2004

Fitting Lanchester Equations to the Battles of Kursk and Ardennes

Lucas, Thomas W.; Turkes, Turker

Wiley Periodicals, Inc.

Naval Research Logistics, Volume 51, pp. 95â 116, 2004
https://hdl.handle.net/10945/44169

Fitting Lanchester Equations to the Battles of Kursk and Ardennes

Thomas W. Lucas,[1] **Turker Turkes**[2]

[1] *Operations Research Department, Naval Postgraduate School, Monterey, California 93943*

[2] *Turkish Army, Ankara, Turkey*

Received 12 June 2001; revised 3 February 2003; accepted 6 June 2003

DOI 10.1002/nav.10101

Abstract: Lanchester equations and their extensions are widely used to calculate attrition in models of warfare. This paper examines how Lanchester models fit detailed daily data on the battles of Kursk and Ardennes. The data on Kursk, often called the greatest tank battle in history, was only recently made available. A new approach is used to find the optimal parameter values and gain an understanding of how well various parameter combinations explain the battles. It turns out that a variety of Lanchester models fit the data about as well. This explains why previous studies on Ardennes, using different minimization techniques and data formulations, have found disparate optimal fits. We also find that none of the basic Lanchester laws (i.e., square, linear, and logarithmic) fit the data particularly well or consistently perform better than the others. This means that it does not matter which of these laws you use, for with the right coefficients you will get about the same result. Furthermore, no constant attrition coefficient Lanchester law fits very well. The failure to find a good-fitting Lanchester model suggests that it may be beneficial to look for new ways to model highly aggregated attrition. © 2003 Wiley Periodicals, Inc. Naval Research Logistics 51: 95–116, 2004.

Keywords: Lanchester equations; Battle of Kursk; combat models; attrition; model validation

1. INTRODUCTION

Combat models provide information that assists decision-makers in making and justifying decisions that involve the expenditure of billions of dollars and impact many lives. For example, the simulation Concepts Evaluation Model (CEM) was used to give senior Army leadership insight into potential courses of action in the planning of Desert Storm [1]. Attrition plays a pivotal role in most combat models, particularly campaign-level simulations, such as CEM. The attrition in CEM, and many other combat models, is based on extensions of the theory developed by Lanchester [11] (see also Osipov [14] and Taylor [17]). Due to a dearth of data, particularly

Correspondence to: T. W. Lucas (twlucas@nps.navy.mil)

two-sided, time-phased data, the validity of Lanchester equations as a model of aggregate attrition remains in question. This research examines how well Lanchester equations fit the only detailed two-sided, time-phased combat data available. In particular, we look at how well Lanchester models describe the battles of Kursk and Ardennes. Moreover, we test whether any one of the common variants of Lanchester's equations fits the battles better than the others.

Lanchester hypothesized that force levels in combat could be characterized by a coupled set of differential equations. A generalized version of Lanchester equations, as described by Bracken [2], is

$$\dot{B}(t) = a(d \text{ or } 1/d)R(t)^p B(t)^q, \tag{1}$$

$$\dot{R}(t) = b(1/d \text{ or } d)B(t)^p R(t)^q, \tag{2}$$

where

$B(t)$ and $R(t)$ are the Soviet (Blue) and German (Red) force levels at time t,
$\dot{B}(t)$ and $\dot{R}(t)$ are the rates at which the Soviet and German force levels are changing at time t,
a and b are constant attrition-rate parameters,
d is a tactical parameter that adjusts the attrition to the defender by a factor of (d) and the attacker by a factor of $(1/d)$,
p is the exponent parameter of the attacking force,
q is the exponent parameter of the defending force.

This generalized Lanchester model has five parameters (a, b, d, p, and q). The model begins at the start of the battle ($t = 0$) with initial force sizes $B(0)$ and $R(0)$. When solved numerically, the force levels are incrementally decreased according to the equations $B(t + \Delta t) = B(t) - \Delta t\, \dot{B}(t)$ and $R(t + \Delta t) = R(t) - \Delta t\, \dot{R}(t)$, until the battle ends. For us, the time step Δt is set equal to one day (the data resolution) throughout the analysis.

Lanchester studied two versions of these equations. The condition $p = q = 1$ (or, more generally, when $p - q = 0$) yields Lanchester's *linear law.* Here, the state equation relating force levels at time t is $b(B(0) - B(t)) = a(R(0) - R(t))$. Lanchester viewed this as a description of combat under "ancient conditions." That is, these are the equations that would result from a series of one-on-one duels between homogeneous forces. The linear law is also considered a good model for indirect fire (sometimes called area-fire) weapons, such as artillery (see [17]). Lanchester contrasted this with combat under "modern conditions," which is defined as $p = 1$, $q = 0$ (or, more generally, $p - q = 1$). The state equation for modern combat is $b(B(0)^2 - B(t)^2) = a(R(0)^2 - R(t)^2)$. This formulation is known as Lanchester's *square law.* Lanchester showed the added importance of force size (relative to force quality) if attrition follows the square law. Specifically, Lanchester used his simple models to show "the principle of concentration" in modern combat. A third version, with $p = 0$, $q = 1$ (or, more generally, $q - p = 1$), is called the *logarithmic law* (see Peterson [15]). This formulation suggests that a force's attrition is a function of its force size, rather than the opponent's. The logarithmic law is sometimes used to model losses due to equipment failure, desertion, disease, and other nonbattle losses.

There are important consequences for force structure, training, and battlefield tactics if one of the Lanchester laws turns out to be a good model of aggregate attrition. Specifically, in the linear law, ability (as measured by the attrition parameters a and b) is as important as numbers. Some have speculated that the linear law may turn out to be the "new modern conditions." In particular, if a force can orchestrate the battle (perhaps by using information superiority and agile maneuvers), such that engagements are typically one-on-one, the force may be able to trade quantity for quality.

There are few detailed two-sided, time-phased databases of historical battles. Time-phased data are needed to estimate the five parameters for a battle because before- and after-battle data (i.e., two time periods) result in an overdetermined system of equations, with an infinite number of solutions. Therefore, for any p and q, we can find a and b that fit perfectly. Thus, as Hartley and Hembold [9] write: "Unless we are able to procure [two-sided, time-phased data] we will not be able to validate the homogeneous square law (or any other attrition law)." Fortunately, thanks to the Center for Army Analysis (CAA) and the Dupuy Institute, detailed time-phased data on the World War II battles of Kursk and Ardennes have recently become available (see [3] and [5]).

The goal of this research is to enhance our understanding of highly aggregated attrition by studying how homogeneous generalized Lanchester models fit the time-phased data on the battles of Kursk and Ardennes. In particular, we want to see if any of the basic laws (square, linear, or logarithmic) stand out. A better understanding of aggregate attrition offers the promise of enhancing the utility of our campaign-level simulations. We emphasize the battle of Kursk, as there have been several previous analyses of the Ardennes campaign. Section 2 reviews some of the very few Lanchester attrition studies using time-phased data. A brief overview of the battle of Kursk is provided in Section 3. Section 4 presents and discusses the data on the battle of Kursk, and introduces a new method for looking at two-sided, time-phased combat data. In Section 5, this method is applied to the battle of Kursk data. Section 6 applies the method to the data on the Ardennes campaign and explains why previous researchers came to different conclusions. Section 7 examines some excursions; in particular, the effects of breaking the battle of Kursk into multiple phases, assessing differences, and fitting only the manpower data. The final section summarizes the key findings.

2. PREVIOUS RESEARCH ON TIME-PHASED COMBAT DATA

This section reviews the previous studies on time-phased combat data that motivated this research. Empirical quantitative validation studies of Lanchester equations involving time-phased data include: Bracken [2], Fricker [8], and Wiper, Pettit, and Young [20], on the Ardennes campaign; Clemens [4] on the battle of Kursk; Hartley and Helmbold [9] on the Inchon–Seoul campaign; and Engle on the battle of Iwo Jima [7]. The last two had reliable time-phased data for only one side. The previous works on the Ardennes data are discussed in some detail because different authors, using the same data, found diverse parameter estimates.

Bracken [2] formulated four different models for the Ardennes campaign using Eqs. (1) and (2). By means of a constrained grid search, he estimated the parameters (a, b, d, p, q) for the first 10 days of the Ardennes campaign with and without the defensive parameter (d) for combat forces and for total forces. Among his conclusions, Bracken found that "the Lanchester linear equation fits the [Ardennes] campaign" and there is an "attacker advantage."

Fricker [8] extended Bracken's analysis of the Ardennes campaign by: (1) using linear regression to fit the data from the entire campaign (32 days) to logarithmically transformed versions of Bracken's generalized Lanchester equations; (2) using air-sortie data; and (3) restructuring the data "to estimate initial force sizes that reflect all of the troops that eventually fought in the campaign and then subtract the casualty attrition from this total on a daily basis."

Fricker found that neither the linear nor the square law fit well. He concludes that a force's losses were more a function of its own forces than of its opponent's forces, similar to the logarithmic law. Fricker's best fits occur with exponent parameter values of q that are quite large—unbelievably large, in fact. For example, one of his models yields a q of 4.6 and a p of 0. This suggests that a force could reduce its casualty rate by a factor of 24 by sending half of its troops home, without affecting the opponent's casualty rate. Of course, this extrapolates way beyond the range of the data. Note: Fricker also fit Bracken's data by regression techniques, obtaining estimates of p and q, for the combat forces with the tactical parameter, of .43 and −.50, respectively.

Wiper, Pettit, and Young [20] used Bayesian methods to reexamine Bracken's and Fricker's Ardennes data. Their model is more general in that it uses two defensive parameters, a separate one for each side. Using Gibbs sampling, they estimate several posterior quantities of interest, in particular the means of p and q. Normal distributions are used to quantify their prior on the parameters (note: they have an extra defensive parameter and additional parameters for the variances) and the residual error. The prior on p and q is an informative one, specifically, (p, q) is distributed as circular normal with mean (.5, .5) and variance $\sigma^2 = .25$. This prior, which has over a .9 probability of (p, q) both being between 0 and 1, is consistent with the conventional wisdom of plausible values. Moreover, the prior places equal *a priori* probability on the logarithmic, linear, and square laws. The most probable prior values of p and q are the prior mean of (.5, .5). This prior also heavily discounts p and q values larger than 1.5, with an *a priori* probability of about 1 in 2000 of both p and q being greater than 1.5. Using the 10 days of combat forces data that Bracken used, Wiper et al.'s estimated posterior mean on (p, q) is (.51, .48)—almost equal to the prior mean. Using 32 days of Bracken's combat forces data, their estimated posterior mean on (p, q) is (.27, .84). Furthermore, using Bayes factors, they conclude that "the logarithmic law fits best [and] the linear laws cannot be rejected, but the square law does seem implausible." Using Fricker's data, and the same prior, the estimated posterior mean on (p, q) is (.69, 1.06). When the variance of the prior on (p, q) is inflated by a factor of 200 (i.e., a much less informative prior, with $\sigma^2 = 50$), the posterior mean on (p, q) shifts to (−0.08, 4.88)—close to what Fricker found. This shows that the conclusions depend critically on the sharpness of the prior, as the mean of the original prior and the mode of the likelihood are about 10 prior standard deviations ($\sigma = .5$) apart.

Clemens [4] fit Eqs. (1) and (2), without the tactical parameter d, to the battle of Kursk. He formatted the data as Bracken did and used two different estimation techniques: (1) linear regression on logarithmically transformed equations, and (2) a nonlinear fit to the original equations, using a numerical Newton-Raphson algorithm. Clemens also found that none of the Lanchester square, linear, or logarithmic laws fit very well. Furthermore, the different estimation techniques gave dramatically divergent estimates of the parameters. Using Newton-Raphson, his estimates were $p = 0$ and $q = 1.62$, while the linear regression gave estimates of $p = 5.32$ and $q = 3.63$.

It is interesting to look across the breadth of Bracken's [2], Fricker's [8], Wiper et al.'s [20], and Clemens' [4] findings. Figure 1 shows a scatter plot of the best-fitting p and q values that these authors found. The point here is to address the whole of the findings rather than the specifics of the individual models. The only clear pattern that emerges from Figure 1 is that the estimates of p and q appear extremely sensitive to the different formulations and assumptions. Note: While Bracken's fits are all close to the linear law, this is primarily due to his constrained grid search, which focused on parameters close to the linear law.

A couple of additional time-phased studies deserve mention. Hartley and Helmbold [9] tested Lanchester's square law using daily manpower numbers from the Inchon-Seoul campaign. They

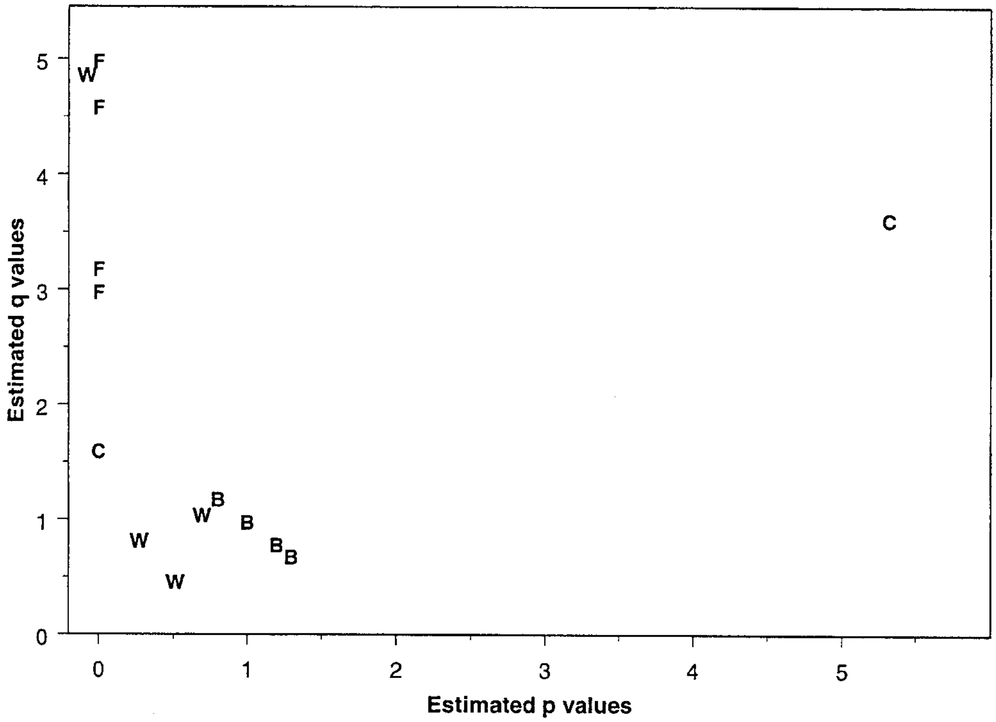

Figure 1. Bracken's (B) [2], Fricker's (F) [8], Wiper, Pettit, and Young's (W) [20], and Clemens' (C) [4] estimates of p and q. The parameters range from below 0 to over 5. Bracken, Fricker, and Wiper et al. were studying the battle of the Ardennes. Clemens' fits are on the Battle of Kursk.

concluded that: (1) "any square law effects are largely masked by other factors"; (2) the data better fit a set of three separate battles (one distinct battle every six or seven days); and (3) "unless we are able to procure data . . . with many periodic casualty figures for both sides, we will not be able to validate the homogeneous square law (or any other attrition law)." The first attempt to use time-phased data to validate Lanchester's square law was Engel [17]. Using daily data on the American forces during the battle of Iwo Jima, as well as the starting and ending figures for the Japanese, he was able to graphically demonstrate that the square law could reasonably track U.S. force levels. Finally, Speight [16] contains a nice summary discussion and additional analysis on the Ardennes, Inchon–Seoul, and Iwo Jima data. Moreover, he has an excellent discussion on the challenges, limitations, and hazards associated with trying to explain something as complex as combat, or even combat attrition, with simple functions, such as Lanchester equations. Inevitably, for highly aggregated data, many important factors (e.g., abilities, training, morale, organization, objectives, terrain, weather, luck, etc.) are not adequately accounted for. Moreover, the nature of combat ensures that there is considerable uncertainty in the data themselves.

3. BATTLE OF KURSK OVERVIEW

After suffering a terrible defeat at Stalingrad in the winter of 1943, the Germans desperately wanted to regain the initiative. In the spring of 1943, the Eastern Front was dominated by a salient, 200 km wide and 150 km deep, centered on the city of Kursk. The Germans planned, in a classic pincer operation named Operation Citadel, to eliminate the salient and destroy the Soviet forces in it. On 2 July 1943, Hitler declared, "This attack is of decisive importance. It must succeed, and it must do so rapidly and convincingly. It must secure for us the initiative. . . . The victory of Kursk must be a blazing torch to the world." [18]

After nearly 2 months of delays, Operation Citadel was launched on 5 July, with a two-front attack on the Kursk salient. Due to good intelligence and the extra time that the Soviets had to get ready, the Germans attacked well-prepared positions. The attack on the northern front ran into stiff Soviet defenses and quickly bogged down. However, on the southern front of the battle, in heavy fighting, the Germans penetrated as much as 46 km by 12 July. This put the Germans in position to capture the town of Prokhorovka and establish a bridgehead over the Psel River, the last natural barrier between them and Kursk. To counter this, the Soviets deployed their strategic armored reserve, the 5th Guards Tank Army, under Lieutenant General Pavel Rotmistrov. On the 12th, outside of Prokhorovka, the German's II SS Panzer Corps, commanded by SS *Obergruppenfuehrer* Paul Hausser, slammed into the advancing 5th Guards Tank Army. The result has been called the greatest armored engagement in history. On that day, the Germans lost 98 tanks,[1] while the Soviets lost 414 tanks. Although, in terms of casualties, Germany seemed to win the day, Hitler gave orders on 13 July to cancel Operation Citadel. For the rest of the battle, the Germans assumed a generally defensive posture. Field Marshal Erich von Manstein, Commander of Army Group South, felt that "[stopping the offensive] at this moment [was] tantamount to throwing victory away [10]." By 23 July, Soviet counterattacks had regained all of the ground lost in the battle. As Manstein prophetically wrote "the last German offensive in the East ended in a fiasco, even though the enemy . . . suffered four times their losses" (see [10]). More information about the battle of Kursk can be found in [3], [10], [18], and [22].

4. DATA AND ANALYSIS APPROACH

4.1 The Data

The data in [3], provided by CAA, detail Army Group South's attack on the southern side of the salient. The data cover 2 weeks of fighting involving over 300,000 German and 500,000 Soviet combat soldiers, including all line and headquarter units. Support personnel are not in the database. While the database contains information on 15 days of fighting (4–18 July), we only use the last 14 days because the battle did not begin in earnest until 5 July. Table 1 shows the German's and Soviet's daily combat manpower (line and headquarter units) and casualties.

Following Bracken [2], Fricker [8], Clemens [4], and Wiper et al. [20], we aggregate the numerous systems on each side into four broad categories: manpower, tanks, armored personnel carriers (APC), and artillery. See Turkes [19] for a specific mapping of systems to categories in this study. This is not an easy task and probably accounts for some of the differences between the studies.

Table 2 displays the German daily on hand and losses for the four categories during the battle. Daily losses included items destroyed, abandoned, and damaged; some of the damaged items were repaired and returned to battle. Table 3 gives the same data for the Soviets.

[1] It is interesting to note that many sources (e.g., [18]) estimate higher German losses. However, those estimates are based more on second-hand battlefield reports and myths about the battle, rather than a careful examination of unit's logs. The developers of the data, The Depuy Institute, speculate that higher German casualties at Prokhorovka may have served the interests of both Soviet and German propaganda.

Table 1. Combat manpower for both sides.[a]

Day	Soviet manpower	Soviet casualties	German manpower	German casualties
5 July	507698	8527	301341	6192
6 July	498884	9423	297205	4302
7 July	489175	10431	293960	3414
8 July	481947	9547	306659	2942
9 July	470762	11836	303879	2953
10 July	460808	10770	302014	2040
11 July	453126	7754	300050	2475
12 July	433813	19422	298710	2612
13 July	423351	10522	299369	2051
14 July	415254	8723	297395	2140
15 July	419374	4076	296237	1322
16 July	416666	2940	296426	1350
17 July	415461	1217	296350	949
18 July	413298	3260	295750	1054

[a] Casualties are those killed, wounded, captured/missing in action, and disease and nonbattle injuries. Daily force levels depend on previous force levels, casualties, and reinforcements. The attacking Germans are outnumbered.

The data in Tables 1–3 classify the many combat systems in the battle into four broad categories (manpower, tanks, APCs, and artillery). Even so, with only 14 days of data, there are not enough degrees of freedom to fit heterogeneous Lanchester models—though Clemens [4], with additional assumptions (constraints), investigated aspects of heterogeneous models. Thus, like the authors above, we fit homogeneous models. Table 4 presents the data on the combat power of the Soviet and German forces. The combat power of a force is defined as a weighted sum of combat manpower, APCs, tanks, and artillery, with weights of 1, 5, 20, and 40, respectively. These are the weights used in the previous studies. Bracken [2] writes that: "Virtually all theater-level dynamic combat simulation models incorporate similar weights,

Table 2. Daily German on hand and loss data for tanks, APCs, and artillery.[a]

Day	Tanks on hand	Tanks lost	APCs on hand	APCs lost	Artillery on hand	Artillery lost
5 July	986	198	1142	29	1166	24
6 July	749	248	1128	14	1161	5
7 July	673	121	1101	27	1154	7
8 July	596	108	1085	16	1213	13
9 July	490	139	1073	14	1210	6
10 July	548	36	1114	42	1199	12
11 July	563	63	1104	16	1206	15
12 July	500	98	1099	12	1194	12
13 July	495	57	1096	4	1187	7
14 July	480	46	1093	6	1184	5
15 July	426	79	1089	5	1183	3
16 July	495	23	1092	1	1179	4
17 July	557	7	1095	1	1182	2
18 July	588	6	1098	5	1182	11

[a] The Germans suffered their heaviest tank losses during the first 2 days of the battle, when they were attacking heavily prepared defenses.

Table 3. Daily Soviet on hand and loss data for tanks, APCs, and artillery.[a]

Day	Tanks on hand	Tanks lost	APCs on hand	APCs lost	Artillery on hand	Artillery lost
5 July	2396	105	507	4	705	13
6 July	2367	117	501	6	676	30
7 July	2064	259	490	11	661	15
8 July	1754	315	477	13	648	14
9 July	1495	289	458	19	640	9
10 July	1406	157	463	3	629	13
11 July	1351	135	462	4	628	7
12 July	977	414	432	30	613	16
13 July	978	117	424	8	606	10
14 July	907	118	418	8	603	5
15 July	883	96	417	1	601	5
16 July	985	27	417	0	600	3
17 July	978	42	417	2	602	0
18 July	948	85	409	8	591	4

[a] On 12 July, the day of the "blood bath at Prokhorovka," the Soviets lost 414 tanks.

either as inputs or as decision parameters computed as the simulations progress." In a sensitivity analysis, Turkes [19] found that, for the battle of Kursk, the broad conclusions are robust to the chosen weights.

In generalized Lanchester models, the attrition (combat power losses) to each side is a function of the size of the forces, as measured in combat power. Figure 2 displays simultaneous pairwise scatter plots on German combat power (GCP), German combat power losses (GCPL), Soviet combat power (SCP), and Soviet combat power losses (SCPL). Each plot in the figure consists of a scatter plot between two of the four variables, which are identified by the labels in the corresponding diagonal elements. The points in the plots are labeled according to the 14 days of the battle.

We can glean several things from Figure 2. First and foremost is that all of the variables appear to be (generally) positively correlated—with the exceptions often being caused by the outlying data point corresponding to Soviet losses on day 8. The ordering of the numbers in the

Table 4. The combat power and losses of the aggregated forces.

Day	Soviet forces	Soviet losses	German forces	German losses
5 July	586353	11167	373411	11257
6 July	575769	12993	364265	9532
7 July	559345	16266	359085	6249
8 July	545332	16472	372524	5702
9 July	528552	18071	367444	6043
10 July	516403	14445	366504	3450
11 July	507576	10754	365070	4415
12 July	480033	28492	361965	5112
13 July	469271	13302	362229	3491
14 July	459604	11323	359820	3290
15 July	463159	6201	357522	3047
16 July	462451	3600	358946	1975
17 July	461186	2067	360245	1174
18 July	457943	5160	360280	1639

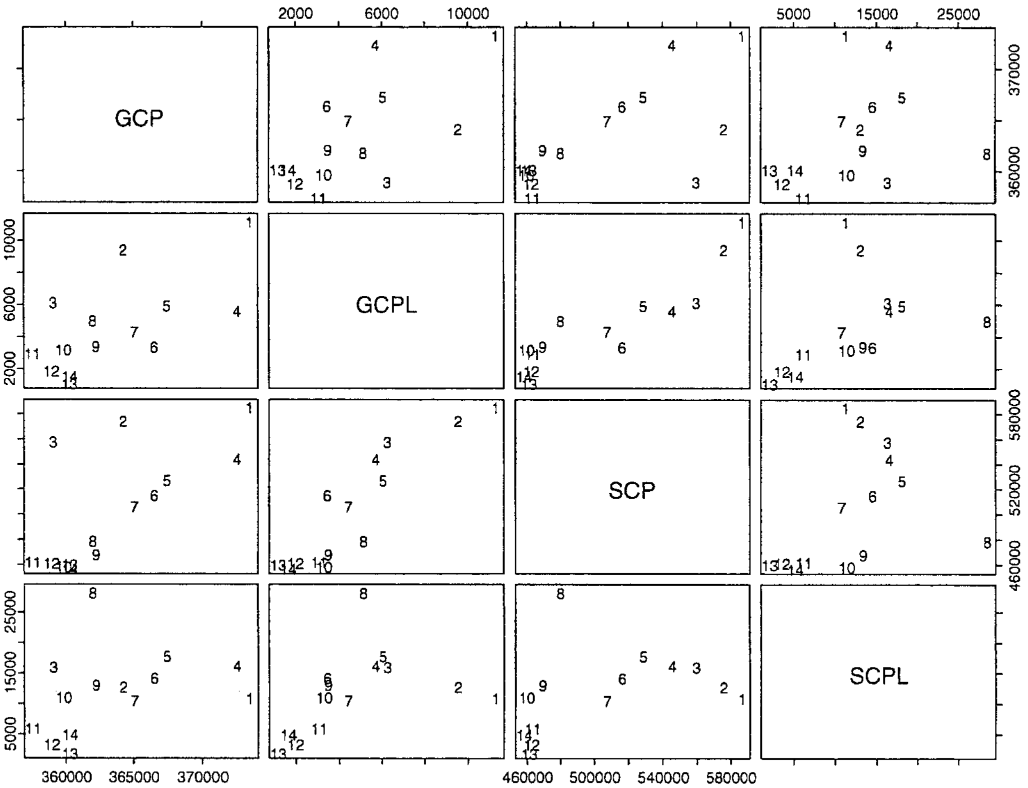

Figure 2. Pairwise scatter plots of German combat power (GCP), German combat power losses (GCPL), Soviet combat power (SCP), and Soviet combat power losses (SCPL), with the points corresponding to the 14 days of data.

plots reveals that the combat forces on both sides, but particularly the Soviets, were steadily declining during the battle. Casualties were also generally decreasing in time, particularly for the Germans. It is important to emphasize that correlations among the variables, and their correlations with time, complicate the analysis by confounding relationships. For example, is the primary cause of the decreasing German casualties the decreasing Soviet force level, the decreasing German force level, a combination of both, or something else that is correlated with time? Also, there is not much variation in GCP; thus, it may be difficult to see the effects that occur as GCP changes. Finally, a few points stand out from the others—in particular, the Soviet losses on day 8 and, to a lesser extent, the German losses during the first 2 days.

4.2. Analysis Approach

Given the aggregated values in Table 4, we investigate what values of the parameters (a, b, p, q, and d) best fit the data. Of course, our focus is on p and q, for they relate to the Lanchester laws of attrition. In particular, we are interested in whether the square, linear, or logarithmic laws fit well. The other parameters (a, b, and d) undoubtedly will vary in different battles, depending on the forces fighting, the intensity of the battle, the terrain, and many other factors. Our measure of fit, taken from [2], is the sum of the squared residuals (SSR) between the estimated and actual attrition. The objective is to find the parameters that minimize SSR—i.e., provide the best fit. Specifically, the objective function that we minimize is

$$\mathrm{SSR} = \sum_{n=1}^{14} (\dot{B}_n - a(d_n^*)R_n^p B_n^q)^2 + \sum_{n=1}^{14} (\dot{R}_n - b(d_n^*)B_n^p R_n^q)^2, \tag{3}$$

where

n indexes the 14 days of the battle,
$d_n^* = d$ if the side (Red or Blue) is on the defensive on day n and $1/d$ if the side is on the offensive. If neither or both sides are clearly on the offensive, then $d_n^* = 1$.

Rather than use a standard optimization method to find the parameters that minimize SSR, such as Bracken's [2] constrained grid search, Fricker's [8] linear regression on logarithmically transformed data, Wiper et al.'s [20] posterior maximum, or Clemens' [4] numerical Newton-Raphson, we look at the best-fitting response surface as a function of p and q. Given p and q, it turns out to be relatively easy to find a, b, and d to minimize SSR. Our approach is as follows: For a fixed d, solve for a and b by regression through the origin (a simple analytic formula exists; see [12] or [13]). Note: B_n, R_n, $\dot{B}_n$, and $\dot{R}_n$ are known—thus, only a and b need to be estimated. This calculation is repeated for many values of d. Specifically, through a one-dimensional line search on d, we find the d, and associated a and b, that minimize SSR for the given p and q. By plotting the contours of the minimum SSR as a function of p and q, not only can we visually assess where the optimum occurs, but we also get a better understanding of how the surface of Lanchester exponent parameters fits the battle.

5. LANCHESTER AND THE BATTLE OF KURSK

This section investigates how the various parameters fit the Kursk data, initially with $d = 1$. Figure 3 displays the contours of the minimum SSR as function of p and q. We find that the surface is relatively flat around the optimum. In fact, the contour with SSR = 6.0×10^8, about 10% above the minimum SSR, contains p values ranging from below 2 to over 10 (off the chart), and q values ranging from below 0 to near 3.

Using a grid search (at four decimals of precision) in the region of the visually obtained optimum, the best fit is obtained at $p = 5.6957$ and $q = 1.2702$, with an SSR of 5.46546×10^8.[2] Note: The precision is needed because small changes in large exponent parameters, as one gets when rounding, dramatically affects the fits. The best-fitting p ($p = 5.6957$) is so large as to be highly implausible. It suggests that a doubling of the attacker's force increases the instantaneous kill rate by greater than 50 times. One must, of course, be wary of extrapolating beyond the range of the data. Nonetheless, most analysts use values of p and q that are within the interval from zero to one—typically zero or one. The best fit is also far from all of the basic Lanchester laws. The a and b corresponding to the optimal fit are 1.466×10^{-35} and 1.201×10^{-36}. This suggests that, individually, the Germans were more effective than the Soviets. The

[2] If we assume that the Germans lost 300 tanks on day 8, as some books suggest, we get almost identical results, with the best fit obtained at $p = 5.01$ and $q = 1.35$.

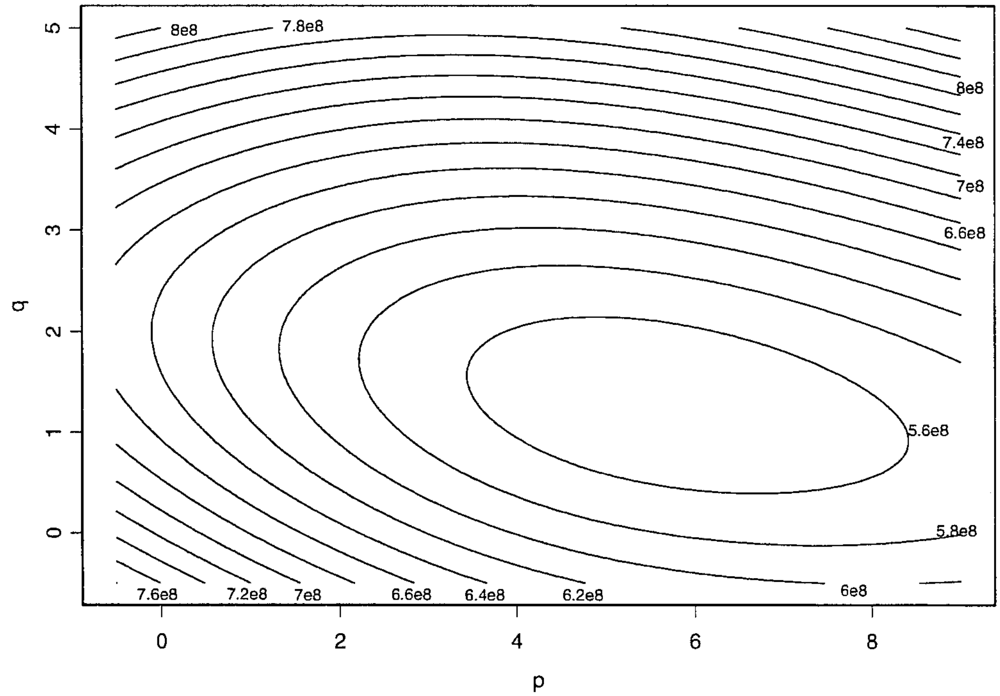

Figure 3. A contour plot of the minimum SSR as a function of p and q, with the tactical parameter (d) not included (i.e., no attacker or defender advantage in this model). Note: The small wiggles and other nonsmooth features in this and the other contour plots are a function of the granularity of the points upon which the surface was evaluated and the software used to generate the plots. In all subregions that the authors have zoomed in on, the contours are smooth.

optimal values of a and b are quite small, as they must be to balance the very high values of the optimal exponent parameters p and q.

In addition to SSR, another measure of fit that we use is R^2, where $R^2 = 1 - \text{SSR/SST}$. SST, the sum squares total, is calculated by

$$\text{SST} = \sum_{n=1}^{14} (\dot{B}_n - \bar{B})^2 + \sum_{n=1}^{14} (\dot{R}_n - \bar{R})^2, \tag{4}$$

with $\bar{B}$ and $\bar{R}$ denoting the mean daily force losses for the Soviets and Germans, respectively. Here, R^2 is defined analogously to how it is in standard least squares regression and is a linear function of SSR. Larger R^2 values correspond to better fits, with a perfect fit (i.e., an SSR = 0) yielding an $R^2 = 1$. Note: In this setup, it is possible to get negative R^2 values, which means that the fitted model yields worse results than using the average daily losses as estimates. An advantage of R^2 over SSR is that it is invariant to linear transforms of the data. This allows us to compare models using different weights and data. In the best-fitting model above, the R^2 is .237. That is, the model explains less than a quarter of the squared variation.

We can see how well the optimum parameter estimates fit the battle by examining the actual and estimated losses over the 14 days. Figures 4 and 5 show the actual and estimated Soviet and

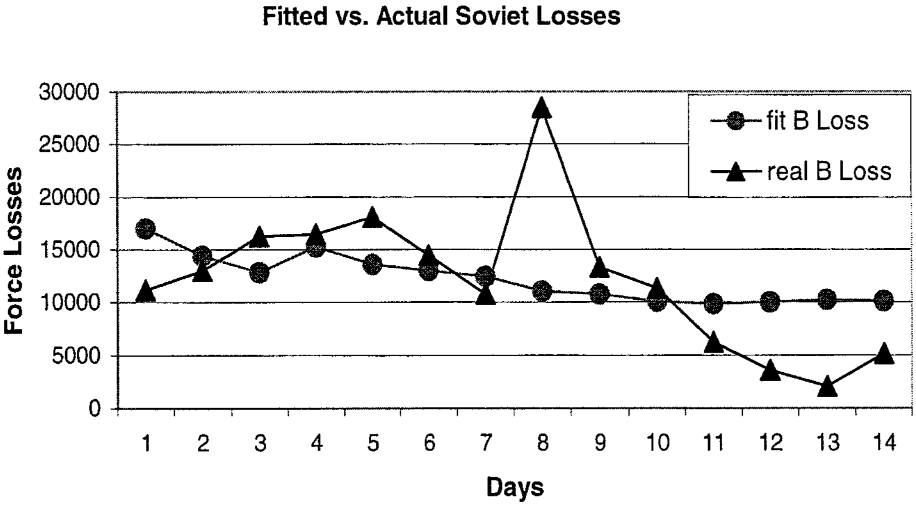

Figure 4. Fitted losses plotted versus real losses for the Soviet forces using the best-fitting Lanchester model (with $d = 1$). There is one outlying data point (corresponding to 12 July). Also, the estimated losses are noticeably high for the last 4 days.

German losses using the best-fitting generalized Lanchester model. The estimated German losses track the real losses quite well, while the Soviet's do not. In fact, the R^2 calculated using only the German forces is .870. The R^2 for only the Soviet forces is .124—only marginally better than using the mean daily loss as the estimate. The Soviet's fit is highly affected by the outlying point on day 8. Also, we see that, for both sides, the model overestimates the casualties for the last few days of the battle, suggesting that the battle had lost some of its intensity. This brings into question the assumption of constant attrition parameters.

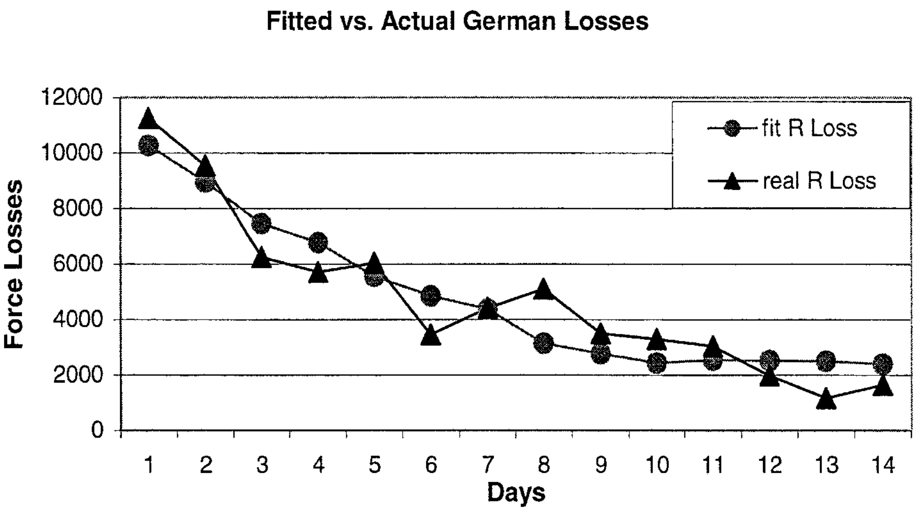

Figure 5. Fitted losses plotted versus real losses for the German forces using the best-fitting Lanchester model (with $d = 1$). The fit to the Germans is much better than that of the Soviets.

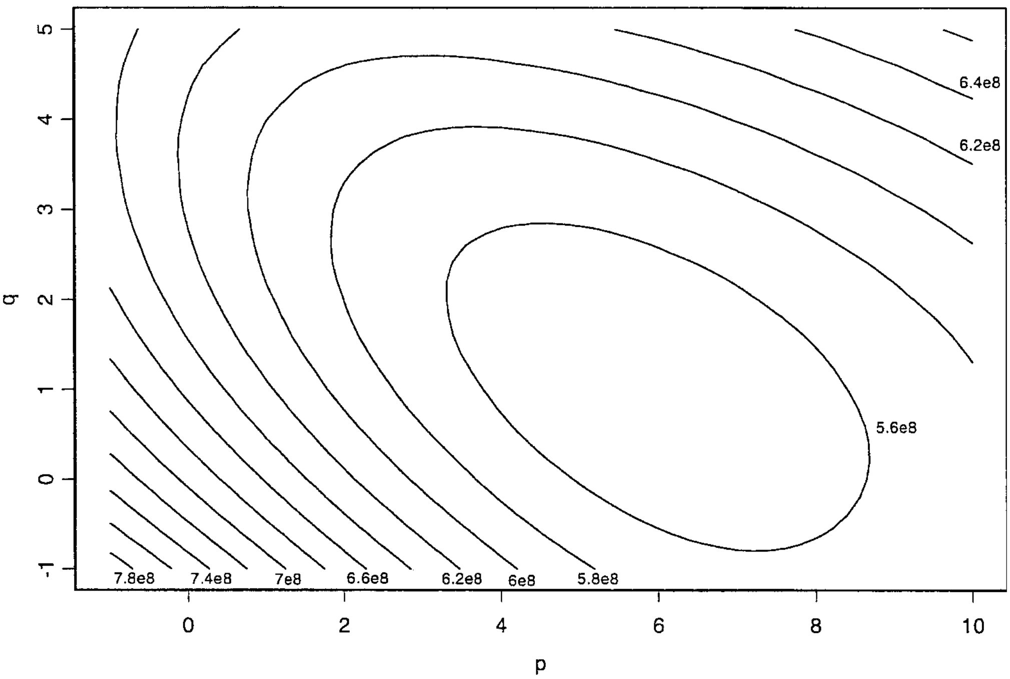

Figure 6. A contour plot of the minimum SSR as a function of p and q with (d) included.

Let us now consider the tactical parameter d, which turns out to be a complex issue. From [3], we see that on most days, over the entire southern front, some units of both the Soviets and the Germans are attacking and defending. However, in general, the Germans were the attackers during the first 7 days of the data. On the 8th day, for the first time, more Soviet than German units were attacking. After the 8th day, the proportion of German units attacking is significantly lower, while the proportion of Soviet units attacking is generally higher. Consequently, we define the Germans as the attackers the first 7 days and the Soviets the last 7 days. With the tactical parameter d so defined, the best-fitting model is $p = 5.8703$, $q = 1.0078$, $d = 1.028$, $a = 4.907 \times 10^{-35}$, and $b = 3.521 \times 10^{-36}$. The new SSR is 5.46202×10^8, a negligible improvement over the no tactical parameter fit. The new R^2 is .238. The Germans are still more lethal per unit of force ($a > b$), and there is a slight attacker advantage ($d > 1$). The contour plot of the minimum SSR (see Fig. 6) as a function of p and q, (now optimized over a, b, and d), is similar in shape to Figure 3, though slightly flatter, as the d parameter allows more flexibility in the fit.

6. LANCHESTER AND THE BATTLE OF THE ARDENNES

In this section, we apply the same approach to the data on the Ardennes campaign (i.e., mapping the surface contour as a function of p and q). See [2], [5], [8], and the references therein, for overviews of the battle of the Ardennes. Specifically, we look at contour plots for both Bracken's [2] and Fricker's [8] data sets on the Ardennes campaign. Figures 7 and 8 display the minimum SSR as a function of p and q for Bracken's [2] and Fricker's [8] data, respectively. In both figures, the tactical parameter d is used in the fit. As before, for each p and q, the minimum SSR (optimized over a, b, d) is displayed.

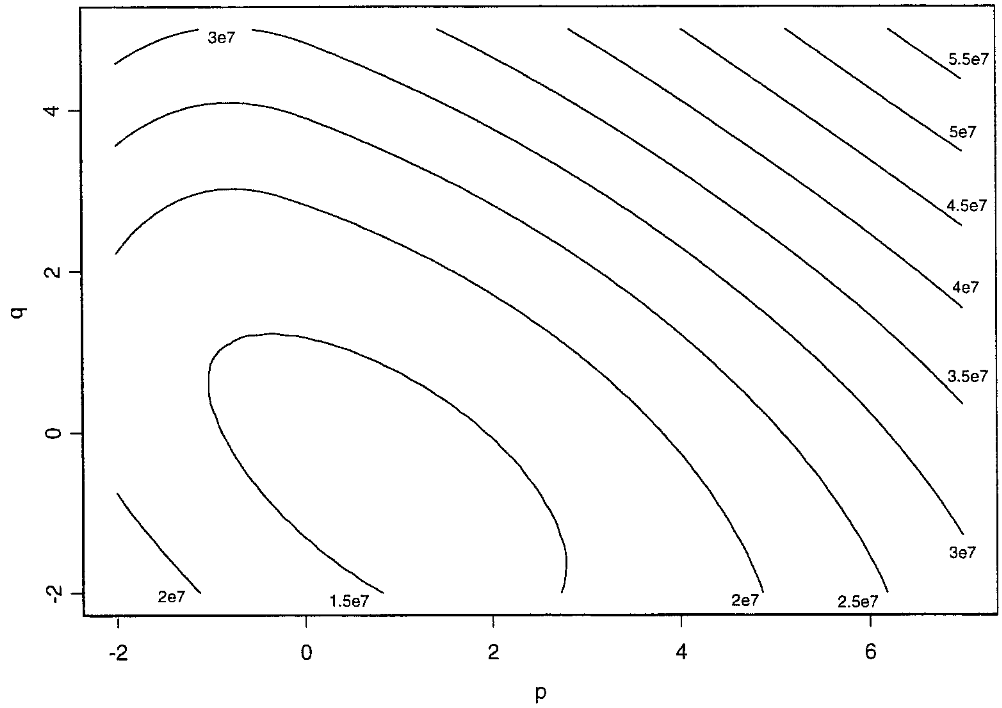

Figure 7. A contour plot of the minimum SSR as a function of p and q, with the tactical parameter d, using Bracken's [2] Ardennes combat forces data.

The best-fitting p and q found while calculating Figure 7 are $p = .91$ and $q = -.61$, with an SSR of 1.36×10^7 and an R^2 of .381. The tactical parameter at the empirically determined best fit is 1.12, implying an attacker advantage. Moreover, we see from the graph that the surface is relatively flat around the optimum, and contrary to what Bracken found with his grid search, the square law fits the best of the major Lanchester laws (i.e., square, linear, and logarithmic). The flatness of the surfaces also suggests that in Bayesian analyses posterior inferences will be strongly influenced by precise priors, as Wiper et al. [20] found.

The best-fitting p and q found while calculating Figure 8 are $p = -.6$ and $q = 5.2$, with a SSR of 2.8×10^7 and an R^2 of .50. These p and q seem questionable, with a side's casualties increasing exponentially (at a power of 5.2) as a function of its force level, and inversely to its opponent's force level. The tactical parameter d, at the empirically determined best fit, is 1.25, again implying an attacker advantage.

The last two figures, based on different approaches to formatting the data, tell very different stories. Which surface should we believe? There are arguments for both approaches. Fricker [8] fit his model to more days; however, over the 10 days Bracken [2] used, we are more likely to get a relatively homogeneous battle. Furthermore, Bracken's model of 10 1-day battles may be easier to interpret than Fricker's formulation, which considers the fighting one big battle, with all of the forces that were involved in the battle calculated in the equations as if they participated from the beginning. Substantial forces, particularly on the Allied side, such as Patton's Third Army, did not participate until several days after the start of the campaign (see [21]). The key point, however, is not which formulation is better, as there are good arguments both ways, but to notice that the surfaces look qualitatively different, depending on how the data are formatted.

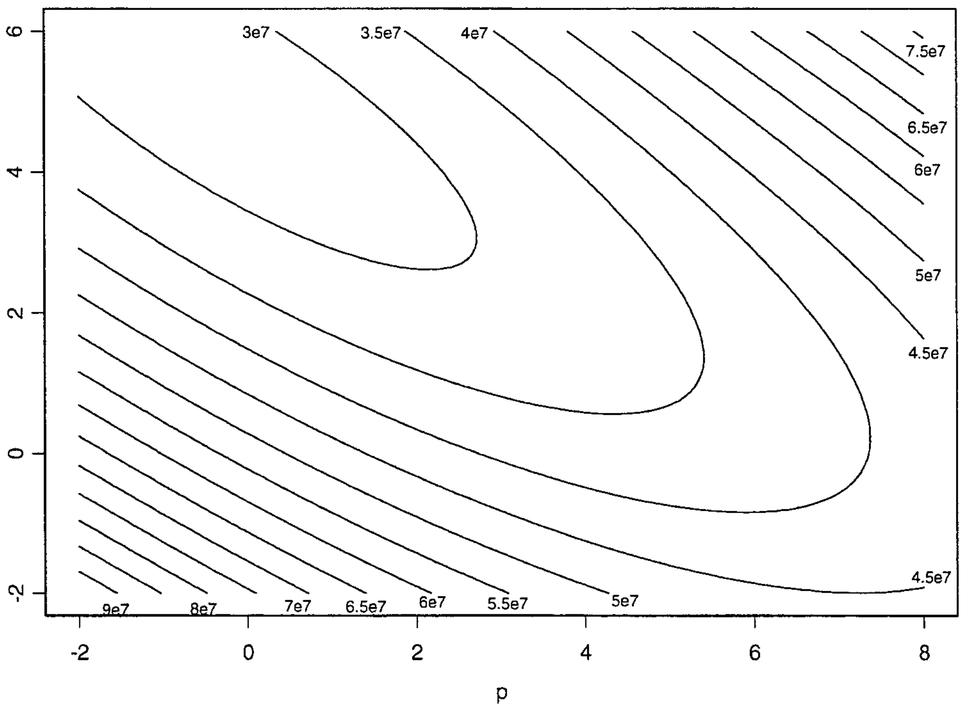

Figure 8. The minimum SSR as a function of p and q, with the tactical parameter d, using Fricker's [8] formulation of the Ardennes combat forces data.

Thus, the conclusions can be very sensitive to these factors. Furthermore, in neither case is the optimum near one of the basic Lanchester laws.

7. EXCURSIONS

In the previous sections we looked at how the generalized Lanchester exponent parameters (p and q) fit the battles of Kursk and Ardennes. In this section, we look more closely at the basic Lanchester models, some of the key assumptions, and variability in the fits.

7.1. Fitting Lanchester Equations for the Battle of Kursk Using Combat Manpower

In our fits above, when computing the combat power of a force, we use the previous authors' subjective weights. The question is: Do the weights chosen qualitatively affect what we found? Figure 9 shows the minimum SSR surface, as a function of p and q, with the tactical parameter d, using only combat manpower. That is, the weights for artillery, APCs, and tanks are all set at zero. When comparing Figure 9 with Figure 6, we see that they are qualitatively similar. Both surfaces have smooth, relatively flat, elliptical contours, with the minimum occurring far from the basic Lanchester laws. Again, the best-fitting p and q values are highly implausible, with the optimal p exponents much greater than the optimal q exponents. The optimum p and q values for the combat manpower data are $p = 7.8$ and $q = 3.4$, with an R^2 of .234. In [19], a variety of other weights are examined (e.g., tanks worth more than artillery, etc.) with the bottom-line

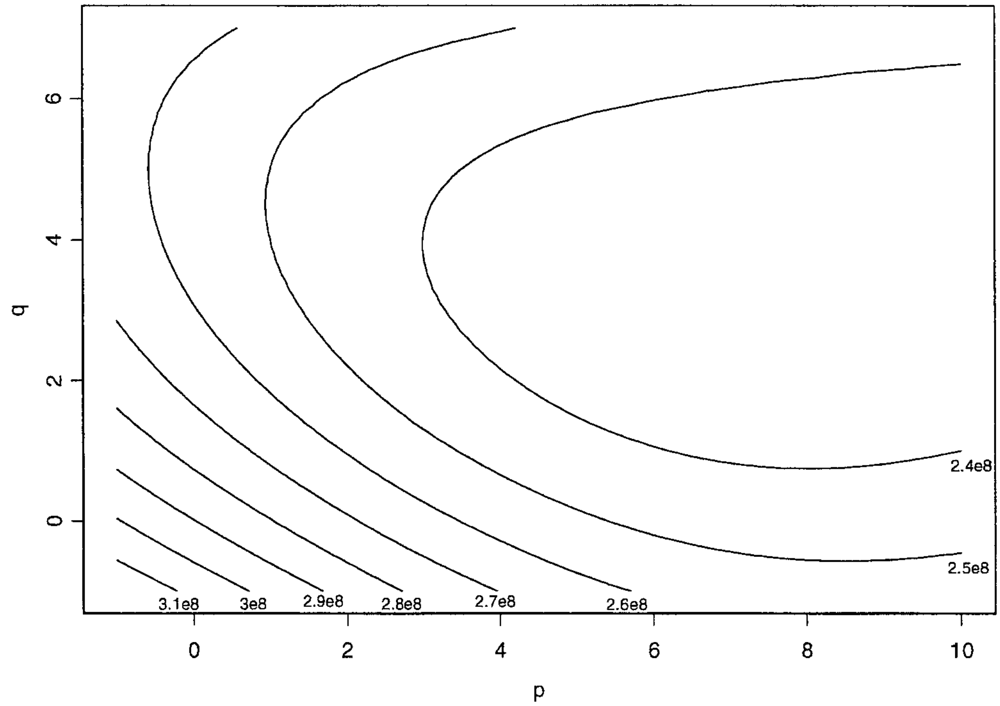

Figure 9. The minimum SSR as a function of p and q for Kursk, with the tactical parameter d, based only on combat manpower.

conclusions remaining the same. Furthermore, adding air-sorties to the calculations does not substantially affect the fits (again see [19]).

7.2. Fitting the Basic Lanchester Models

We have seen that the best-fitting generalized Lanchester models are far from the basic Lanchester laws (i.e., square, linear, and logarithmic). Several questions remain: How do these laws compare to one another? How do they compare to the best-fitting models? Does one of the

Table 5. Lanchester law fits for the battle of Kursk, with the tactical parameter d, for both Bracken's [2] weights and combat manpower.[a]

Lanchester law	Weights	p	q	d	R^2
Square	Bracken's	1	0	1.09	.081
Linear	Bracken's	1	1	1.02	.131
Logarithmic	Bracken's	0	1	1.02	.085
Optimum fit	Bracken's	5.87	1.01	1.03	.238
Square	Manpower	1	0	1.11	.074
Linear	Manpower	1	1	1.04	.116
Logarithmic	Manpower	0	1	1.04	.086
Optimum fit	Manpower	7.74	3.41	.86	.234

[a] In both cases, the linear law is the best fitting of the basic laws; however, the linear law is not close to the optimal fit.

Table 6. Lanchester law fits for the battle of the Ardennes using Bracken's [2] combat manpower model and the tactical parameter d.[a]

Lanchester law	Weights	p	q	d	R^2
Square	Bracken's	1	0	1.14	.367
Linear	Bracken's	1	1	1.17	.291
Logarithmic	Bracken's	0	1	1.23	.330
Optimum fit	Bracken's	.91	−.61	1.12	.381
Square	Manpower	1	0	1.04	.079
Linear	Manpower	1	1	1.06	−.226
Logarithmic	Manpower	0	1	1.10	.025
Optimum fit	Manpower	.15	−.90	1.05	.280

[a] Bracken's weights provide much better fits than the manpower data. Also, the square law is the best-fitting basic Lanchester law.

basic laws consistently fit better across the breadth of models examined? This subsection examines these questions. Here, we consider only models with the tactical parameter d. Similar results hold for models without d.

Table 5 shows how the basic Lanchester models (i.e., the square, linear, and logarithmic laws) fit the battle of Kursk using Bracken's [2] weights (i.e., 1, 5, 20, and 40) and the manpower data. We see that, for both sets of weights, the linear law fits best among the basic laws, but it is significantly inferior to the optimum fit, with an R^2 of .131, compared to .238. While the linear law fits better than the square and logarithmic laws, its fit is not substantially better. Thus, this suggests that with the right coefficients, any of the basic laws give about the same result. It is also interesting to note that we get similar results using either weighting scheme, with Bracken's weights generally fitting slightly better.

Table 6 displays similar information from the Ardennes campaign, using Bracken's 10 days of data [2]. There are several interesting features in the table. First, much better fits are achieved with Bracken's weights than with combat manpower. Furthermore, while the best-fitting basic Lanchester law is the square law, with Bracken's weights, all of the basic laws fit the data about as well. The R^2 (.367) for the square law is only 3.6% below the optimum fit's R^2 (.381). It is interesting to note that Bracken concluded that the best-fitting model was the linear law. Why the discrepancy? Because Bracken used a grid search, which, as he noted in [2], by necessity, will not consider all of the parameter combinations for the various Lanchester laws. When comparing Tables 5 and 6, we see that the Ardennes data fit all of the models significantly better than the Kursk data when using Bracken's weights. However, much of the difference is due to the Soviet losses on 12 July, which is an outlier. Moreover, none of the basic laws stand out as the best fitting.

7.3. Breaking the Battle of Kursk into Several Phases

An important assumption we (and the previous researchers) made when fitting the models is that the attrition parameters (a and b) remain constant throughout the campaign. Surely, over the course of the battle, the attrition rates waxed and waned to some extent. Ideally, we would estimate new coefficients for each day of the fighting, as the data are available daily. Unfortunately, if we do so, the system is overdetermined. That is, on each day, for any p and q, there exist a and b such that the fitted casualties equal the real data.

To overcome this limitation, we break the battle into several phases, each of which is believed relatively homogeneous. While several different partitions of the data were examined (see [19]),

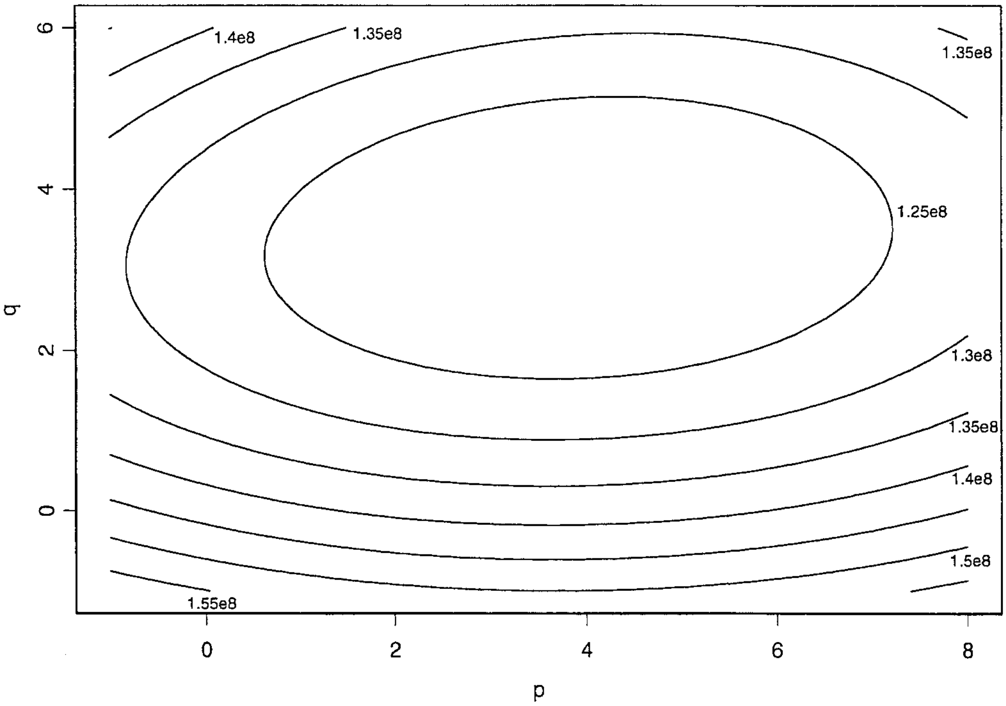

Figure 10. The SSR surface when the attrition coefficients are fit separately for four different phases (days 1–2, 3–7, 8, and 9–14). Due to the additional parameters, the fits are much better; however, the shape of the contour plot is similar to Figure 3. Here, too, the minimum occurs at implausibly large exponent parameters p and q.

we focus here on what seems to be the most natural one based on the Kursk data and other historical accounts. In our first phase—the first 2 days of the campaign—the Germans generally attack prepared defenses. Our second phase contains days 3–7, which, by and large, had the Germans attacking hasty defenses. The 8th day of fighting, "the bloodbath at Prokorovka," is unique and is considered a phase by itself. Of course, since this phase is only a single day, there is a perfect fit (i.e., no residual error); thus, this removes the (outlying) eighth day from the fits. The fourth and last phase is days 9–14, in which the Soviets were more on the offensive, and the battle intensity was fading, as can be seen in Figures 2, 4, and 5.

Table 7. Lanchester law fits for the battle of Kursk, using Bracken's [2] weights, when the attrition coefficients are estimated in four separate phases.[a]

Lanchester law	p	q	R^2
Square	1	0	.804
Linear	1	1	.816
Logarithmic	0	1	.812
Optimum fit	3.92	3.38	.832
Four phase means fit	Not applicable	Not applicable	.800

[a] Almost all of the improvements in R^2 are due to the differences in means between the four phases.

Fitting the model over multiple phases must result in a better overall fit because there are additional parameters (degrees of freedom) to explain the variation in casualties. In our earlier fits there were five parameters (p, q, a, b, and d). The new fits have 10 parameters: the exponent parameters p and q, and different attrition coefficients, a and b, in each of the four phases. The defensive parameter d cannot be used here, because in each of the phases one side is on the offensive throughout. Thus, within any phase, the effect of being on the defensive, as measured by d, is confounded with the phase's attrition parameters.

The optimization routine, as before, searches for the p and q that minimize the sum of the squared residuals between the fitted and actual casualties. For all p and q that are evaluated, in each phase, the best-fitting a and b are determined by regression through the origin using the data for that phase. This results in eight attrition coefficient parameters being calculated for every p and q examined. For each p and q, the sum of the squared residuals (SSR) is calculated by summing over the entire campaign (14 days) the squared difference between the fitted and actual casualties. Figure 10 displays the SSR surface as a function of p and q. This surface is strikingly similar in shape to the surface in Figure 3, generated by the constant attrition coefficient model, in that the contours are elliptic in shape, with the axes almost aligned to the ordinate and the abscissa. In both figures, the ellipses are wider than they are tall, and the minimums are in the first quadrant at implausibly large exponent parameters. Some important differences are: (1) the SSR surface for the four-phase model is much lower, i.e., a significantly better fit; and (2) the four-phase model surface is significantly flatter, implying that a wider range of generalized Lanchester models are "near optimal."

Table 7 displays the optimum and basic Lanchester model fits for the four-phase model. The optimum Lanchester exponent parameters for the four-phase model are $p = 3.92$ and $q = 3.38$, with an R^2 of .832. This is a much better fit than any of the single-phase models we looked at. However, most of the improvement comes from partitioning the battle into the four phases. The R^2 that is obtained by using the mean loss in each phase, for each side, as the sides' estimated losses, is .800. Consequently, the optimum four-phase Lanchester model, with 10 free parameters, has an R^2 of only .032 above the R^2 obtained by using the four phase means. Thus, accounting for the varying intensity, terrains, postures, etc. explains significantly more of the variation in losses than the Lanchester models do. This is consistent with what Hartley and Hembold found in their study of the Inchon–Seoul battle [9]. Also, we can see from Table 7 that the basic Lanchester models fit only marginally better than the phase means, with the linear law fitting slightly better than the square and logarithmic laws.

7.4. Assessing Differences in Fits

In the discussions above on the differences between various fits, as measured by R^2, we loosely refer to some differences as significant and others as insignificant. In this section we use the bootstrap (see Efron and Tibshirani [6]) to formally assess the differences in R^2 between Lanchester models. We focus on identifying differences that are of sufficient size that the ordering of the estimates (different R^2 values) is unlikely to be effected by the natural variation in the data. Keeping in mind that all battles are nonrepeatable events, we define the natural variation as the variation that would occur (due to the inherent randomness of combat) if many essentially identical forces fought similar 14-day battles and the inevitable errors associated with recording and collecting decades old combat data.

In fitting the Lanchester models above, the 14 days of the battle are essentially treated as 14 minibattles. That is, for a given law (i.e., specific p and q), we find the values of a, b, and d which minimize R^2 (or equivalently SSR) over the 14 days. Here, we quantify the variability in

Table 8. Bootstrap estimate of the standard error and 90% bootstrap percentile intervals in Kursk R^2 values for the basic Lanchester laws and the optimum law.

Lanchester law	R^2	Bootstrap mean	$ESE(R^2)$	90% bootstrap percentile interval
Square	.081	.054	.064	(−.012, .185)
Linear	.131	.087	.102	(−.059, .265)
Logarithmic	.085	.065	.087	(−.045, .234)
Optimum fit	.238	.254	.158	(.020, .533)

R^2, for each law, nonparametrically, by resampling the empirical daily attrition coefficients, from the 14 days, as follows. For each of the three basic Lanchester laws and the Lanchester optimum fit, using the law's optimum d value, we calculate $\hat{a}_i$ and $\hat{b}_i$ for $i = 1, 2, \ldots, 14$, where $\hat{a}_i$ and $\hat{b}_i$ are the daily attrition parameters that achieve equality in Eqs. (1) and (2). That is, the 14 ($\hat{a}_i$, $\hat{b}_i$) pairs are the attrition rates that actually occurred in the battle (according to the data set) if the Lanchester law being used (to generate the resampled battles) held exactly. A "bootstrap battle" is created by sampling with replacement from the 14 ($\hat{a}_i$, $\hat{b}_i$) pairs and generating 14 daily "bootstrap casualties" by multiplying the 14 ($\hat{a}_i$, $\hat{b}_i$) pairs with the actual force levels and the law's optimum d (or $1/d$, as appropriate). Thus, all of the bootstrap battles have the same daily force levels, attacker, and defender as in the real battle. The bootstrap battles are different from the real battle in that the daily casualties are generated as just specified.

Using the above procedure for the three basic Lanchester laws and the optimum Lanchester fit, 1000 bootstrap battles (from the 14^{14} possible ones) are independently generated (sampled). In each of the 1000 bootstrap battles we find the best-fitting model (maximum R^2) over a, b, and d, as before. This gives us a sample of 1000 bootstrap R^2 values, which we label R_i^2, for $i = 1, 2, \ldots, 1{,}000$. For each law, the R_i^2 values are sorted from smallest to largest. The interval from the 50th largest R_i^2 to the 950th largest R_i^2 constitutes a 90% bootstrap percentile interval (see [6]). Table 8 shows the bootstrap percentile intervals for the battle of Kursk. From the table, we see that the differences in R^2 estimates from the real data are small relative to the variability in them. This is especially true for the basic laws. In particular, there is substantial overlap in the percentile intervals.

The variability associated with the R^2 estimates makes it seem as though the Kursk data do not support one model over the others. However, there are substantial positive correlations between the bootstrap R^2 values for the different Lanchester models. That is, bootstrap battles that fit well to one law usually do likewise to the other laws. In our resampling, we use the same resampled days for all four of the Lanchester laws that are compared. Thus, we can count how often one law fits better than another. Table 9 displays this information for all six pairs of laws. Only the optimum fit consistently fits better than the other laws. Thus, we conclude that the

Table 9. Proportion of bootstrap battles (out of 1000) in which the Lanchester Law specified by the row fit better (i.e., higher R_i^2) than the Lanchester law specified by the column for the battle of Kursk.

Lanchester Law	Square	Linear	Logarithmic
Linear	.636	—	—
Logarithmic	.584	.291	—
Optimum fit	.970	.980	.968

Table 10. Proportion of bootstrap battles (out of 1000) in which the Lanchester Law specified by the row fit better (i.e., higher R_i^2) than the Lanchester law specified by the column for the battle of the Ardennes using Bracken's data.

Lanchester law	Square	Linear	Logarithmic
Linear	.604	—	—
Logarithmic	.376	.268	—
Optimum fit	.525	.446	.638

natural variation in the data is such that for any given realization all of the basic Lanchester laws had a reasonable chance (greater than a quarter of the time) of being a better fit than any other basic law.

Table 10 is similar to Table 9, only it uses the first 10 days of Bracken's Ardennes data. Here, for any given bootstrap realization, all of the laws have a nontrivial empirical probability ($>.25$) of having a greater bootstrap R^2 than any other (including the optimal) law. Thus, here we conclude that the observed ordering of the laws can easily be explained simply by the variation inherent in the data.

8. CONCLUSIONS

Generalized Lanchester models and their extensions are used to model highly aggregated attrition in many important analyses despite little empirical validation. The recent work by CAA in developing the detailed two-sided, time-phased Ardennes and Kursk databases offers new hope to those trying to validate Lanchester's models and determine which variant best applies to real aggregated combat. Our efforts focused on answering the following questions: (1) Why had the previous studies on the Ardennes data given such discordant best-fitting models? (2) Can our new method and the Kursk data add something to the long running debate about which Lanchester law is "best?"

By recognizing that, given p, q, and d, the optimum a and b can be found by regression through the origin, our approach reduces the dimensions that need to be searched from five to three. Furthermore, this approach allows us to easily find the contours of the sum of the squared residuals' surface as a function of p and q. Consequently, we are able to visually identify where the optimum occurs. Moreover, we get a better understanding of how the fit varies with the parameter settings. An examination of SSR surfaces, in particular, their relative flatness, reveals why Bracken [2], Clemens [4], Fricker [8], and Wiper, Pettit, and Young [20] obtained such varied optimal parameter estimates.

As in previous studies, we were searching to see which of the basic Lanchester laws best fit the data. We found that there is little difference in fits between the square, linear, and logarithmic Lanchester laws—with those observed differences explainable simply by chance variation! For analysts, this implies that it doesn't matter which of the basic Lanchester laws you use. With the right coefficients, you will get about the same result with whichever variant is used. More importantly, we found that no constant attrition coefficient generalized Lanchester model fit very well in either battle! Indeed, much more of the variation in casualties is explained by the phases of the battle. Therefore, if there is an underlying Lanchester law, at least for highly aggregated data, its effects appear to be dominated by other factors.

Lanchester's intuitive and parsimonious equations have been, since their introduction the better part of a century ago, the most common tool for modeling aggregate attrition. While we are wary about making too much from two battles (though these are all we have of this type), this research adds to the evidence that Lanchester equations may be too blunt of an instrument for modeling the attrition of highly aggregated forces. Indeed, it is asking a lot to address most of the complexities of combat attrition in a model with only a handful (four or five in this paper) of parameters. The failure to find any good-fitting Lanchester model suggests that it may be beneficial to look for new approaches to model highly aggregated attrition.

ACKNOWLEDGMENTS

This work was inspired and shaped by the efforts of Jerry Bracken and Ron Fricker, to whom thanks are given. The authors are also grateful to E. B. Vandiver III and his staff at the Center for Army Analysis, who provided the data and excellent feedback. Finally, the reviews provided by the associate editor and a referee significantly improved the content and clarity of this paper.

REFERENCES

[1] J. Appleget, “The combat simulation of Desert Storm with applications for contingency operations,” Warfare modeling, J. Bracken, M. Kress, and R. Rosenthal (Editors), Military Operations Research Society, Alexandria, VA, 1995, pp. 549–571.
[2] J. Bracken, Lanchester models of the Ardennes Campaign, Naval Res Logist 42 (1995), 559–577.
[3] Center for Army Analysis, Kursk operation simulation and validation exercise—Phase II (KOSAVE II), The U.S. Army’s Center for Strategy and Force Evaluation Study Report, CAA-SR-98-7, Fort Belvoir, VA, September 1998.
[4] S. Clemens, The application of Lanchester Models to the Battle of Kursk, unpublished manuscript, Yale University, New Haven, CT, 5 May 1997.
[5] Data Memory Systems Inc., The Ardennes Campaign simulation Data base (ACSDB) Final Report, Center for Army Analysis, Fort Belvoir, VA, February 1990.
[6] B. Efron and R. Tibshirani, An introduction to the bootstrap, Chapman and Hall, New York, 1993.
[7] J. Engel, Verification of Lanchester’s Law, Oper Res 2 (1954), 163–171.
[8] R. Fricker, Attrition models of the Ardennes Campaign, Naval Res Logist 45 (1998), 1–22.
[9] D. Hartley III and R. Helmbold, Validating Lanchester’s Square Law and other attrition models, Naval Res Logist 42 (1995), 609–633.
[10] D. Glantz and J. House, The Battle of Kursk (Modern War Studies), University Press of Kansas, Lawrence, KS, 1999.
[11] F. Lanchester, Aircraft in warfare: The dawn of the fourth arm, Constable, London, 1916.
[12] R. Myers, Classical and modern regression with applications, Duxbury, Boston, 1986.
[13] J. Neter, M. Kutner, C. Nachtsheim, and W. Wasserman, Applied linear statistical models, 3rd edition, Irwin, Chicago, IL, 1996.
[14] M. Osipov, The influence of the numerical strength of engaged sides on their casualties, Voenniy Sb (Military Coll) (6–10) (1915).
[15] R. Peterson, On the “Logarithmic Law” of combat and its application to tank combat, Oper Res 15 (1967), 557–558.
[16] R. Speight, Lanchester’s equations and the structure of the operational campaign: Within-campaign effects, Mil Oper Res 6(1) (2001), 81–103.
[17] J. Taylor, Lanchester Models of warfare (two volumes), Operations Research Society of America, Arlington, VA, 1983.
[18] P. Tsouras, The Great Patriotic War, Presidio Press, Novato, CA, 1992.
[19] T. Turkes, Fitting Lanchester and other equations to the Battle of Kursk data, Masters Thesis, Department of Operations Research, Naval Postgraduate School, Monterey, CA, March 2000.
[20] M. Wiper, L. Pettit, and K. Young, Bayesian inference for a Lanchester type combat model, Naval Res Logist 47 (2000), 541–558.
[21] P. Young (Editor), Great generals and their battles, The Military Press, Greenwich, CT, 1984.
[22] *http://dspace.dial.pipex.com/town/avenue/vy75/*, last accessed 15 January 2003.