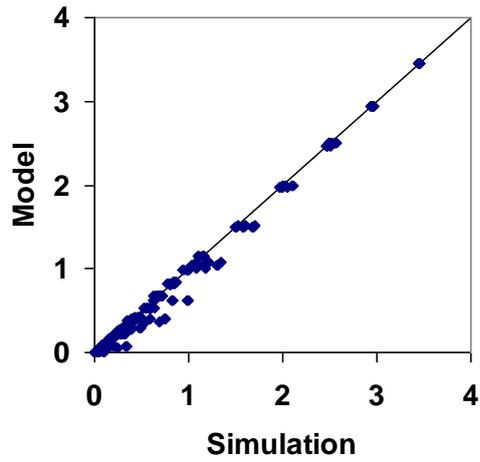


Figure 4. Scatter plot of probability of zero loss for all 486 scenarios.

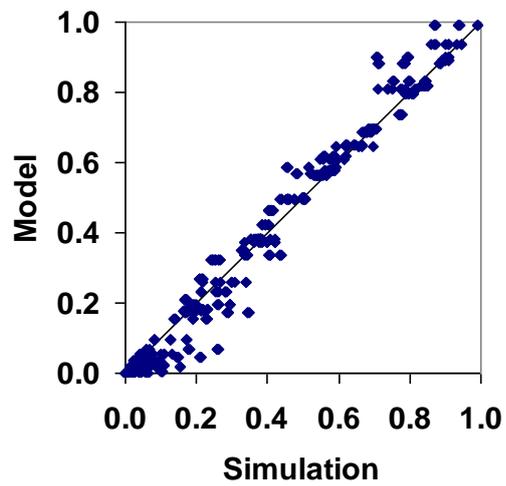


Figure 5. Scatter plot of 95th percentile of loss for all 486 scenarios.

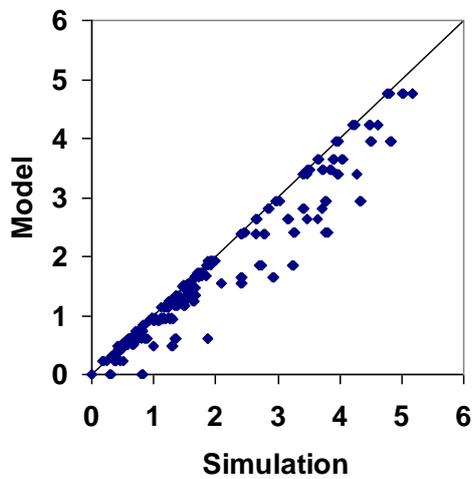


Figure 6. Scatter plot of standard deviation of loss for all 486 scenarios.

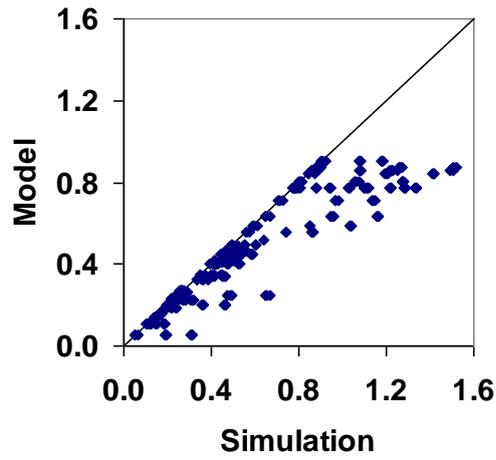


Figure 7. Plot for standard deviations of 162 scenarios with independent missiles.

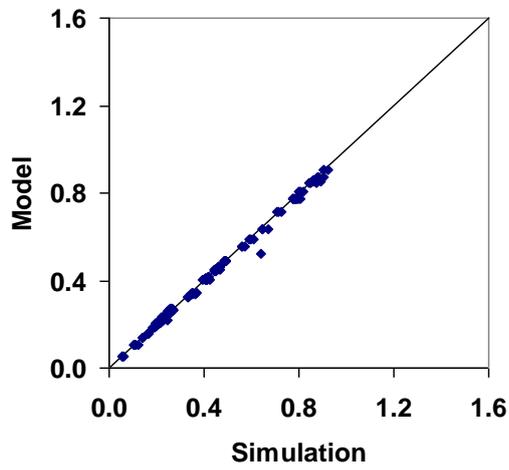


Figure 8. Plot for standard deviations of 324 scenarios with correlated missiles.

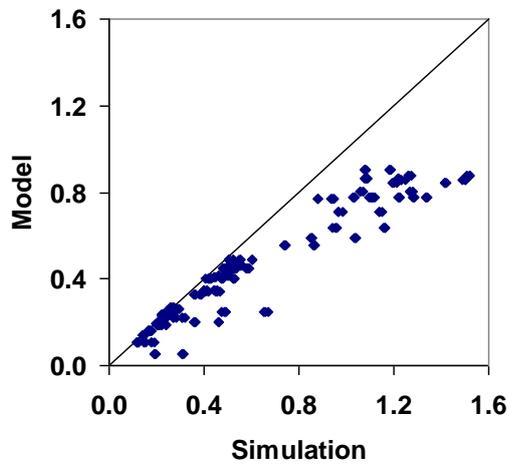


Figure 9. Plot for standard deviations in modified model of 324 scenarios with correlated missiles.

