Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

A DATA FARMING ANALYSIS OF A SIMULATION OF ARMSTRONG'S STOCHASTIC SALVO MODEL

Gökhan Kesler

Operations Division
Turkish Air Force
Çankaya, Ankara 06420, TURKEY

Thomas W. Lucas
Paul J. Sanchez

Department of Operations Research
Naval Postgraduate School
Monterey, CA 93943, USA

ABSTRACT

In 1995, Retired Navy Captain Wayne Hughes formulated a salvo model for assessing the military worth of warship capabilities in the missile age. Hughes' model is deterministic, and therefore provides no information about the distribution of outcomes that result from inherently stochastic salvo exchanges. To address this, Michael Armstrong created a stochastic salvo model by transforming some of Hughes' fixed inputs into random variables. Using approximations, Armstrong provided closed-form solutions that obtain probabilistic outcomes. This paper investigates Armstrong's stochastic salvo model using data farming. By using a sophisticated design of experiments to run a simulation at thousands of carefully selected input combinations, responses such as ship losses are formulated as readily interpretable regression and partition tree metamodels of the inputs. The speed at which the simulation runs suggests that analysts should directly use the simulation rather than resorting to approximate closed-form solutions.

1 INTRODUCTION

Nations around the world use models to assist in determining how their militaries are equipped, organized, and, if necessary, will fight. Two early such models were proposed by Lanchester (1916). Lanchester's models involve coupled differential equations that modify force levels due to differing forms of attrition. The first is his square law, where a side's rate of losses is proportional to the opposing force's numbers. This occurs when combatants can aim and coordinate their fire. The second is his linear law, where the loss rate is proportional to the product of the force levels. This happens when forces fire into areas rather than at specific targets. Lanchester's models, which have been staples for military analysts for over a century, identify conditions in which the combat strength of a force is proportional to the square of its numbers. In essence, Lanchester provides a mathematical explanation for the well-known military principle of concentration of forces. Moreover, his models are relatively simple and transparent.

Informed by Lanchester's research, Hughes (1995) formulated a new analytical model designed specifically for examining the elements of ship design in modern naval surface combat. This adopts a view that military worth is determined by maximizing "the quantity of accurately delivered lethality, or ordnance, over the combat life of the warship." Hughes' salvo equations are famous for their relative simplicity and ease of implementation. Instead of continuous attrition, he proposed using discrete salvos of missiles. In his analysis, he illustrated the interactions among four main factors on each side that are vital for analysts and decision makers—ship numbers, and their offensive, defensive, and staying powers. However, one of the limitations of Hughes' salvo model is that it is deterministic, so it provides no information about variability. This can be misleading to decision makers facing real-life problems (Lucas 2000).

To address the deterministic limitation to Hughes' salvo model, Armstrong (2005) developed an analytic stochastic salvo model (SSM) that provides outputs in terms of means, variances, and probabilities.

Armstrong's SSM replaced the fixed inputs (e.g., the amount of damage caused by offensive missiles) to Hughes' deterministic salvo model with random variables. Using distributional approximations, such as the normal approximation to the binomial, Armstrong solved his SSM to yield distributional outputs in closed-form. This allows military planners to consider distributions of possibilities rather than point values. Armstrong (2011) assessed the accuracy of his closed-form solution by comparing it to a simulation that solves the conceptual model exactly as it is specified. Over the ranges explored, the closed-form solutions perform quite well, except when the likelihood of offensive missiles hitting their targets has substantial positive correlation.

Our research avoids the distributional approximations Armstrong used by studying the exact formulation using simulation and data farming. We efficiently vary all of the factors in a simulation of Armstrong's SSM over thousands of carefully selected design points (i.e., input combinations). The output from the experiments is used to construct regression and partition tree metamodels of key responses as functions of the inputs. These metamodels are readily interpretable, and typically explain over 90 percent of the SSM's variability.

2 DETERMINISTIC AND STOCHASTIC SALVO MODELS

This section details Hughes' deterministic salvo model and Armstrong's stochastic extension of it.

2.1 Hughes' Salvo Model

Hughes' salvo equations offer a transparent, highly aggregated, and straightforward method for understanding critical ship design elements in modern naval surface combat (Hughes 1995). In discrete time increments, two opposing sides (*side A* and *side B*) simultaneously engage each other using *offensive power* with anti-ship missiles (ASMs) and defensive power with surface-to-air missiles (SAMs). Hughes also takes into consideration staying power, which is a ship's ability to keep fighting after being struck by enemy weapons. Among his findings, Hughes concludes that naval warfare can be unstable, especially when ships have weak staying power. He also concludes that numerical superiority is consistently critical in achieving an advantage. Hughes' model has also been used to explore force design, warfighting tactics, and the value of information—see Hughes (2018), Armstrong (2003), and Lucas and McGunnigle (2002). The eight inputs to Hughes' model and the primary output (losses to each side) are defined in Table 1.

Table 1: Hughes' model's variables using Armstrong's (2005) notation.

Symbol	Description
A, B	Number of available ships at the beginning of the salvo
α, β	Number of available ASMs per salvo by each ship (*offensive power*)
y, z	Number of available SAMs per salvo by each ship (*defensive power*)
w, x	Number of hits needed to put a single ship out of action (*staying power*)
$\Delta A, \Delta B$	Number of ships neutralized by the adversary in the salvo exchange

Hughes' mathematical formulation for the losses (in terms of ships) for the two sides in a simultaneous salvo exchange is presented in (1).

$$\Delta B = \frac{\alpha A - zB}{x}, \quad 0 \le \Delta B \le B \qquad \text{and} \qquad \Delta A = \frac{\beta B - yA}{w}, \quad 0 \le \Delta A \le A \tag{1}$$

In (1), ΔA and ΔB represent the number of ship losses for *side A* and *side B*, respectively, in a single salvo. If the calculations for ΔA or ΔB yield values below or above the permissible ranges, they are truncated to zero or the actual number of ships, respectively. If we remove the bounds in (1), negative values for ΔA or ΔB represent a situation in which the defensive power is more than enough to counter all incoming ASMs. We define this condition as *over-defense*. Similarly, values of ΔA or ΔB that are greater than A or

B, respectively, represent a situation in which at least one side is completely destroyed with an excess of *offensive power*. We define this situation as *over-kill*. All other situations are designated *intermediate*.

Hughes used *Fractional Exchange Ratio* (FER) to quantify the ratio of the combat power of the given forces. In layman's terms, FER is the proportion that *side B* loses divided by the proportion that *side A* loses in a salvo. FER's mathematical expression is given in (2).

$$\text{FER} = \frac{\Delta B/B}{\Delta A/A}, \quad 0 \leq \text{FER} \leq \infty \tag{2}$$

FER ranges from zero to positive infinity. Zero occurs when $\Delta B = 0$ and $\Delta A > 0$ (i.e., no *side B* ships are damaged and some *side A* ships are put out of action). Infinity occurs when $\Delta A = 0$ and $\Delta B > 0$ (i.e., no *side A* forces are lost and some *side B* forces are put out of action). When each side loses the same proportion (other than zero) of their current force level, FER equals one. Hughes considered an FER of one as equating to force parity. Values greater than one imply that *side A* has an advantage, while values less than one favor *side B*. It is worth noting that in *over-kill* situations one side can have a decisive advantage with respect to FER, but both sides suffer mutual annihilation (Hughes 1995).

2.2 Armstrong's Stochastic Salvo Model (SSM)

One of the drawbacks of Hughes' deterministic salvo model is that it provides no information on the randomness and variability that is inherent in warfare. This fact inspired Armstrong (2005) to introduce a stochastic version of the basic salvo model by representing Hughes' *offensive power*, *defensive power*, and *staying power* as random variables. Conceptually, each ASM and SAM is represented as a Bernoulli random variable (RV). Let N_o be the number of offensive missiles that impact a ship if not countered. Since the offensive missiles are assumed independent and identically distributed, N_o is a binomial RV. Similarly, if N_d is the maximum number of successful defensive intercepts, it also has a binomial distribution. Thus, the number of ASMs that hit their target is $\max(N_o - N_d, 0)$. The damage inflicted (proportion of a ship taken out of action) given an ASM hits a ship is also assumed to be random—with Armstrong assuming that the damage follows a normal distribution.

Solving Armstrong's conceptual SSM in closed-form is infeasible. However, by using the normal approximation to the binomial and assuming a normal damage function, Armstrong is able to derive formulas that approximate the expected number and variance of ships lost, as well as some probabilities (such as the probability of annihilation). To assess the accuracy of his assumptions and approximations, Armstrong (2011) compared his closed-form estimated solutions to what is obtained by simulation of the conceptual model across 486 scenarios (input combinations) using a full-factorial design of experiments consisting of six ship level combinations, three offensive missile success probabilities, three defensive missile success probabilities, three damage function shapes, and three levels of correlation. He concludes that the approximations mostly perform reasonably well. Each scenario was simulated 50,000 times in order to make the resulting estimates extremely precise.

3 DATA FARMING A SIMULATION OF ARMSTRONG'S SSM

This section describes the simulation model and provides the design of experiments we use to data farm the conceptual SSM. This includes a list of all the factors explored as well as their types and ranges. We also show that the design has excellent space-filling and correlation properties.

3.1 An R Simulation of the Conceptual Model of Armstrong's SSM

A simulation of Armstrong's SSM was implemented in the R programming environment (R Core Team 2019). The simulation is an exact implementation of the conceptual model specified in Armstrong (2005) and employed in Armstrong (2011). Both a flowchart of the simulation experiments and the R code are available in Kesler (2019). As in Armstrong (2011), 50,000 replications of the simulation are made at each design point. This ensures that the standard error associated with any probability estimate is less than 0.0023. It is worth noting that 50,000 replications of the simulation can be made in under a second on a standard desktop computer. Across all of the input combinations we explored, there is no discernible difference between the time required to calculate Armstrong's closed-form but approximate solution using a spreadsheet and the time to more accurately and precisely estimate those answers by using the true distributions in a simulation.

3.2 The Design of Experiments

In Armstrong's SSM, the four factors for missiles fired per ship from Hughes' model (α, β, y, and z) are represented with two parametric inputs apiece, as required by the binomials used to replace fixed values. In addition, the two fixed factors for losses per hit ($u = 1/w$ and $v = 1/x$) are modeled as normal distributions with means u and v and standard deviations sd_u and sd_v, respectively. The force levels, A and B, remain as fixed inputs. Thus, Armstrong's SSM has 14 factors that we wish to explore. Table 2 shows the factors, as well as the ranges over which they are varied. The 14 factors are explored over ranges similar to those used in Armstrong (2011). Generally, this focuses on battles that are not too one-sided—and hence of greater interest.

Table 2: The factors in the SSM using Armstrong's notation, and their ranges.

<table>
<tr><th>Factor</th><th>Explanation</th><th>Low</th><th>High</th><th>Type</th></tr>
<tr><td>A</td><td rowspan="2">Beginning force strength</td><td rowspan="2">9</td><td rowspan="2">18</td><td rowspan="2">Integer</td></tr>
<tr><td>B</td></tr>
<tr><td>n_α</td><td rowspan="2">Maximum number of ASMs per ship per salvo for A and B, respectively</td><td rowspan="2">2</td><td rowspan="2">4</td><td rowspan="2">Integer</td></tr>
<tr><td>n_β</td></tr>
<tr><td>p_α</td><td rowspan="2">Probability of successful strike of single ASM for A and B, respectively</td><td rowspan="2">0.5</td><td rowspan="2">1.0</td><td rowspan="2">Continuous</td></tr>
<tr><td>p_β</td></tr>
<tr><td>n_y</td><td rowspan="2">Maximum number of SAMs per ship per salvo for A and B, respectively</td><td rowspan="2">1</td><td rowspan="2">2</td><td rowspan="2">Integer</td></tr>
<tr><td>n_z</td></tr>
<tr><td>p_y</td><td rowspan="2">Probability of successful intercept for a single SAM for A and B, respectively</td><td rowspan="2">0.5</td><td rowspan="2">1.0</td><td rowspan="2">Continuous</td></tr>
<tr><td>p_z</td></tr>
<tr><td>u</td><td rowspan="2">Mean losses per hit for A and B, respectively</td><td rowspan="2">0.25</td><td rowspan="2">0.50</td><td rowspan="2">Continuous</td></tr>
<tr><td>v</td></tr>
<tr><td>sd_u</td><td rowspan="2">Standard deviation of mean losses per hit for A and B, respectively</td><td rowspan="2">0.1</td><td rowspan="2">0.2</td><td rowspan="2">Continuous</td></tr>
<tr><td>sd_v</td></tr>
</table>

The factors in the design are a combination of six integer-valued (from 2 to 10 levels) and eight continuous input variables. Since the model runs quickly, we chose a design that provides us with plenty of degrees of freedom to fit diverse complex metamodels while maintaining nearly independent coefficient estimates. Moreover, the design has good space-filling properties for the continuous factors so that we can identify thresholds and change points if they are present. The design we used is called a mixed NOB (nearly orthogonal and balanced) design, see Vieira et al. (2011). Using the 512 design point spreadsheet available at the SEED Center website (SEED Center 2019), which was rotated and stacked 10 times, we ran the SSM simulation at 5,120 design points (i.e., unique input combinations). With 50,000 replications per design point, a total of 256 million stochastic salvo battles were simulated. This took less than half an hour on a modern desktop computer.

A visual representation of the design space is shown in Figure 1 as a pairwise scatterplot of the factors. We see that the mixed NOB design fills the interior for the continuous factors of the SSM. For the discrete predictors, the design fills much like a full factorial, at least in the pairwise projections. The absolute values of the pairwise correlations among columns in the design matrix are all on the order of 0.01 or less.

Moreover, the maximum of the absolute value of the 6,786 pairwise correlations for a full second-order model—all main effects, quadratic effects (except for 2-level factors), and 2-way interactions—is 0.13, with no other pair above 0.07. Therefore, we expect minimal confounding effects for even complicated metamodels.

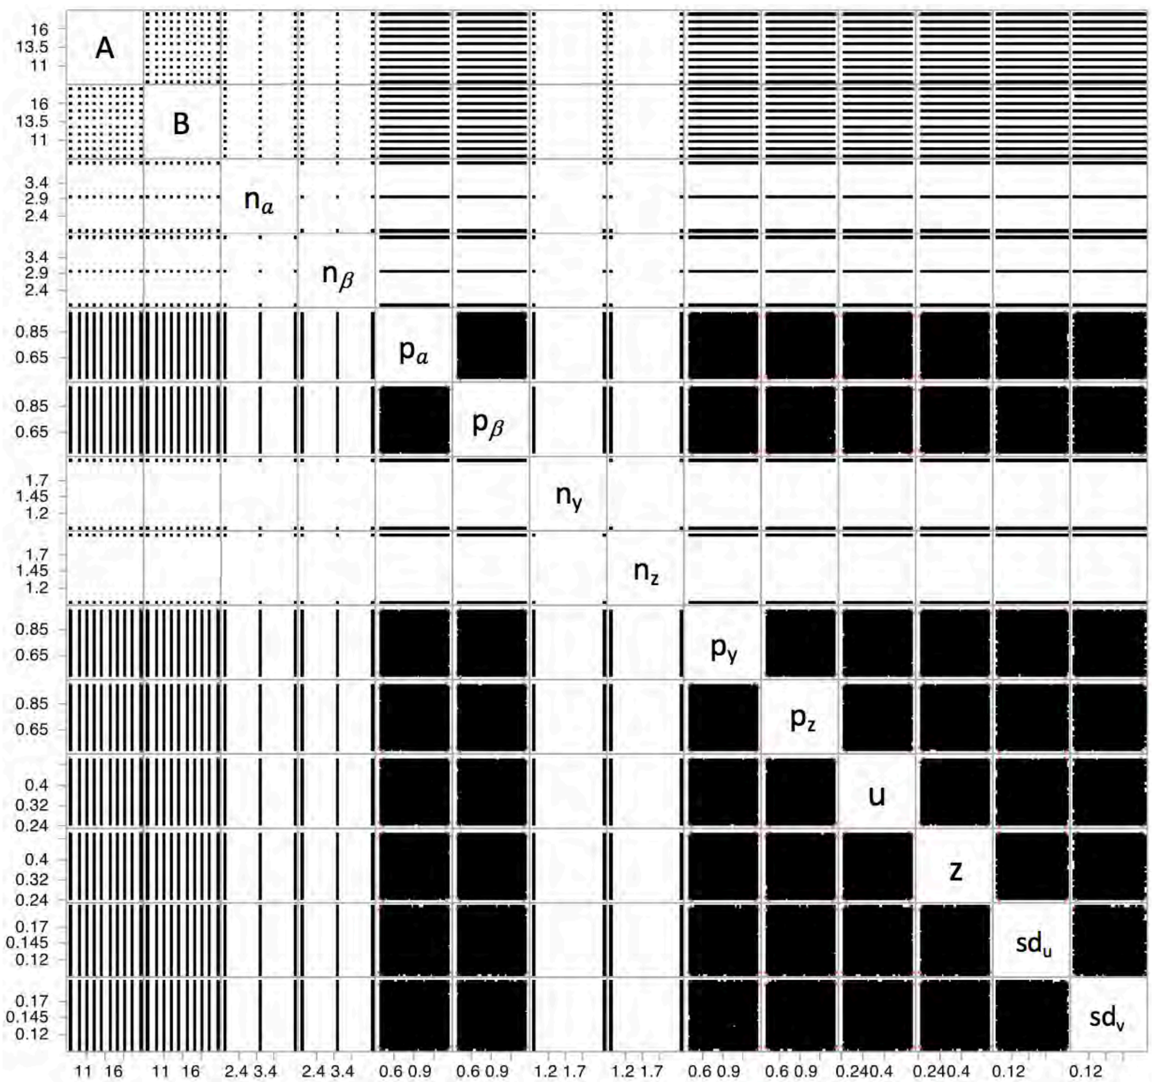

Figure 1: Pairwise scatter plot of the NOB's factors.

One of the important benefits of using a space-filling design rather than one of the so-called "optimal" design choices is that the latter are only optimal for a particular parametric form of the response surface. Unless you have solid *a priori* knowledge about the true nature of the response surface, an optimal design may fail to help you fit an appropriate model. If you are fitting metamodels for two or more outputs, it's a near certainty that their parametric representations will differ. A design chosen to be optimal for one of the outputs would likely be sub-optimal for others. While space-filling designs—such as the NOB—are not optimal, they are extremely robust in the sense that they work very well across a broad variety of response surface geometries.

4 ANALYSIS

This section uses stepwise regression and partition trees to quantify in easily understandable forms how the SSM's factors affect expected ship losses (ΔB) and a modified FER. Kesler (2019) examines additional responses, using metamodels and graphical representations. We wish to strongly emphasize that these and other insights obtained from this set of experiments cannot be guaranteed to extend beyond the regions covered by our design.

4.1 What Causes Ship Loss?

Since the SSM and our design are both symmetric with respect to the two sides, we chose to fit our models to *side B*'s losses (i.e., ΔB). In addition, for this effort we initially focus our search on the input factors themselves rather than functions of them (e.g., force ratio A/B). Figure 2 shows a partition tree generated using the R statistical programming language. Partition trees recursively bifurcate the data into subsets that minimize the sum of squares in the two resultant subgroups. All possible splits for each individual factor are considered. No distributional assumptions are required in building the tree. Moreover, interactions emerge naturally and the trees are easy to interpret.

The top node contains 100% of the 5,120 design points, and has a mean ΔB of 5.3. The first split occurs on the number of ASMs per *side A* ship. When $n_\alpha < 2.5$, which occurs in 35% of the design points, the mean ΔB is 2.6. Conversely, when $n_\alpha \geq 2.5$, which occurs in 65% of the design points, the mean ΔB is 6.8. The minimum mean ΔB in a node occurs when ASMs per *side A* ship are low ($n_\alpha < 2.5$) and SAMs per *side B* ship are high ($n_z \geq 1.5$). These conditions, summarized across the other 12 factors, yield a mean ΔB of 1.4. This happens in 18% of the design points. The maximum mean ΔB in a node occurs when ASMs per *side A* ship are high ($n_\alpha \geq 3.5$), they have a high probability of hit ($p_\alpha \geq 0.73$), and there are a high number of *side A* ships ($A \geq 13$). When this occurs, across the other 11 factors, the mean ΔB is 11. This happens in 10% of the design points.

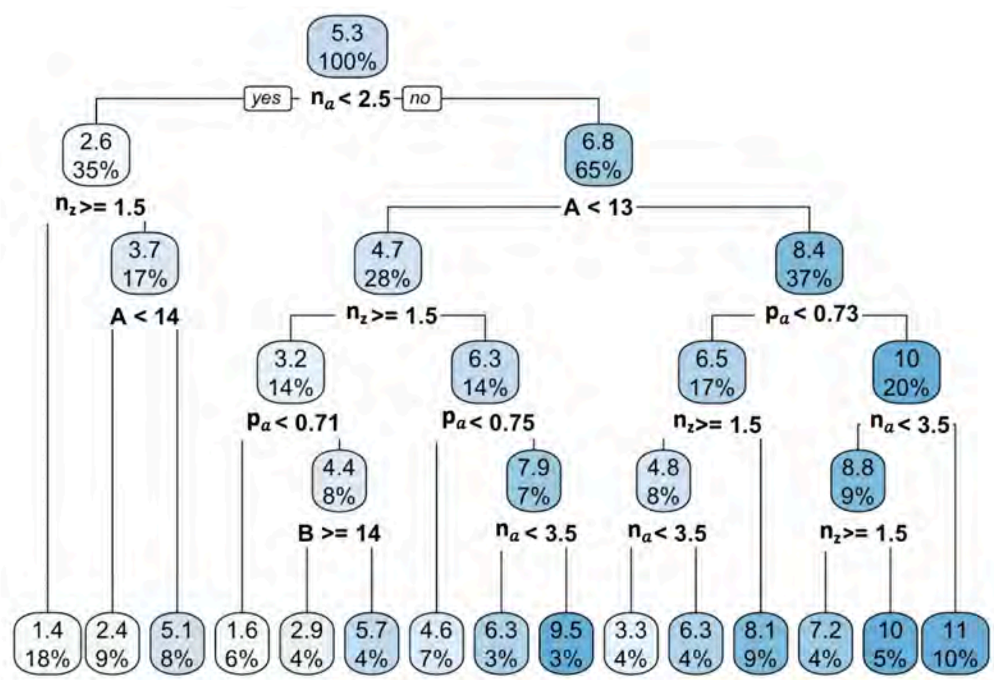

Figure 2: Partition tree for ΔB.

Several additional interesting takeaways can be gleaned by looking at the whole of Figure 2. First and foremost is that the factors at the top of the tree are predominantly related to *side A*'s *offensive power* (A, the force level; n_a, the offensive missile salvo size; and p_a, *side A*'s ASM probability of hitting a target when unopposed). The number of targets B is included in only two nodes at the bottom of the tree which comprise just 8% of the design points. The effectiveness of *side B*'s SAMs does not appear in the tree, nor does either of its two damage function parameters.

The variable importance heuristic from R's `rpart` function is displayed in Table 3. This reaffirms our observation that the dominant factors determining *side B*'s losses are those associated with the lethality of *side A*. We also evaluated ΔB using a bootstrap forest. This approach bootstraps many sampled subsets of the data to avoid overfitting (which is often a temptation with partition trees) and the potential masking of variables in splits. The rankings from both techniques are similar, so we restrict our presentation to the partition tree.

Table 3: Variable importance ranking from R's `rpart` analysis.

n_α	A	p_α	n_z	B	sd_u	u	p_z	v	p_y	p_β	sd_v	n_β	n_y
24,004	12,909	9,361	9,245	1,483	746	682	597	574	455	400	322	142	128

The most commonly used method for constructing metamodels is to fit a regression model of a response on the factors (Kleijnen et al. 2005). Furthermore, fitting smooth regression curves to the data nicely complements partition trees. Of course, there is much art to building informative and parsimonious metamodels. We started by using stepwise regression with the default settings in JMP (SAS Institute Inc. 2019) to construct up to a full second-order model—which produces a model with 31 terms and an R-square of 0.927. That is, nearly 93% of the SSM simulation's variability is explained by the regression fit.

To simplify the metamodel, we manually removed terms from this initial fit and ended up with the 10-term model shown in Figure 3. We chose to retain 10 terms because there is a large gap in the magnitudes of the t-Ratios after that point. This parsimonious model has an R-square of 0.903, only slightly lower than the 31-term model. Note that large-scale designs provide a large number of degrees of freedom, and tend to yield regression models with far too many terms. To enhance understanding without noticeably compromising explanatory power, practical significance should supersede statistical significance in model selection.

Term	Estimate	Std Error	t Ratio	Prob>|t|
n_α	2.8002825	0.02133	131.28	<.0001*
A	0.6024809	0.006029	99.93	<.0001*
p_α	10.636594	0.120692	88.13	<.0001*
n_z	-2.786162	0.034908	-79.81	<.0001*
p_z	-5.508274	0.120734	-45.62	<.0001*
v	10.725175	0.241416	44.43	<.0001*
B	-0.185127	0.00603	-30.70	<.0001*
(A-13.3559)*(n_α-2.97559)	0.1508624	0.007394	20.40	<.0001*
(n_α-2.97559)*(p_α-0.75)	2.744932	0.148112	18.53	<.0001*
(B-13.3598)*(n_z-1.49629)	-0.221891	0.012065	-18.39	<.0001*

Figure 3: Sorted parameter estimates for a parsimonious regression metamodel on ΔB.

The metamodel in Figure 3 tells a story that reinforces the main findings from the partition tree analysis. Throughout the following discussion, we focus on main effects interpretation. The reader should bear in mind that these effects are modified by interactions. Note that the regressions in this paper were done using JMP software (SAS Institute Inc. 2019), with their default practice of centering interaction and quadratic terms at their means—hence the numeric constants subtracted from these terms.

The most significant factor once again is n_α—the number of ASMs per salvo per ship on *side A*. Each additional ASM per salvo results in an expected increase of 2.8 ships in ΔB. The three most significant terms are those associated with *side A*'s offensive power (n_α, A, and p_α). Each additional ship on side A results in an expected 0.6 more ship losses for *side B*. Every 0.1 increase in *side A*'s ASM probability of hitting a *side B* ship increases *side B*'s expected ship losses by about one ship. *Side A*'s *offensive power* factors are followed in significance by four terms related to *side B*'s *defensive power* and *staying power*. These also have readily interpretable meanings. For example, each additional SAM that a *side B* ship can launch in a salvo results in an expected reduction of 2.8 ships in ΔB. The two most significant interactions relate to additive effects of *side A*'s offensive parameters. The final term shows an increasing benefit to *side B* when more SAMs are available. Taken together, these results from our data farming analysis reinforce the primary tactical advice from Hughes (2018) to "attack effectively first." This requires a force with a scouting advantage and sufficient offensive power.

While Armstrong's SSM is reasonably well characterized by the 10-term metamodel, the residuals show that the model suffers in the tails, i.e., when there is an *over-kill* or *under-kill situation*. Much better fits are obtained for the subset of battles that exclude *over-kill* or *under-kill* conditions. These cases amount to 65% of our design points. A stepwise regression fit using the Bayesian Information Criterion (BIC) with default settings in JMP yields a 27-term metamodel with an R-square value of 0.995. For more information about this approach of fitting different regression metamodels for the different kill categories—*over-kill*, *under-kill*, and *intermediate*), see Kesler (2019).

4.2 What Drives FER?

FER is the measure that Hughes (1995) focused on, so we felt that it would be worthwhile to explore it using the data farming approach as well. As mentioned in Section 3.2, the use of a space-filling design enabled this without the need for additional designs or simulation runs.

Initial regression models for FER showed a distinct S-shape in the actual versus predicted diagnostic plots. Because of this, we performed a logit transformation of the response. Inspection of (2) shows that we will end up with infinity in the numerator or denominator if B or A are zero, respectively. To avoid this, we added 0.01 to both the numerator and denominator for each design point. Doing so results in the adjusted FER's range becoming $[0.0099, 101]$. This slight transformation changes the outcomes minimally but avoids the infinite extremes. We then rescaled by the maximum FER + 0.01 to yield results in the range (0,1), where the smallest scaled FERs are close to 0 and the largest are close to 1. This enabled us to apply the logit transformation to all data and capture the S-shaped behavior noted in our initial multiple regression model. An initial model was fit using stepwise regression, and subsequently reduced by dropping all terms with t-ratios less than 10.

As can be seen in Figure 4, the scaled and transformed model does an excellent job capturing the tails, where the bulk of the data reside. The greatest variability occurs surrounding the range where FERs transition fairly rapidly from low to high. The model achieves a good fit as indicated by an R-square value of 0.90. Figure 5 shows that this model is not as parsimonious as the one for ΔB, but the lowest t-ratio in our model has a magnitude greater than 11. We were unable to reduce the model further without significant impact on the R-square. Figure 6 shows that all remaining factors have significant slopes, many (*ForceRatio*, n_b, p_b, n_y, and p_y) with non-linear impact on FER. Additionally, there are significant interactions amongst the same set of factors which manifest the non-linear behaviors, as can be seen in Figure 7.

While the coefficients in Figure 4 are harder to interpret for the transformed and scaled FER response, we see that the initial force ratio (i.e., A/B) is by far the most significant (and nonlinear) predictor. This relates directly to one of Hughes (1995) main findings: "[n]umerical superiority is the force attribute that is consistently most advantageous." The next set of force attributes that appear in the sorted estimates list relate to *offensive power* and *defensive power*, and they do so in natural ways—e.g., more *offensive power* for A increases the adjusted and scaled FER. Of less importance are the damage function (or *staying*

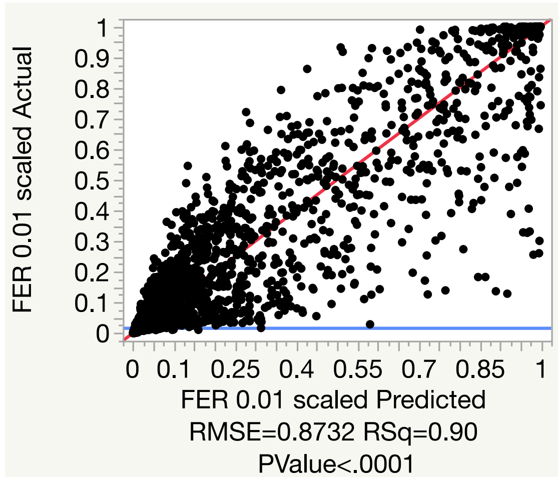

Figure 4: Actual versus predicted plot for logit transform fit of FER.

Term	Estimate	Std Error	t Ratio	Prob>|t|
ForceRatio	5.2583344	0.036587	143.72	<.0001*
n_β	-1.124698	0.014978	-75.09	<.0001*
n_y	1.6010959	0.024418	65.57	<.0001*
n_α	0.95086	0.014916	63.75	<.0001*
p_β	-4.698272	0.084421	-55.65	<.0001*
n_z	-1.155887	0.024423	-47.33	<.0001*
p_α	3.6510859	0.084492	43.21	<.0001*
p_y	3.0021488	0.084419	35.56	<.0001*
(*ForceRatio*-1.05019)*(n_y -1.51055)	2.3862992	0.073289	32.56	<.0001*
p_z	-2.381414	0.084464	-28.19	<.0001*
(n_β -3.00664)*(n_y -1.51055)	-0.823109	0.029974	-27.46	<.0001*
(*ForceRatio* -1.05019)*(n_β -3.00664)	-1.011372	0.045005	-22.47	<.0001*
(n_y -1.51055)*(p_y -0.75)	3.0505322	0.169035	18.05	<.0001*
v	3.0009759	0.16883	17.78	<.0001*
(p_β -0.75)*(n_y -1.51055)	-2.912014	0.169072	-17.22	<.0001*
(*ForceRatio*-1.05019)*(p_y -0.75)	4.3188227	0.25432	16.98	<.0001*
(n_β -3.00664)*(n_β -3.00664)	0.3712101	0.025862	14.35	<.0001*
(*ForceRatio*-1.05019)*(p_β -0.75)	-3.49983	0.251308	-13.93	<.0001*
(n_β -3.00664)*(p_y -0.75)	-1.29493	0.104152	-12.43	<.0001*
u	-1.966487	0.168883	-11.64	<.0001*

Figure 5: Sorted parameter estimates for logit transform fit of FER.

power) parameters. In fact, the standard deviations do not make it into the model. These insights can also be visually gleaned from Figure 6.

The visual depictions of the four most significant interactions in Figure 7 also provide insight. The most important *side B* attributes in interaction terms relate to its offensive missile capabilities (n_β and p_β), while for *side A* it is its defensive missile capabilities (n_y and p_y). The most pronounced interactions involve the initial force ratio (A/B). In all of the interactions involving force ratio, the effect of the missile attributes is much more pronounced at high force ratios. For example, increasing the number of *side B*'s SAMs (n_β) has no discernible impact on our adjusted and scaled FER at low force ratios, but results in a strong reduction at high force ratios.

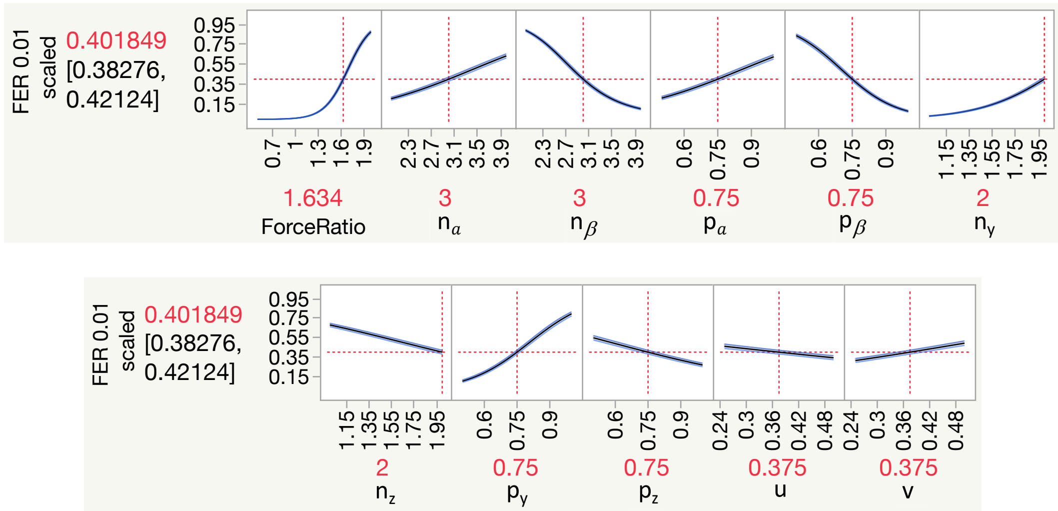

Figure 6: Prediction profiler shows non-linearity of the response surface for logit transform fit of FER.

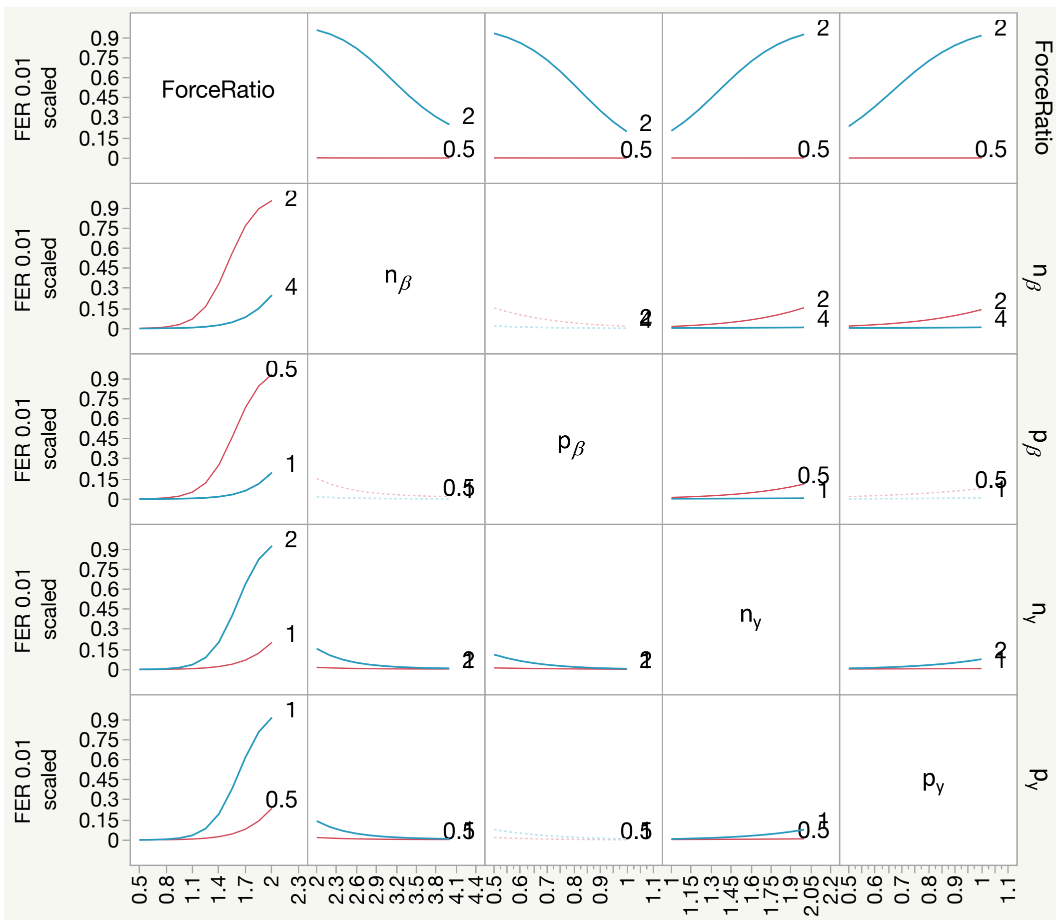

Figure 7: The initial ForceRatio dominates, but its impact is significantly affected by secondary factors in the logit transform fit of FER.

5 CONCLUSIONS

Hughes (1995), Armstrong (2005), and many others have created models to help researchers better understand complex systems or phenomena. By Department of Defense standards, Hughes' and Armstrong's models are exceedingly simple, respectively having 8 and 14 inputs. However, when trying to extract insights even this degree of dimensionality can be challenging. For example, with 14 inputs there are potentially 91 two-way, 364 three-way, and over 1000 four-way interactions. In practice, many of the insights about models are obtained graphically by fixing all but one or two of the inputs and seeing how responses vary as the selected factors are varied. This can be informative, but requires the analyst to pick out the right subsets from the set of all inputs. It also requires the results to be invariant to the inputs that are being held fixed. If the fixed inputs are involved in significant interactions, conclusions based on this approach may be misleading in other contexts.

Data farming is a holistic approach. It lets us quantify the relationships between numerous inputs and multiple responses using efficient designs that enable us to fit a diverse set of metamodels involving numerous factors. By using partition trees and regression models, for example, we are able to identify the driving factors, determine key interactions, find thresholds or change-points, and construct readily interpretable metamodels whose component contributions can be understood by decision makers. In short, we develop an understanding of the model that is unobtainable by looking at low-dimensional subsets. For this study, over the ranges explored, we are able to identify and quantify how all 14 factors in Armstrong's SSM impact losses and FER.

The data farming approach is even more valuable when the dimensionality—and hence the complexity—of our models increases. In fact, most Navy and other services models are simulations with many hundreds or thousands of inputs, and thus potential factors. The intent of those models is to allow analysts to explicitly study many of the details that are subsumed in the aggregated inputs to Hughes' equations. For models with so many input factors, the ability to reveal the relationships between inputs and outputs in readily comprehensible forms is essential.

Technology is on the side of the data farmer. Our ability to data farm continues to advance with increasing computational power, better design of experiments, and improving data analysis methods. Technology is also on the side of the simulator, who can now use simulation to solve conceptual models to whatever precision is desired rather than having to rely on unrealistic simplifications, approximations, and asymptotic limiting behaviors. As Lucas et al. (2015) advocate, the combination of simulation and data farming should now be viewed as a method of first resort.

ACKNOWLEDGMENTS

The authors would like to thank Captain Wayne Hughes, USN (Ret.) for his tremendous contributions to operations research and naval combat modeling. This study is rooted in the simple, elegant, and useful model Hughes created that has influenced navies around the world. His greatest legacy, however, is his impact on thousands of naval officers, modelers, and analysts—including us.

REFERENCES

Armstrong, M. J. 2003. "Effects of Lethality in Naval Combat Models". *Naval Research Logistics* 51(1):28–43.

Armstrong, M. J. 2005. "A Stochastic Salvo Model for Naval Surface Combat". *Operations Research* 53(5):830–841.

Armstrong, M. J. 2011. "A Verification Study of the Stochastic Salvo Combat Model". *Annals of Operations Research* 186(1):23–38.

Hughes, W. P. 1995. "A Salvo Model of Warships in Missile Combat Used to Evaluate Their Staying Power". *Naval Research Logistics* 42(2):267–289.

Hughes, W. P. 2018. *Fleet Tactics and Naval Operations*. 3^{rd} ed. Annapolis: Naval Institute Press.

Kesler, G. 2019. "A Data Farming Analysis of Multiple Salvo Equations". Master's thesis, Operations Research Department, Naval Postgraduate School, Monterey, California.

Kleijnen, J. P. C., S. M. Sanchez, T. W. Lucas, and T. M. Cioppa. 2005. "A User's Guide to the Brave New World of Designing Simulation Experiments". *INFORMS Journal on Computing* 17(3):263–289.

Lanchester, F. W. 1916. *Aircraft in Warfare: The Dawn of the Fourth Arm*. London: Constable & Co.

Lucas, T. W. 2000. "The Stochastic Versus Deterministic Argument for Combat Simulations: Tales of When the Average Won't Do!". *Military Operations Research* 5(3):9–28.

Lucas, T. W., W. D. Kelton, S. M. Sanchez, P. J. Sanchez, and B. L. Anderson. 2015. "Changing the Paradigm: Simulation, Now a Method of First Resort". *Naval Research Logistics* 62(4):293–303.

Lucas, T. W., and J. E. McGunnigle. 2002. "When is Model Complexity Too much? Illustrating the Benefits of Simple Models with Hughes' Salvo Equations". *Naval Research Logistics* 50(3):197–217.

R Core Team 2019. "R: A Language and Environment for Statistical Computing". https://www.r-project.org.

SAS Institute Inc. 2019. *JMP® Version 14*. Cary, NC, USA.

SEED Center 2019. "SEED Center for Data Farming". https://harvest.nps.edu/.

Vieira, Jr., H., S. M. Sanchez, K. H. Kienitz, and M. C. N. Belderrain. 2011. "Improved Efficient, Nearly Orthogonal, Nearly Balanced Mixed Designs". In *Proceedings of the Winter Simulation Conference*, edited by S. Jain, R. R. Creasey Jr., J. Himmelspach, K. P. White, and M. C. Fu, 3605–3616. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

GÖKHAN KESLER is an Operations Research Analyst and a First Lieutenant at the Turkish Air Force (TuAF) Command Headquarter in Ankara Turkey. He has a BS in Industrial Engineering from Turkish Air Force Academy (TuAFA), and MS in Operations Research from Naval Postgraduate School (NPS). He works on requirements and planing of air-delivered munitions and is a member of Aerospace Capability Group 2 (ACG2) on Effective Engagement in NATO. His email address is gkesler@hvkk.tsk.tr. His website is http://www.gokhankesler.com/.

THOMAS W. LUCAS is a Professor of Operations Research at the Naval Postgraduate School (NPS) in Monterey California. His primary research interests are simulation, warfare analysis, computational design of experiments, and robust Bayesian statistics. In addition to his regular faculty duties, Professor Lucas is a co-director of NPS's Simulation Experiments & Efficient Designs (SEED) Center for Data Farming—see https://harvest.nps.edu/. His email address is twlucas@nps.edu. His website is http://faculty.nps.edu/twlucas/.

PAUL J. SANCHEZ is a Research Associate Professor of Operations Research at the Naval Postgraduate School, and a member of the Simulation Experiments & Efficient Designs (SEED) Center for Data Farming. He has an SB in Economics from MIT, and MS and PhD degrees in Operations Research from Cornell University. His research interests include design of experiments, simulation output analysis, and object-oriented modeling, and he actually enjoys programming. His email address is pjsanche@nps.edu.