To big wing, or not to big wing, now an answer

Matthew Oldham

Computational Social Science Program
Department of Computational and Data Sciences
George Mason University, 4400 University Drive, Fairfax, VA 22030
moldham@gmu.edu

Abstract. The Churchillian quote "Never, in the field of human conflict, was so much owed by so many to so few" [3], encapsulates perfectly the heroics of Royal Air Force (RAF) Fighter Command (FC) during the Battle of Britain. Despite the undoubted heroics, questions remain about how FC employed the 'so few'. In particular, the question as to whether FC should have employed the 'Big Wing' tactics, as per 12 Group, or implement the smaller wings as per 11 Group, remains a source of much debate. In this paper, I create an agent based model (ABM) simulation of the Battle of Britain, which provides valuable insight into the key components that influenced the loss rates of both sides. It provides mixed support for the tactics employed by 11 Group, as the model identified numerous variables that impacted the success or otherwise of the British.

1 Introduction

1.1 The Battle of Britain

The air war that raged over Britain between the 10th of July and 31st of October 1940 is colloquially known as the Battle of Britain. The battle's significance comes from the fact that not only did the Germans fail to achieve either of their objectives, but it is seen as the first major campaign to be fought entirely by air forces [2]. The initial phase of the battle revolved around the German's attempt to gain air superiority prior to their planned invasion of England – Operation Sea Lion. After September 6th, the Germans shifted to bombing civilian targets, a period that has become known as the 'Blitz', as they attempted to force Britain into surrender.

At the commencement of the battle, the RAF was at a numerical disadvantage having only 754 front line fighters spread across the entire country to combat the combined Luftwaffe force of 2,288 (1,029 fighters and 1,259 bombers) [2]. Despite this numerical disadvantage, the RAF managed to match or exceed the daily sortie rate of the Luftwaffe [5], achieved with some pilots flying up to four sorties a day. The cost of the battle was high for both sides, with FC losing over 1,000 aircraft and 544 of the approximate 3,000 aircrew that participated. Luftwaffe losses totaled nearly 1,900 aircraft and more than 2,600 of their airmen killed [2].

Prior to World War II (WWII) the RAF developed its fighter defense strategy in line with the principles of concentration [7], which stemmed from the Lanchester equation [6] of aimed fire. When it came time to defend Britain, there were two implementations of the FC doctrine. Air Vice Marshall Keith Park, who controlled 11 Group, the Group which bore the brunt of the action in the Battle of Britain, tended to send single or pairs of squadrons (12 aircraft per squadron) to intercept the enemy. This allowed Park to confront the enemy while denying the Luftwaffe a major engagement. Air Vice Marshall Leigh-Mallory who controlled 12 Group, which was typically held in reserve, preferred to form a 'Big Wing' of 5 or more squadrons before engaging [7]. The main issue with this tactic was the time it took to arrange the 'Big Wing', which in turn limited the time the wing had to search for the enemy, and ultimately engage it. Another negative of the tactic for the RAF was that a larger formation was what the Luftwaffe was seeking, as it improved its chances of inflicting greater losses [5].

1.2 The Lanchester model

The advent of air warfare during the First World War (WWI) necessitated a rethink of existing military doctrine. One such attempt was provided in [6], where Lanchester developed a mathematical model addressing the implications of various combat scenarios, including directed fire. Equation 1 illustrates the general form of the model, where a force's loss rate $\left(\frac{dB}{dt}\ or\ \frac{dG}{dt}\right)$ is dependent on $g(b)$, the kill rate/effectiveness of the opposition, the strengthen of the opposition as given by $G(B)$, raised to a particular power $g1(b2)$, and the strengthen of your force $B(G)$, raised to a particular power $b1(g2)$[1].

$$\frac{dB}{dt} = -gG^{g1}B^{b1}, \qquad \frac{dG}{dt} = -bB^{b2}G^{g2}. \tag{1}$$

One particular form of the model is the aimed fire model, where $g1 = b2 = 1$ and $b1 = g2 = 0$. These values allow Equation 1 to be simplified and after setting the conditions by which both forces suffer the same proportional losses $\left(\frac{dB}{dt}\right)/B = \left(\frac{dG}{dt}\right)/G$, the following equation is derived:

$$gG^2 = bB^2. \tag{2}$$

The importance of Equation 2 is that when forces are using aimed fire; their fighting strength becomes proportional to a weapon's effectiveness multiplied by the square of the number of weapons employed. The implication being, the concentration of force becomes a vital consideration in military strategy [5].

An analysis of the Battle of Britain utilizing the Lanchester model was undertaken by [5] in an attempt to understand whether the 'Big Wing' approach was the correct approach. The conclusion of [5] was that the model was right about British losses, a large German force meant greater losses, but not about German losses. Therefore, the 'Big Wing' appeared to fail as massed battles weakly favored the Germans [5].

[1] [5] make the point that the $g1$, $b1$, $g2$ and $b2$ have no justification and are used solely to facilitate modeling.

1.3 Agent Based Models

The evidence provided by [5] in support of the strategy employed by Park and 11 Group came from fitting the actual daily data from the Battle of Britain to the Lanchester model via regression analysis. While this provided insight in terms of the relevance of the Lanchester model, the results do not provide insight into the dynamics that produced the result. In particular, there is no insight into how Park achieved the 'defender's advantage'. A source of this problem, as [5] points out, is that "the Lanchester models are spatially and temporally homogenous, allowing for no variation in unit type, terrain or tactics, command or control, skill or doctrine". These assumptions appear inconsistent with modern warfare, which is ultimately dynamic and heterogeneous.

The approach utilized by [5] saw the force size estimated by the number of sorties flown by each side on a particular day. [5] indicates that ideally the data would be per raid. However, this was not possible due to the lack of records[2]. While the data shows the proportional loss rate of both sides, importantly it does not convene the loss rate per battle contact, as many sorties did not engage the enemy for a variety of reasons. Therefore, the true performance of the RAF against the Luftwaffe is lost.

An alternate approach is to implement an ABM that is capable of creating a virtual Battle of Britain. The model can be designed to explore the various tactics, and in particular whether the tactics of Park were indeed more effective than those of Leigh-Mallory. Agent-based models (ABMs) allow for the interaction of individual agents (aircraft/pilots in the proposed model), who undertake actions based on the context of their environment using basic rules. [4] successfully demonstrated the ability of an ABM to analyze air combat by creating a model of the Falklands War air battle. The model produced results consistent with what was observed in the conflict and tested various scenarios by varying the capabilities of both the U.K. and Argentinean forces.

For this paper certain abstractions were made to ensure that the research questions could be addressed in a timely manner. To achieve this, agents perform simplified actions that are supportable by fact or theory. The justification for the abstractions is that the aim of the model was to better understand the consequences of changing the number of squadrons per wing while removing noise from other factors. While this approach may not be fully authentic, it is more realistic than the Lanchester model and further iterations of the model can enhance the level of authenticity.

The level of abstraction means the model is not a one for one simulation, with a tick accounting for approximately 30 seconds. Determining the actions of each plane within a 30 second window is all but impossible, hence the simplifications and assumptions. Other abstractions include the weather having no impact,

[2] While [5] were able to provide supporting evidence that binning data by day rather than raid did not invalidate the approach, this author feels an alternative approach is warranted.

there being a 100% chance of the RAF making contact with an incoming raid[3], and the dogfight algorithm being simplified with concepts such as the role of a wingman removed.

The level of abstraction and available data did present a problem in terms of calibrating the model. Using the data from [5], the Luftwaffe and RAF losses and the number of sorties for each day are known. However, the actual loss rates per combat interactions are not known, which is what the model is actually simulating. An alternate approach was to review the diaries of the individual RAF squadrons. However, these tended to overstate the success rate of the pilots.

2 Model Design

2.1 RAF Forces

The objective of the implemented model was to have two forces; the RAF and the Luftwaffe, engage in an aerial battle over the English Channel, with the RAF fighters defending and trying to disrupt the incoming Luftwaffe attack. The various variables and agents associated with the RAF are summarized in Table 1. To allow altitude to be a consideration, the model was implemented in the 3D version of Netlogo [11].

2.2 German Forces

Table 2 summarizes the variables and agents associated with the Luftwaffe.

2.3 Model Functionality

The model's objective, and therefore its functionality, is centered around being able to answer the question of how the RAF could best arrange their forces to maximize the damage to the Luftwaffe, while minimizing their own damage. Therefore, at a high level the model must account for a defending force finding and then engaging the enemy, plus an offensive force that moves towards their assigned targets, that also has the ability to defend itself. In addition, the output of the model needs to provide key insights into the dynamics involved in producing the results in a more meaningful manner than the Lanchester model.

[3] In reality, this was not the case, as some RAF sorties were patrols that did not make contact with the enemy or were scrambled to meet a raid but failed to make contact. However, given the intent was to analyze actual combat performance the decision was made to ensure contact was made between the two forces.

Table 1: RAF Variables and Agents.

Variable	**Purpose**
	Variables
number_of _wings (wings)	The user selects the number of wings, between 1 and 5, that the FC scrambles. Each wing is assigned to a rally point and then has the fighters of a squadron(s) deployed to it. As per RAF standards, a squadron consists of 12 aircraft.
squadrons_per _wing (sPW)	The user decides how many squadrons are assigned to each wing, as determined above. The option is again between 1 and 5. Therefore, the user can test Leigh-Mallory's single 'Big Wing' (5 squadrons per wing, which means scrambling 60 fighters) compared to Park's smaller multiple wings (1–2 squadrons per wing or 12–24 fighters per wing).
number_of_ home_bases	The user sets the number of home bases that the FC forces are spread across. This allows the model to test for the implications of forming a large wing with fighters from multiple bases. This was a key consideration of the 'Big Wing' approach [10]. In the actual battle, 11 Group had 27 squadrons who had access to 25 airfields, while 12 Group had 15 squadrons spread across 12 bases [2].
	Agents
Rally points	When initializing the RAF force, rally points are created first being spread evenly across the y axis (longitude) while having the same x (latitude) and z coordinates (altitude). These settings are independent from the coordinates of the incoming raid. As part of their initialization, RAF fighters are "hatched" by their rally point and allocated to a home base. Fighters form up at their rally point coordinates via the scramble routine. Rally points act as radar stations and direct their aircraft towards the enemy via the search routine, which covers 40 patches in a 360-degree arc.
Home bases	Home bases are spread evenly across the y axis within the British mainland. RAF fighters are assigned to a home base(s) nearest to their rally point. All fighters start a simulation at their home base.
Hurricanes	The RAF has two classes of fighters. Hurricanes attack incoming bombers while Spitfires attack enemy fighters. This is consistent with the tactics of 11 Group [5]. Two key variables the fighters own are `status` and `evading?` The combination of these two determines the actions of a fighter. `evading?` has the value of **true** or **false**, while status can take the following values: **scramble**, **formation**, **engaged**, **searching**, **homebound** or **shot down**.
Spitfires	Spitfires pursue enemy fighters. Note in this version of the model the different performance characteristics of the two fighters was not included.

Table 2: Luftwaffe Variables and Agents.

Variable	Purpose
	Variables
ratio_to_RAF (r2RAF)	The size of the Luftwaffe force is set as a ratio to the RAF force that has been scrambled to meet the incoming raid. In reality, the reverse would be true, but it ultimately makes no difference to the model. The ratio varied greatly throughout the battle as both sides altered tactics [1].
ration_fighters _bombers (rationGFGB)	The composition of a raid can be varied by the ratio of enemy fighters to enemy bombers. During the battle this ratio ranged from 3 fighters per bomber up to 5 [1].
number_of _waves (waves)	Sets the number of waves that the Luftwaffe is sending during a particular raid. [2] provides support for this value ranging between 1 and 3.
number_of _targets (targets)	This sets the number of targets that the Luftwaffe pursue. The combination of targets and waves influences how compact or otherwise the raid is.
	Agents
Targets	Targets are initialized with randomly created x and y coordinates within Britain. The coordinates are then assigned to the Luftwaffe force who track towards their allocated target.
Waves	The initial class created for the Luftwaffe is a wave. Based on waves, their coordinates are set evenly along the y axis. Next aircraft are created for each of the waves. The number per wave is based firstly on the ratio of the Luftwaffe forces to the RAF (r2RAF), and then the number of waves. The aircraft's x and y coordinates are spread out in a formation around their wave's coordinates, while the fighters have a higher altitude (z coordinate) than the bombers.
Bombers	The bomber class has the role of tracking to their target and once they reach it, dropping their bombs before returning home. The number of bomb hits is recorded in an attempt to judge the success of a raid. Bombers are also capable of defending themselves against the fighters. Bombers do not own the evading? variable and have formation, homebound or **shotdown** as their possible status settings.
Fighters	The fighter's role is to defend the bombers. Each fighter has a tolerance variable and when the number of RAF fighters within a 3-patch radius exceeds their tolerance, they will break off and attack the RAF. The Luftwaffe fighters have the evading? variable. Their status includes **formation**, **engaged**, **shotdown** and **homebound**.

The RAF fighters' role was to intercept the incoming Luftwaffe wave(s) and destroy as many aircraft as possible, while avoiding being shot down, before returning to base. Figure 1 provides a flow chart of how the behavior of the RAF fighters was designed to meet this requirement along with other considerations of air combat.

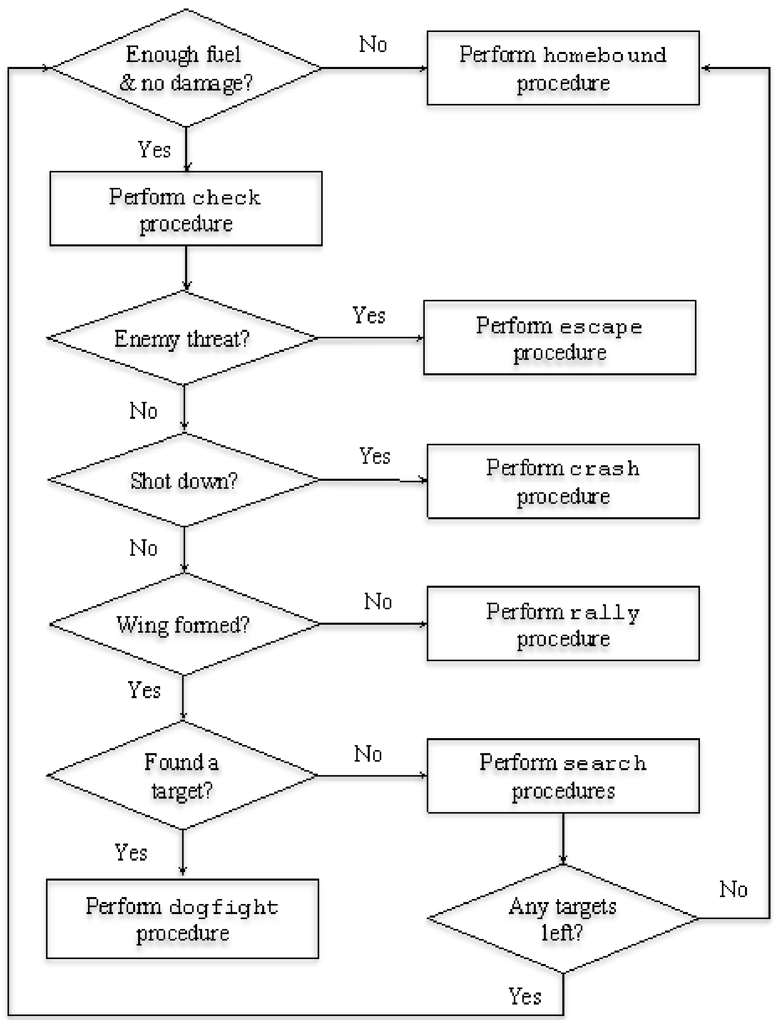

Fig. 1: Flow chart for the RAF fighters.

The Luftwaffe fighters' primary role was to escort and defend their bombers. As mentioned above, the fighters would remain in formation until the number of

RAF fighters around them exceeded their tolerance, at which point they would break off and attack the RAF. The Luftwaffe bombers' key objective was to find their target, drop their bombs and then return home. An abbreviated illustration of how the Luftwaffe fighters' operated is provided in Figure 2.[4]

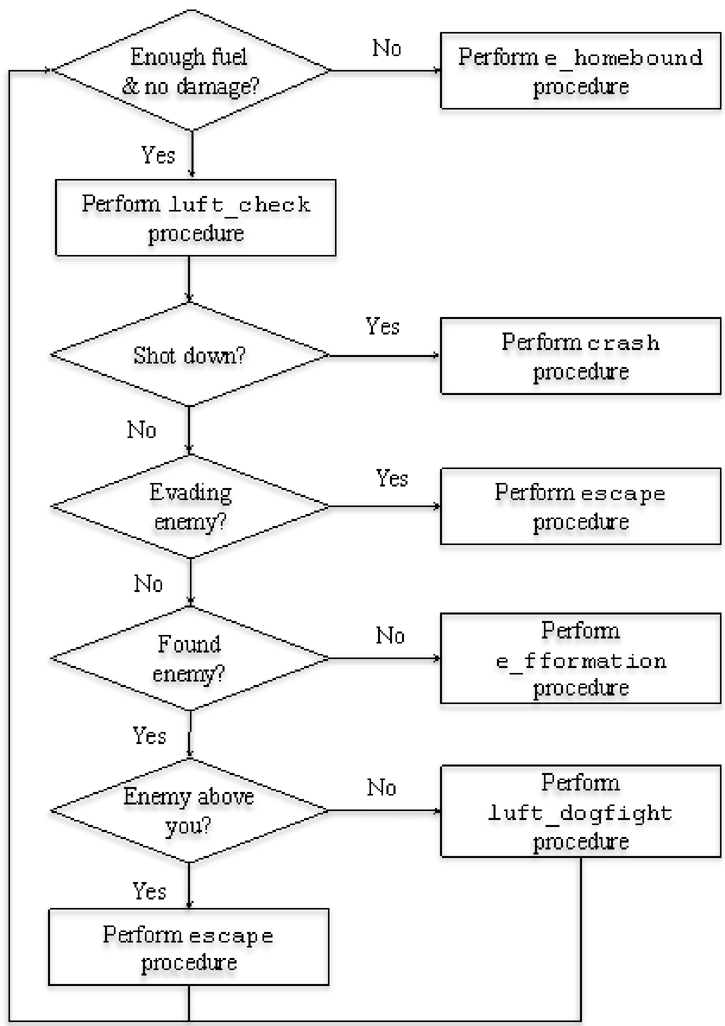

Fig. 2: Flow chart for the Luftwaffe fighters.

[4] Given the simplicity of the bombers role, it was felt a flow chart was unnecessary.

The behavior of the agents is controlled by a combination of their `status` and `evading?` Variables. With the two forces having different objectives and procedures they require different set of statuses. The RAF fighters maintain more possibilities because they are required to scramble, form their wing, find the enemy and engage. In contrast, the Luftwaffe fighters are already airborne and only need to find the enemy, before engaging. The requirements for the bombers are simpler again; as they head towards their target and defend themselves when attacked but do not change course. The agent sets do have some similar statuses, such as; shotdown and homebound. They also have some similar procedures, such as; checking whether they are being attacked and whether they have enough fuel to return home.

A high level flow of the model and therefore a description of Figs. 1 and 2 is;

- Both sides are initialized as per the user settings (see Sects. 2.2 and 2.3),
- At each step each plane will check their fuel and damage. Their `status` is updated to **homebound** and they will head home if their fuel levels are just sufficient to get them home or their damage level is over 0.5. The model reduces the fuel for the agents by 1 unit at each tick. The agents also check that their status is not **shotdown**. If they are shot down they perform the `crash` procedure,
- At each tick fighters also check they do not have an enemy plane on their tail via their `check` procedure. If they do, the `evading?` variable is set to **true** and their opponent is set to 'nobody'. If a plane is required to evade it will implement one of 5 strategies to move away from their threat. A plane cannot attack while they are evading the enemy. This procedure is consistent with the standard air battle tactic of breaking off an attack if you are in immediate danger of being shot down,
- Through the `rally` procedure the RAF scrambles their forces with the fighters climbing towards their assigned wing's rally point. All fighters have the `status` of **scramble** at this point and cannot engage the enemy,
- Once all the RAF fighters have reached their assigned rally point, the fighters have their `status` updated to **formation** and the wing is directed towards the incoming enemy wave(s) by their rally point and they are now able to attack. Rally points become aware of the exact coordinates of the incoming wave via their `search` routine which covers an area of 40 patches in a 360-degree arc,
- Once the rally point confirms the raid, through its `search` procedure, it will assign an enemy opponent to each of their fighters. The `status` of the fighters is updated to **engaged**.
- Meanwhile the Luftwaffe fighters have the `status` **formation**. At this point their `e_fformation` procedure has them moving towards the allocated targets of their bombers, while also being on the "look out" for the RAF. If they spot the RAF and the number of RAF fighters within a 3-patch radius is greater than their tolerance, their `status` is updated to **engaged**. However, if the RAF is above them they will take evasive action, through the `escape` procedure, before pressing an attack on the enemy. This is true

of all offensive actions for both fighter forces and is an example of how the model utilizes basic air combat procedure,

- Once a fighter's `status` is updated to **engaged** their `dogfight` procedure is called with the sequence of an attack being; the attacking plane will set their heading and pitch to intercept their assigned opponent. This will either be a head on attack or the fighter will chase down their opponent from behind. When the attacker is close enough, within a 1 patch radius and within a 10-degree cone of sight, they will fire upon their opponent. The damage inflicted upon the opponent in the attack is a random float up to the value of 0.9. If a plane sustains damage above their survival level (.9) their `status` is updated to **shotdown**. The amount of damage inflicted per attack was determined by calibrating the model such that a representative battle of 24 RAF aircraft against 36 Luftwaffe aircraft recorded a similar loss ratio as an actual battle, which was around 5–10 percent,
- After an attack, the attacking fighter's `status` is changed. The RAF fighters are changed to **searching** while the Luftwaffe fighters change to **formation**, a status that ensures that the Luftwaffe is searching for the RAF. There are some minor differences in the search routines for the two fighter forces. While the RAF fighters check themselves for enemy fighters within 4 patches, if they cannot find any, they ask their rally point whether they are tracking any enemy planes, remembering the rally point has a broader search arc. If the rally point is tracking enemy targets, then it provides the coordinates of an enemy aircraft to their fighter, otherwise the battle is considered over and the RAF returns home via the **homebound** procedure. This process is consistent with the RAF having the advantage of radar to assist in finding the enemy. In contrast, the Luftwaffe fighters are solely responsible for finding their own target and will return home after reaching the target, unless have sustained damaged or are running low on fuel,
- Enemy bombers continue towards their target with a number of fighters remaining in support. Bombers are able to protect themselves at each tick through their `defend` procedure. In this procedure a bomber selects any two fighters within a one-patch radius and fires upon them. To reflect the lower probability of hitting a fighter the damage inflicted is a random float up to .05, and
- When bombers reach their target they "unload" their bombs and return home. It should be noted that this step is simply a checkpoint with no consideration given to the amount of bombs that hit the target in this iteration of the model.

As detailed above, to achieve the objectives of the research question there was a certain amount of abstraction undertaken. However, while the actual movements of the aircraft might not match the exact characteristic of an air battle, the actions are supportable given the objectives of the agents and basic air battle tactics. There will always be the need for some abstraction in an ABM; otherwise you have you have moved beyond an ABM into an engineering model.

Verification of the model's behavior was undertaken by performing parameter sweeps on extreme values, with the results analyzed to ensure that the model was performing as per design. Extensive visual inspections were also undertaken, with the various agent classes color-coded based on status to ensure updating occurred as per the design.

3 Experiments

To understand the possible influences on the losses for both sides, a full factorial experiment (Experiment 1) combined with an analysis of variance (ANOVA) as outlined in [8] was undertaken. The design matrix, seen in Table 3, was used and generated 128 combinations, with each combination run 50 times in the simulation. From these settings the largest battle was 300 RAF fighters up against 1,200 Luftwaffe planes. While a battle like this did not occur, it highlights the benefits of creating a simulation capable of exploring the outcomes of such a battle. The mean of both the input and output variables from each combination was taken to form the values used in the ANOVA model. A principle component analysis (PCA) was undertaken on the data as well.

Table 3: Design matrix for the full factorial experiment.

Variable	**Low setting**	**High setting**
`numbaer_of_waves`	1	3
`number_of_wings`	1	5
`squadrons_per_wing`	1	5
`ration_fighters_bombers`	3	5
`ratio_to_RAF`	1	4
`number_of_home_bases`	2	4
`number_of_targets`	1	3

The settings in Table 3 are supported by the descriptions of various battles provided in [2], [5] and [1]. In particular, the following points are relevant:

- The 'Big Wing' debate is all about determining whether 1 or 5 squadrons was the correct number of squadrons per wing. In addition, 11 Group had the flexibility of sending multiple wings, while 12 Group was restricted to 1. To maintain symmetry the range was set at 1 and 5, but it must be acknowledged that the RAF never deployed 5 wings of 5 squadrons,
- The size of the Luftwaffe force to the RAF varied throughout the course of the battle. Small raids, a `ratio_to_RAF` of 1, were used at the commencement of the battle before larger raids (a ratio of 4) were employed later in the battle,
- As mentioned previously the German's varied the ratio of fighters to bombers within the range of the experiment,

- The range of waves and targets is consistent with records of the battle, and
- Each RAF group had their planes spread across multiple bases, meaning that an intercepting wing was unlikely to be all from the same base; hence this variable ranges from 2 to 4. The author will concede that a combination of 1 squadron scrambling from 4 bases would not have occurred. However, the implications are minor, if any in this version of the model.

To assess the effectiveness of the 'Big Wing' approach, the results from 1 wing of 5 squadrons (1W5Ss) was compared against results from 5 wings of 1 squadron (5Ws1S) in Experiment 2. Each strategy was tested against an increasing German force ratio (`r2RAF`), with the ratio beginning at 1, moving to 4 in increments of 0.1. Fifty runs of the model were made at each ratio setting[5]. This scenario may not be 100% consistent with how 11 Group used their forces. However, to create a valid comparison, the author felt it was appropriate and necessary to ensure the British force size was consistent at 60 RAF fighters.

An ordinary least squares (OLS) model as per Equation 3 was fitted to the output of Experiment 1 and 2. Equation 3 provides the model for the British loss rate[6]:

$$\log\left(-\frac{dB}{dt}\right) = \log g + g1 \log GermanForces + b1 \log BritishForces \quad (3)$$

This approach is consistent with fitting the data to the aimed fire Lanchester model and replicates the analysis provided in [5]. An analysis of covariance (ANCOVA) was undertaken to establish whether a statistical difference existed between the resulting models from Experiment 2. All analysis was undertaken in R [9].

4 Results

Figure 3 presents the Biplot resulting from the PCA analysis. From the chart it can be seen that the first component (PC1), which had an explanatory power of 33%, relates mostly to the size of the forces engaged in the battle (RAFSize and GermanSize). Both the number of wings deployed and the number of squadrons per wing make a contribution. The second component (PC2), which explained 16.3% of the data, relates primarily to the ratio of the two force (`r2RAF`) contrasted against the number of German losses. This indicates that as the ratio of German planes to the British increased, their losses tended to decrease. This finding is inconsistent with what the Lanchester model prescribes. Further tests explore this finding.

Also from Figure 3 it can be seen that there is a clear division between where the data for raids with a 1:1 ratio sits (bottom half) compared to a 4:1 ratio (top

[5] The other settings used were 3 `homebases`, `ratio_fighters_bombers` 3, `number_of_waves` 2, `number of targets` 1 and `ratio_spitfires_hurr` 1.

[6] For Experiment 2 the British force was held constant, therefore the $b1$ term was dropped.

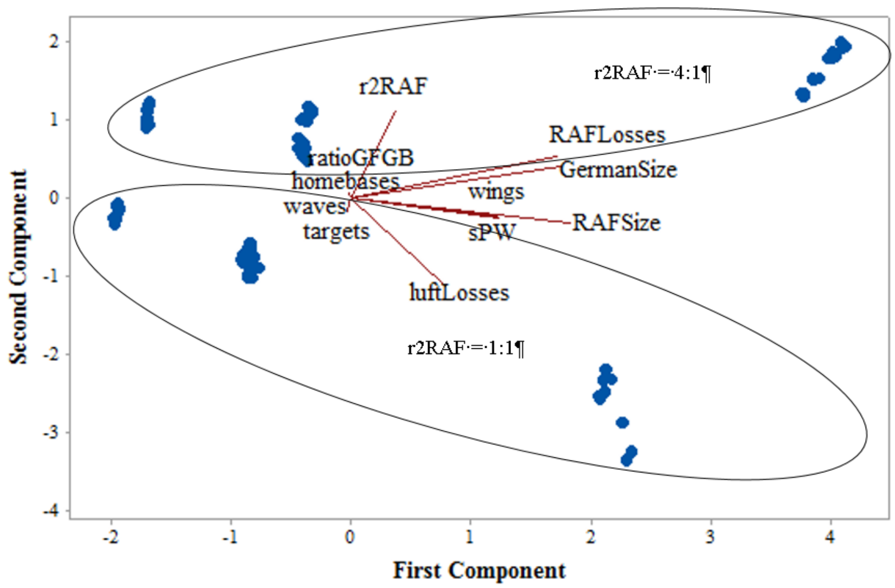

Fig. 3: Biplot resulting from the PCA analysis of the data from Experiment 1.

half). The implications being that the British needed to match the force size of the Germans because while it increased the overall size of the battle, their losses were relatively lower when their ratio was closer to the Germans.

Table 4 presents the results of applying Equation 3 to the full experimental data set. It should be noted that despite using the logs of the variables, both models failed the test for normality with regards to their residuals; hence the results are not robust. The British model, which returned an R^2 of 96.3%, is consistent with the Lanchester model in that the British losses scale positively with increasing force sizes from both sides, albeit at a rate less than one (the assumption of the aimed fire model). The results for the German losses are not consistent with the Lanchester model, with the model returning a negative coefficient for the impact of an increase German force. The interpretation of this result is that the German's benefited from safety in numbers – an increasing return to scale for safety. This result was most likely driven by the improved defensive performance of massed bombers and is consistent with the interpretation of the PCA analysis. The R^2 for the German model was 60.2%.

The results of the ANOVA and the subsequent OLS coefficients as per the approach of [8] are contained in Table 5. The dependent variable for the model was the British loss ratio (actual losses / the number of sorties), not the log of the actual losses, as per the previous model. The rationale for the change is that the ratio normalizes the outcome across the various settings, thus enabling the identification of the key drivers. The R^2 of the regression model was 98.1%, with residuals meeting normality requirements.

Table 4: The results of fitting an OLS model to the data from Experiment 1.

Side	Variable	Estimate	Std.Error	t value	Pr(>|t|)
British	(Intercept)	-2.2732	0.0584	-38.90	0.0000
losses	g1	0.8784	0.0502	17.49	0.0000
(Log)	b1	0.7854	0.0588	13.36	0.0000
German	(Intercept)	-0.0529	0.1158	-0.46	0.6488
losses	b2	1.5201	0.1165	13.04	0.0000
(Log)	g2	-0.8842	0.0995	-8.88	0.0000

From Table 5 it can be seen that there is significant interaction between `r2RAF` and the other variables, supporting the hypothesis that the FC needed to consider more than just force size in determining their strategies. This result is consistent with the influence of `r2RAF` in (PC2), as seen in Figure 3. Other observations from Table 5 are:

1. Negative values for both `r2RAF:waves` and `r2RAF:targets` suggest that if an incoming raid is spread out, it benefited the British through a lower loss ratio;
2. Increasing the number of `wings` and `sPW` increased the British loss ratio as the Germans increased their force. This result provides mixed evidence in answering the 'Big Wing' debate. Experiment 2 provides greater insight on this point, and
3. The composition of the raiding party, the ratio of German fighters to bombers, the number of RAF home bases, were not significant factors.

Table 5: Results of the effect model fitted to the full factorial data set.

	Estimate	F-Value	t value	Pr(>|t|)
(Intercept)	0.2361	–	86.33	0.0000
ratio_to_RAF (r2RAF)	0.1290	2224.59	47.17	0.0000
number_of_wings(wings)	0.0945	1194.25	34.56	0.0000
squadrons_per_wing (sPW)	0.0999	1335.22	36.54	0.0000
number_of_waves (waves)	-0.0520	361.59	-19.02	0.0000
number_of_targets (targets)	-0.0192	49.48	-7.03	0.0000
r2RAF:wings	0.0459	282.33	16.80	0.0000
r2RAF:sPW	0.0564	425.88	20.64	0.0000
r2RAF:waves	-0.0352	165.83	-12.88	0.0000
r2RAF:targets	-0.0108	15.67	-3.96	0.0001
wings:sPW	0.0205	56.07	7.49	0.0000
waves:targets	0.0083	9.11	3.02	0.0031

Table 6 provides the results of fitting an OLS model explaining British losses as per Equation 3 for the different strategies. The data was generated from Experiment 2 with the data illustrated in Figure 4.

Table 6: Results of the effect model fitted to the full factorial data set.

Model	Variable	Estimate	Std.Error	t value	R^2
5Ws1S	(Intercept)	-0.6933	0.0622	-11.14	0.9682
	g1	0.8584	0.0289	29.69	
1W5Ss	(Intercept)	-1.0761	0.0616	-17.47	0.9777
	g1	1.0195	0.0286	35.64	

From Table 7 the results from the analysis of covariance (ANCOVA) and indicates that the interaction of the German force's size, and the number of wings (GF:Wing), is significant. This supports the hypothesis that the number of wings employed did indeed impact the loss rate of the British.

Table 7: Results of the ANCOVA testing for the significance of the two strategies.

	Estimate	F-Value	t value	Pr(>|t|)
(Intercept)	-1.0761	0.0619	-17.38	0.0000
g1	1.0195	0.0288	35.45	0.0000
Number of Wings (NW)	0.3828	0.0876	4.37	0.0001
LGF:NW	-0.1612	0.0407	-3.96	0.0002

Figure 4 illustrates the results of Experiment 2 by showing the relationship between the British and German losses versus the size of the German force, remembering that the German force increased against a set number of British fighters (60). Consistent with the findings from Experiment 1, an increasing German force results in greater British losses but lower German losses.

The results of Experiment 2 suggest that the rate of British losses scales at greater than one on average when the 'Big Wing' (1W5Ss) is employed, and less than one for the smaller wings (5W1S). This finding in isolation is indeed supportive of the strategy of Park. However, the smaller wing approach has a higher intercept value, with the interpretation being that the 'Big Wing' has a lower fixed cost yet higher variable cost of doing battle, while the smaller wing has the opposite. A similar analysis was undertaken on the German losses (Figure 4) and there was no significant difference in the damage the British inflicted on the Germans by the two formations. The conclusion being that the British were limited in their ability to inflict greater losses on the Germans. Indeed, from the findings of Experiment 1, increased British successes were reliant on the German's waves spreading out.

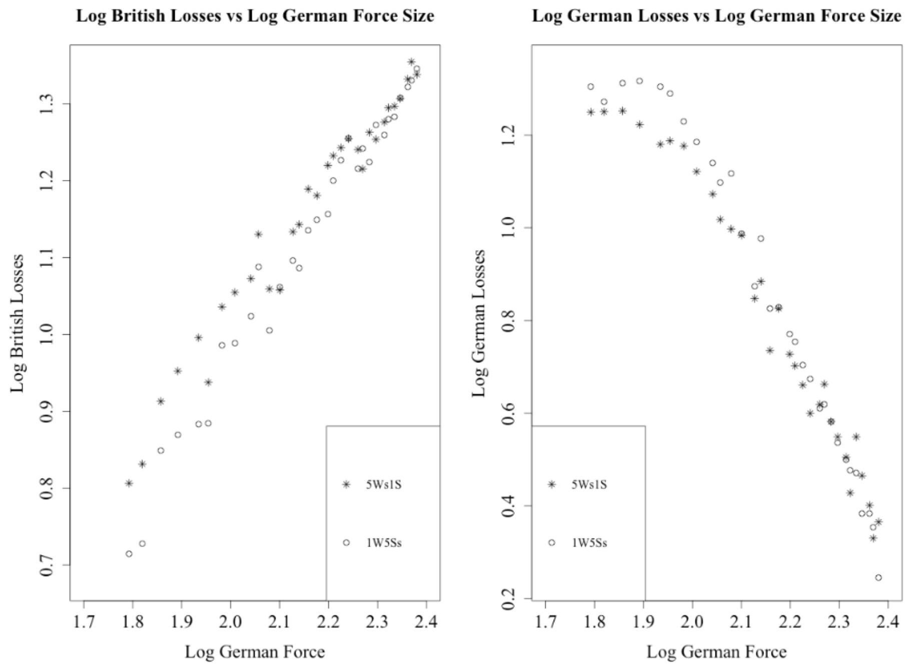

Fig. 4: **(Left)** Log plot of British Losses vs the log of the German forces. **(Right)** Log plot of German Losses vs the log of the German forces.

Combining this inference with those from Experiment 1, it can be concluded that for the ranges that were tested for, the loss ratio of the British was higher under the smaller wing approach due to the average cost of the smaller wings being higher. Additionally, given there was no benefit from attacking in a smaller wing, it appears that when facing a German force equal to four times larger than the RAF force on average, the 'Big Wing' was the right approach.

5 Conclusion

The results obtained from the first known ABM of the Battle of Britain are supportive of the 'Big Wing', albeit only in terms of minimizing British losses. On the flip side, causing greater damage to the Luftwaffe rests on denying them the opportunity to achieve safety in numbers, achieved through matching force size or ensuring the incoming formation is forced to spread out. Further investigation should focus on the role of the cost functions identified for the different wing formations. The need also exists to investigate the possibility that the actual time a wing is engaged in combat may be a factor, something that would have been detrimental to smaller wings, who engaged quicker and longer in this model. Following this, the optimal combination of wings and squadrons per wings that reduces British losses and maximizes damage to the Luftwaffe can be solved for. Importantly, such an investigation is beyond the Lanchester model but well within the abilities of an ABM.

References

1. Bickers, R.T.: Battle of Britain. Salamander Books, London (2000)
2. Bungay, S.: The most dangerous enemy: a history of the battle of Britain. Aurum, London (2009)
3. Churchill, W.: The few. http://www.winstonchurchill.org/resources/speeches/1940-finest-hour/113-the-few (1940)
4. Gowlett, P.: Assessing the military impact of capability enhancement with Netlogo using the Falklands War as a case-study. http://www.mssanz.org.au/modsim2013 (2013)
5. Johnson, I.R., MacKay, N.J.: Lanchester models and the battle of Britain. Naval Research Logistics 58(3), 210–222 (2011)
6. Lanchester, F.W.: Aircraft in warfare: The dawn of the fourth arm. Constable (1916)
7. MacKay, N., Price, C.: Safety in numbers:. Ideas of concentration in Royal Air Force Fighter Defence from Lanchester to the battle of Britain 96(323), 304–325 (2011)
8. NIST/SEMATECH: Nist/sematech e-handbook of statistical methods. http://www.itl.nist.gov/div898/handbook/ (2015)
9. R Core Team: R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna (2015), http://www.R-project.org
10. Sarkar, D.: Bader's Duxford fighters: The Big Wing controversy. Ramrod Publications, St Peter's, Worscester (1997)
11. Wilensky, U.: NetLogo. http://ccl.northwestern.edu/netlogo/ (1999)