

Spatio-temporal patterns of IED usage by the Provisional Irish Republican Army

STEPHEN TENCH^{1,2}, HANNAH FRY² and PAUL GILL¹

¹ *UCL Jill Dando Institute of Security and Crime Science, 35 Tavistock Square, London, WC1H 9EZ*
email: stephen.tench.13@ucl.ac.uk

² *UCL Centre for Advanced Spatial Analysis, 90 Tottenham Court Road, London, W1T 4TJ*

(Received)

In this paper a unique dataset of improvised explosive device (IED) attacks during “The Troubles” in Northern Ireland (NI) is analysed via a Hawkes process model. It is found that this past dependent model is a good fit to IED attacks yielding key insights about the nature of terrorism in NI. We also present a novel approach to quantitatively investigate some of the sociological theory surrounding the Provisional Irish Republican Army (PIRA) which challenges previously held assumptions concerning changes seen in the organisation. Finally we extend our use of the Hawkes process model by considering a multidimensional version which permits both self and mutual-excitations. This allows us to test how the PIRA responded to past IED attacks on different geographical scales from which we find evidence for the autonomy of the organisation over the six counties of NI and Belfast. By incorporating a second dataset concerning British Security Force (BSF) interventions, the multidimensional model allows us to test counter-terrorism (CT) operations in NI where we find subsequent increases in violence.

1 Introduction

Terrorism is a major international concern which shows little signs of abating. There is therefore great importance in developing scientific approaches to understand the behavioural underpinnings of terrorism in order to prevent and disrupt these activities. The ability to gain such insights through real world experimentation is questionable due to the risks associated with unsuccessful approaches. One of the cheapest and most adaptable methods of research in this area is mathematical modelling [33]. Such modelling provides not only a vast number of well-developed tools and techniques but also the opportunity to experiment freely without unnecessary safety risks or ethical concerns.

Several advances have been made in this field of late, leading to a variety of conclusions with policy making implications. Braithwaite and Johnson [7] studied the interactions of insurgent attacks and Coalition counter-insurgency operations in Iraq. The authors were able to conclude from space-time patterns that indiscriminate counter-insurgency operations resulted in a backlash effect by insurgents whilst discriminate operations had the opposite effect. Along similar

lines Lewis et al. [29] apply self-exciting point process models to study violent civilian deaths in Iraq during the U.S.-led invasion. They found a two to six month timescale for violent deaths which correspond to a series of related attacks. Hence, quick interventions could help to drastically lower the problem of violent deaths in Iraq. In a similar spirit a study by Mohler [35] found evidence that terrorist attacks in Northern Ireland followed a pattern of self-excitation lasting 9.3 weeks. The line of investigation taken by these sorts of studies have important consequences for tackling the types of terrorism seen in NI where insurgency and civilian deaths were major issues [20]. These approaches marked a large departure from political science-inspired methodologies that generally linked the quantity of terrorist attacks to “root causes” like socio-economic indicators that are quasi-static and fail to provide insight into the triggers needed for a strategic intervention.

The methods employed in these terrorism studies [7, 20, 29, 35] along with several others [14, 22, 55] share similarities with the modelling of spatio-temporal phenomena in crime pattern theory [8]. In that context it is assumed that criminal activity forms a series of quantifiable patterns at the macro scale [16, 44]. Mathematical descriptions of these patterns, or more generally ‘crime hotspots’, can be explored and exploited in real-time. By targeting susceptible areas with preventative measures there is potential for great reductions in subsequent crime. This approach has since been shown to be effective in a number of real-world applications including policy [2, 17, 31], and predictive policing [6, 27].

It is from this point of view that we aim to approach the issue of terrorism during the conflict in NI. In this paper we seek to add to the literature concerning spatio-temporal patterns of terrorism by studying a unique dataset of IED attacks in NI between 1970-1998. Specifically for the case of NI this paper is focused on the group known as the Provisional Irish Republican Army (PIRA). Although there are extensive historical accounts and a growing body of social science research related to this group [3, 19, 25, 26, 51] there is a gap for a wider scope of mathematical investigations of their activities.

The PIRA was predominantly formed from members of the Catholic community in NI [56] and saw itself as “the legal representatives of the Irish people, [who] are morally justified in carrying out a campaign of resistance against foreign occupation forces and domestic collaborators” [39]. The active period of the PIRA between 1969-1998 can be traced out in five phases [3]. These phases and their historical context are described below.

- **1969-1976 - Phase 1:** During this phase the organisation was arranged in a military style having brigades, battalions and companies.
- **1977-1980 - Phase 2:** A cell-based structure was adopted. This approach was characterised by PIRA fracturing into small groups of members known as Active Service Units (ASUs) [25]. This approach aimed to improve the organisation’s secrecy by making it harder to infiltrate.
- **1981-1989 - Phase 3:** The Republican campaign moved into the political arena

through the Sinn Féin party who now had similar levels of prestige as their militant wing, PIRA.

- **1990-1994 - Phase 4:** Secret meetings involving top ranking PIRA leaders negotiating a ceasefire with the British Government.
- **1995-1998 - Phase 5:** Peace talks announced with a ceasefire ratified in the Good Friday Agreement signalling for many the end of “The Troubles”.

During its active phases the PIRA successfully developed a large arsenal of IEDs which it employed with devastating effects for both the security forces and civilians [53]. In particular, the impact of violence was felt heavily by the civilian population which constituted approximately 54% of all deaths [20]. Moreover, the deaths caused by the NI conflict were highly concentrated in Belfast, where approximately 47% of fatalities occurred [20].

In analogy to the studies of Braithwaite and Johnson [7], Lewis et al. [29] and Mohler [35] we make an attempt to understand the driving forces behind IED attacks in NI. Our extension to the present literature revolves around the access we have to a unique dataset of IED events which allows us to study a specific type of terrorist activity in great detail at a fine temporal scale. The dataset we use provides ample evidence for any past dependence on insurgent attacks and this study attempts to understand these dynamics in greater detail. The model chosen to explore this question is a Hawkes self-exciting point process. It makes use of a response function (or kernel) which holds information pertaining to the long-term influence of previous events and has been shown to well represent a number of past dependent processes including gang related violence [18, 24, 52], email exchanges to infer organisation leadership [21], burglary [34, 50] and violent deaths in conflicts as previously mentioned [29]. It also provides an opportunity to examine each of the five phases of PIRA activity separately, as is done in Model 1 and all subsequent models, offering direct quantitative insights into how the group behaved and reacted through the stages of the Republican campaign.

As well as studying temporal patterns of insurgency Lewis et al. [29] also compare Hawkes processes in different regions of Iraq to understand spatial influences seen during waves of violent attacks. Similarly in this paper we undertake spatial disaggregation of PIRA attacks according to the six counties of NI and we also separate Belfast due to its significance during “The Troubles” as discussed above [20]. This additional spatial information is to yield insights about the extent to which PIRA units in NI acted autonomously as suggested by Horgan and Taylor [25]. This latter study forms the theoretical basis for Models 2 and 3.

The PIRA did not act in isolation however, and in this contribution, we also aim to explore the interplay between PIRA and the British Security Forces (BSF) by employing a multidimensional Hawkes process. Previous studies into British counter-terrorism (CT) strategies in NI [12, 28, 40] found evidence that the actions of BSFs could undermine the effort to curb Republican terrorism and even result in a negative backlash increasing the number of attacks. For instance, in 1988 an operation by BSFs resulted in the deaths of three PIRA

members in Gibraltar. An analysis of subsequent PIRA attacks found positive increases related to the incident 36 months after it occurred [28]. The inclusion of additional mutually exciting terms in the Hawkes process has been seen to represent other interacting systems in the past including multiple gang networks [50], but, to the best of our knowledge, has yet to be applied to yield insights into CT strategies.

To test the CT strategies employed in NI two types of events will be assessed in this paper. Firstly, from the discussion above concerning the background of the PIRA it was pointed out the organisation drew the majority of its volunteers from the Catholic community [56]. Consequently deaths of Catholic civilians resulting from BSF operations may be expected to prompt a significant retaliation from the PIRA. We investigate this effect in Model 4. Moreover, following the case study findings of Lafree, Dugan and Korte [28], as discussed above, the backlash effect of BSF actions which killed PIRA members will be considered in Model 5. With such insights we aim to uncover a useful methodology by which CT practitioners and academics can judge the efficacy of past strategies to combat terrorism.

Alongside the contributions this paper aims to make concerning the use of the Hawkes process we also present a novel approach to deal with the issue we refer to as edge effects. This effect is the result of events outside the observation period influencing those inside [46]. Whereas in previous studies using the Hawkes process the data analysed is in a single time series [18, 29] with the segmentation of our data according to the five phases of PIRA it is possible that events in previous phases may have influenced those in future ones. As a consequence of this a moving time window approach was considered whereby the data points from adjacent phases were combined to find the quantitatively best fitting model.

The structure of presentation will be in 6 sections. In Section 2 a discussion of the datasets used for this paper will be provided. Then in Section 3 the mathematical models that have been studied will be introduced in more detail. This section will also contain further information about the method used to obtain model parameters. Next in Section 4 the numerical results of the paper will be provided. This will begin with a description of the novel approach we have taken to examining edge effects in the data analysed. Alongside the results a discussion will be provided about how the model parameters can be interpreted. Finally in Section 5 the findings of the paper will be discussed with their potential impacts for both the academic community and practitioners demonstrated as well as a brief overview of future research topics.

2 Data

This paper utilises a unique dataset of PIRA IED events from 1970 to 1998. This dataset was collected through an exhaustive coding of newspaper reports and other open source outlets. Please see Asal et al. [3] for a full outline of the data collection and verification process. In total, the dataset spans 5461 IED

Table 1. *PIRA IED Dataset Event Fields*

Field	Values
Date	Year (1970-1998)/Month (1-12)/Day (1-31)
Location	{Antrim, Armagh, Belfast, Derry, Down, Fermanagh, Tyrone}
Target	{Political, Military, Police, Paramilitary, Government, Transport, Civilian, Foreign}

Table 2. *BSF Dataset Event Fields*

Field	Values
Date	Year (1970-1976)/Month (1-12)/Day (1-31)
Religion	{Catholic, Protestant}
Status	{Civilian, British Security, Republican Paramilitary, Loyalist Paramilitary, Irish Security}
Organisation Responsible	{British Security, Republican Paramilitary, Loyalist Paramilitary, Irish Security, Unknown}
Geographical Location	{Belfast North/East/South/West, County Antrim/Armagh/Derry/Down/Fermanagh/Tyrone, Derry, Britain, Europe, Republic of Ireland}

events. For each event there are numerous details concerning the IED attack and groups involved. Details of the event fields relevant to our study are given in Table 1. It should be noted here that when discussing BSFs we refer to both military and police targets. For further details on all the information contained in the dataset the reader is referred to Asal et al. [3].

Alongside the study of PIRA related attacks in isolation, an additional investigation was made into how BSF attacks impacted upon further PIRA attacks. The dataset concerning BSF attacks was obtained from the Conflict Archive on the Internet (CAIN) [53]. In Table 2 the event fields and possible values for this dataset are presented. In this study only the events which occurred in NI were considered. There were a total of 131 Catholic civilian deaths recorded in the final dataset of BSF events. Of these entries 78 were found to correspond to Phase 1 of PIRA activity. This lead to a only a small number of data points being available in the other phases with 12, 34 and 7 points for Phases 2, 3 and 4 respectively and 0 for Phase 5. Due to this distribution of data, results were only obtainable for Phase 1. The number of PIRA IED events in Phase 1 targeting BSFs was 144. See Sections 4.8 and 4.9 for further details.

3 Methodology

3.1 Poisson Process

The first model studied serves as a baseline to compare with the Hawkes process [18]. The specific baseline model tested was a Poisson process model specified via a single parameter N/T , where N = number of events and T = time of the final event measured from time 0 [47]. The Poisson process assumes that each event is independent and thus that the system has no memory of attacks in the past.

3.2 Univariate Model

Moving to a more complex model this paper will consider the influence of self-excitations of PIRA IED attacks. The method we used is based on studying an intensity function which describes the rate of IED incidents as a function of time and conditional on the past history of events. The intensity function takes the following form [23] for a given set of event times $\{t_i\}_{i=1}^N$.

$$\lambda(t) = \mu + k_0 \sum_{t > t_i} g(t - t_i; \omega) \quad (3.2)$$

The response function g is taken to be of the form $g(t) = \omega e^{-\omega t}$. The exponential form for the response function is routinely used in studies of crime and insurgency data [18, 29, 52]. It makes good intuitive sense for events which are clustered in time and allows for a physical interpretation of each component of the intensity function. In the example of IED attacks the constant μ can be considered as a background rate at which IED events occur. After an initial IED attack there may be further attacks, for example, a PIRA unit may wish to follow up on the success of a previous attack, and the constant k_0 captures the jump in the IED event rate. However, an indefinitely higher rate is unrealistic and eventually the rate will return to the background rate. The rate of decay is controlled by the term ω . The additional ω preceding the exponential term acts as a normalisation constant so that the jump factor multiplied by the response function can be viewed as the number of offspring after an event and the density of the time interval for the increase in activity [46].

Parameter estimation for the intensity function can be undertaken via the method of maximum likelihood estimation (MLE) [42]. This process involves finding the parameters which maximise the following log-likelihood function derived by Rubin [48]:

$$\log L(\{t_i\}; \mu, k_0, \omega) = \sum_{i=1}^N \log(\lambda(t_i)) - \int_0^T \lambda(t) dt,$$

where $t_N = T$ will be taken as the final time of observation in a similar approach taken by Ozaki [42]. For the form of the intensity function given in

(3.2) the log-likelihood becomes

$$\log L = \sum_{i=1}^N \left[\log \left(\mu + k_0 \sum_{t_i > t_j} \omega e^{-\omega(t_i - t_j)} \right) + k_0 \left(e^{-\omega(T - t_i)} - 1 \right) \right] - \mu T.$$

A number of assumptions are built into this formulation. First, all parameters used in the intensity function should be positive [29] to ensure the model remains realistic. Second, the set of points $\{t_i\}_{i=1}^N$ should be measured from time zero. However, since the Hawkes process depends on the infinite past, this assumption is not achievable in a real world setting and it may be difficult to eradicate the influence of events outside the observation period on those inside [46]. More details on how this problem was handled for this paper will be discussed in Section 4.1.

A further assumption is that the set of times should be unique [30]. Our dataset contains several simultaneous events and thus to satisfy the requirement for uniqueness, events in the same county, or multiple events in Belfast, on the same day were regarded as a single event. However, to avoid losing too much detail, events in different counties, or events inside and outside of Belfast, on the same day were distinguished via the addition of a random timestamp (as in Bowsher [5]). This is justified by looking at the command and functional structure of the PIRA which reveals that at the county and Belfast levels IED attacks were fairly autonomous [25].

Finally, a restraint on the response function g ensures that the model is non-explosive (see Varadhan [54] for further details concerning explosive stochastic processes). This assumes that the integral of g over t should be strictly less than unity [42]. Our choice of g satisfies this condition.

3.3 Multidimensional Model

After examining the past dependent nature of IED attacks based on self-excitations the second type of model investigated will also include mutual-excitations. For example, such models will be used to consider the influence of PIRA attacks and BSF attacks on further PIRA attacks.

The multidimensional Hawkes process model can be defined in a similar way as was done for the one-dimensional case. Here, with two adversaries, we require a two-dimensional model. Now there are two sets of event times which will be labelled $\{t_a\}_{a=1}^N$ and $\{t_b\}_{b=1}^M$ and two counting processes, $N_r(t)$, $r \in \{1, 2\}$, which form a two-dimensional counting process $\mathbf{N}(t)$. Each individual process has intensity function defined by [23]

$$\lambda_r(t) = \mu_r + k_0 \sum_{t > t_a} g(t - t_a; \omega) + s_0 \sum_{t > t_b} h(t - t_b; \nu), \quad (3.3)$$

where the two response functions are defined by $g(t) = \omega e^{-\omega t}$ and $h(t) = \nu e^{-\nu t}$. The form chosen for the response functions is chosen by analogy to the research of Short et al. [50]. In particular, in extending the work of Egesdal et al. [18],

to study interactions between multiple gangs, results were obtained indicating that exponential response functions may prove useful to study mutual-excitations between gangs.

Similar to the univariate case described in Section 3.2 we can interpret the model in a real world setting. In the case of modelling influences on PIRA attacks one could, for example, take event times $\{t_a\}$ to represent times of PIRA IED attacks and times $\{t_b\}$ as BSF attacks. The background rate μ_r , jump rate k_0 and response function g have the same interpretation as that given in the one-dimensional case. Similarly the parameter s_0 represents the jump in IED attacks following a mutual-excitation, such as, retaliation against a BSF attack whilst ν controls the temporal scale over which this mutual-excitation persists. The parameter ν also acts as a normalisation constant for the response function h so that the product of the jump factor s_0 and the response function h can be interpreted as the number of offspring events and the density for the increase in activity following a mutually exciting incident [46]. In this two-dimensional system the second intensity function would model the influence of past BSF and PIRA IED attacks on BSF attacks.

To compute parameter estimates the MLE can again be employed. This MLE takes the following form [50]

$$\log L(\{t_a\}; \mu, k_0, \omega, s_0, \nu) = \sum_{a=1}^N \log(\lambda_1(t_a)) - \int_0^T \lambda_1(t) dt,$$

where, $T = \max\{\max_a\{t_a\}, \max_b\{t_b\}\}$. A similar formula holds for $\lambda_2(t)$ and $\{t_b\}$.

As in the one-dimensional case, all parameters must be positive [23] to make the model realistic and events in the infinite past should be considered [46]. Likewise, event times must be unique [30]. The condition necessary for the model to be non-explosive is reformulated in the higher dimensional case. In particular, consider the 2x2 matrix \mathbf{G} whose entries are formed of the integrals

$$\int_0^\infty g_{rs}(t) dt, \quad r, s \in \{1, 2\},$$

for each response function occurring in the definitions of λ_r . Then the condition for the model to be non-explosive is that the spectral radius defined as

$$\rho = \max_i \{|e_i|\} < 1,$$

where e_i represent the eigenvalues of \mathbf{G} [9]. Again this assumption has been checked and found to be satisfied for the models studied in this paper with the exception of Model 5 (see Table 3 below) which has a spectral radius of 1.0124. This case should therefore be treated with care and may be the result of a small dataset for this model (see the model analysis in Section 4.9 for further details).

Table 3. *Models*

Model Number	Event Times Interpreted	Dataset
0	Number of $\{t_i\}$ = IED events in NI	PIRA Events
1	$\{t_i\}$ = IED events in NI	PIRA Events
2	$\{t_i\}$ = IED events in Belfast	PIRA Events
3	$\{t_a\}$ = IED events in Belfast $\{t_b\}$ = IED events in the six counties of NI	PIRA Events PIRA Events
4	$\{t_a\}$ = IED events targeting BSFs in NI $\{t_b\}$ = BSF events which killed Catholic civilians	PIRA Events BSF Events
5	$\{t_a\}$ = IED events targeting BSFs in NI $\{t_b\}$ = BSF events which killed PIRA members	PIRA Events BSF Events

3.4 Table of Models

Having given the general form of the models in this paper the specific models that were studied are summarised in Table 3. In particular, this table presents an interpretation of the event times used in each model. For clarity the datasets being studied in each model are also listed corresponding to the information in Tables 1 and 2. Times t_i correspond to the univariate model presented in Section 3.2 whilst times t_a and t_b correspond to the multidimensional model from Section 3.3. In Section 4, where the numerical results of this paper are provided, the order of presentation will correspond to the ordering of models shown in Table 3.

3.5 Computational Methodology

Finding the parameters which maximise the log-likelihoods can be undertaken in numerous ways [42]. For this paper optimisation of the log-likelihood functions were undertaken in the Python programming language using the SciPy Optimize package Nelder-Mead [37, 49]. The Nelder-Mead algorithm was chosen based on previous observations of its effectiveness when applied to point process models [45] and also its performance during preliminary coding. It should be noted at this point that this optimisation procedure finds the minimum value, hence the equivalent problem of finding the minimising parameters of $-\log L$ was considered. To obtain further computational efficiency a recursive algorithm described by Liniger [30] was used to compute values of the intensity function.

Another important point, which is made by Egesdal et al. [18], is that due to the nonlinear nature of the minimisation it is not guaranteed that a global minimum will be found. Therefore, there is a need to begin the optimisation

procedure at multiple points and take the parameters yielding the lowest value of $-\log L$ and subject to the conditions given earlier.

Having derived each model's parameters we then go on to assess its goodness of fit. To determine the overall model fit residual analysis was employed. The basic ideas of this approach can be found in Brown et al. [10] and are also summarised below.

Consider a point process formed of the set of event times $\{t_i\}$ with intensity function λ . Perform the following integrals which transform the set $\{t_i\}$ to the set $\{\tau_i\}$

$$\tau_i = \int_0^{t_i} \lambda(t) dt.$$

If the model is a good fit then the residuals $\{\tau_i\}$ are independent and distributed according to a stationary Poisson process with unit rate [43]. Therefore, the inter-arrival times given by

$$Y_k = \tau_k - \tau_{k-1}$$

are exponentially distributed. Setting $\tau_0 = 0$ [10] and applying these procedures for the intensity function given in (3.2) the following formula is obtained

$$Y_1 = \mu t_1,$$

$$Y_k = \mu(t_k - t_{k-1}) - k_0 \sum_{i=1}^{k-1} e^{-\omega(t_k - t_i)} - e^{-\omega(t_{k-1} - t_i)}, \quad 1 < k \leq N.$$

If the inter-arrival times are exponentially distributed then

$$U_k = 1 - e^{-Y_k} \tag{3.5}$$

form a set of independent uniform random variables over $[0, 1)$. Therefore, to test the goodness of fit of the Hawkes process it remains to determine if the corresponding U_k do indeed come from a uniform distribution.

A quantitative test that can be used to check this assumption on the distribution of the U_k values is the Kolmogorov-Smirnov (KS) test [32]. The KS test in this case works by comparing the value of the test statistic $D_n = \max_k (|U_k - \frac{k-1}{N}|, |\frac{k}{N} - U_k|)$ [57] to a critical value D_α (see O'Connor and Kleyner [41] for a table of critical values). Statistical significance is obtained if the condition $D_n < D_\alpha$ is found to hold in which case there is evidence to suggest goodness of fit of the model.

Another method we used, which compares the fit of different models, is the Akaike Information Criterion (AIC) [1]:

$$\text{AIC} = 2k - 2 \log L,$$

where, k is the number of parameters being fitted in the model and $\log L$ is the maximum of the log-likelihood function. The model yielding the lowest value for AIC is deemed the better fit: more parameters are penalised whilst a greater value for the log-likelihood is rewarded. Burnham and Anderson [11] point out the AIC difference is not a significance test in the sense of critical values and

requires some judgement. For general guidance Burnham and Anderson [11] suggest that $0 - 2$ shows little difference between models, $4 - 7$ considerably more evidence for a difference and > 10 is classified as a significant difference. It is important that the AIC should only be used to compare models which are fitted with the same dataset [11], as was done for the comparisons in this paper.

4 Numerical Results

4.1 Determining Phase Boundaries

Initially we will be focused on PIRA IED attacks across all of the counties of NI and Belfast. Each phase of the events as outlined in Asal et al. [3] will be explored separately to offer a quantitative description of the changes in the Republican campaign.

The phases may not be treated in isolation, however. As discussed in Section 3.2, the historical dependence of the Hawkes process means that events outside of the observation period may influence those inside leading to spurious parameter values [46]. Thus, Phase 1 may influence events in Phase 2, Phase 2 may influence Phase 3 and so on. To avoid this edge effects issue, a systematic approach is required to deal with the phase boundaries.

We propose a novel approach to examine the impact of edge effects (see also Nichols and Schoenberg [38] for another similar approach in the field of seismology). Specifically, a moving time window was used to include the influence of events from Phase i in Phase $i + 1$ for $i = 1, 2, 3, 4$. This method is best illustrated via example. Consider Phases 1 and 2. First, we calculated the MLE parameters that result from the dataset consisting of Phase 2 only. Then adding one point from Phase 1 the MLE parameters were recalculated. This was then repeated with two points from Phase 1 and so on until the MLE parameters for all the data in Phases 1 and 2 combined had been calculated. Finally we determined which of these models provided the best fit compared to the others using the KS test. It was decided that a more positive difference $\Delta = D_\alpha - D_n$ gives more certainty that statistical significance has been reached. Thus the choice of the best fitting model, and hence the most sensible phase boundaries to use, was decided based on maximising Δ .

Since Phase 1 was effectively the start of the conflict, and the beginning of the PIRA as an organisation, we do not assume there will be any substantial effects from previous events outside the dataset. Therefore, the boundaries for Phase 1 were unchanged from their original definition as in Asal et al. [3]. In Table 4 the number of points that were required to maximise the difference between the critical value and KS test statistic are given for the remaining four phases. It can be seen that very few points were required to fix the correct mathematical boundaries for Phases 2 and 3 and it was found that with so few points there was little change in the parameter values compared to the unadjusted phases. However, in Phase 4 it was necessary to include 65% of data points from Phase 3 which significantly changed the parameter values from

Table 4. *Edge Effects Results*

Phase	Number of Data Points Added	New Boundary
1	-	27/01/1970
2	11	24/10/1976
3	11	22/08/1980
4	494	11/4/1984
5	41	22/5/1994

those found for the original Phase 4 boundary. This is illustrated clearly by the plots in Figure 1 which show the variation in the values of Δ and the three model parameters as each data point is added from Phase 3 to Phase 4. A similar observation was also made for the case of edge effects in Phase 5.

The boundaries found from this edge effects analyses will be used throughout the remainder of this paper. This ensures consistency and enables comparisons between models. From here on in reference to a model's MLE parameters means the parameters found using the boundaries stated in Table 4.

4.2 Comparing Models

4.3 Table of Results

For brevity, all model results are listed concurrently for each phase in Tables 5 to 7, but care should be taken when making a comparison of the different models. A comparison may be made within each of the three investigations: PIRA events across NI, PIRA events in Belfast and outside, PIRA and BSF events in NI, but not directly between them, except for model fit. See Section 3.5 for more detail.

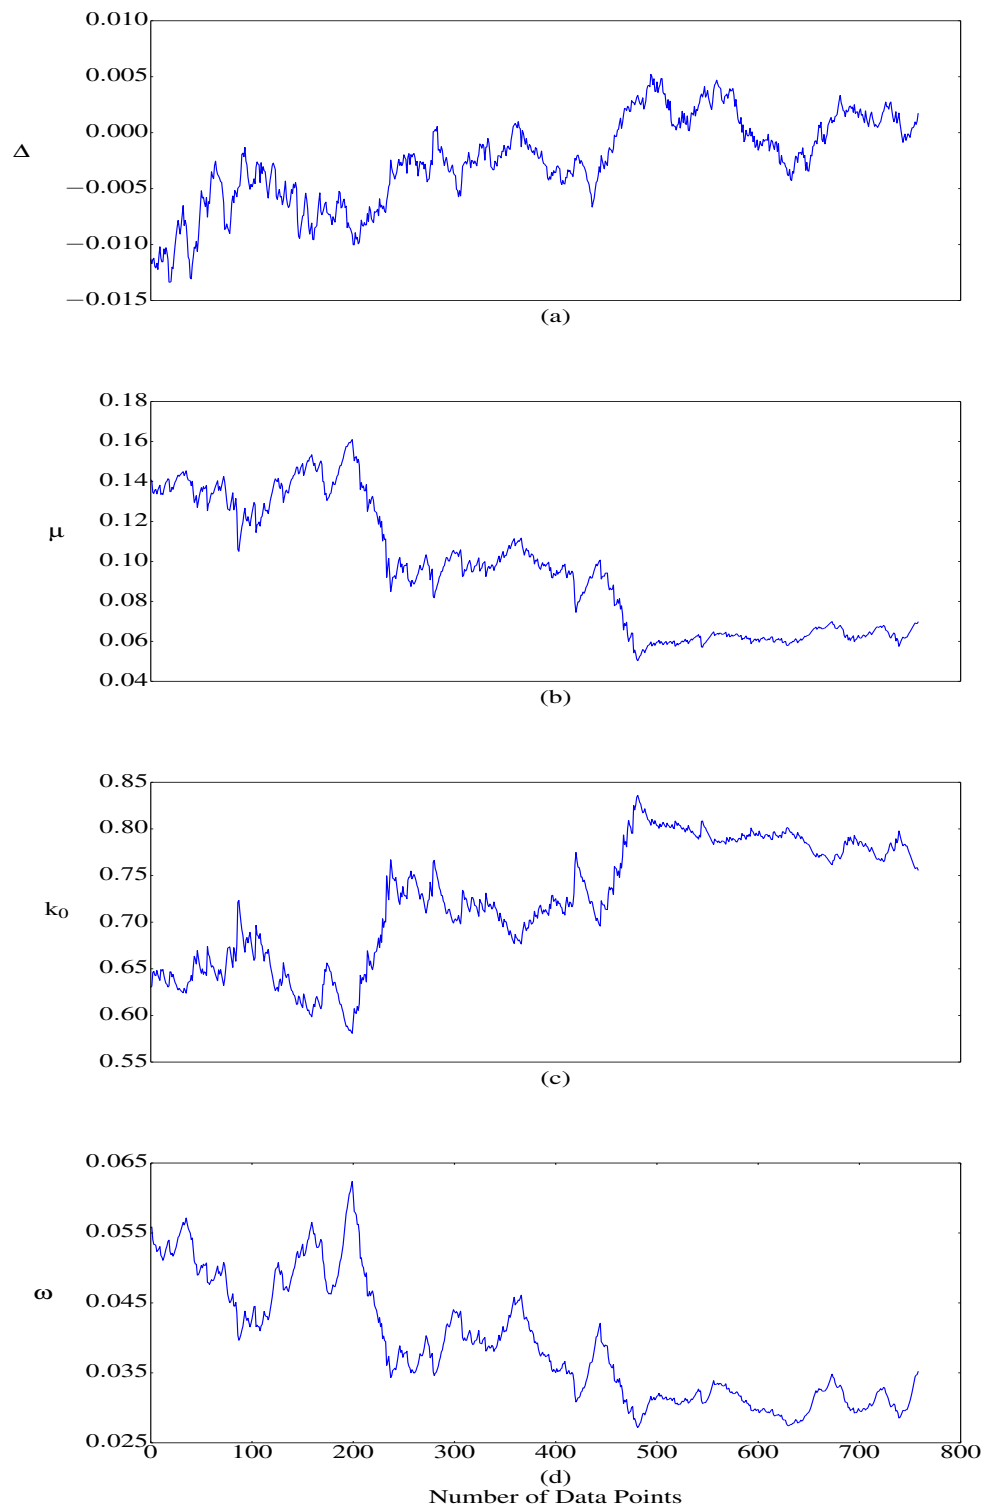


FIGURE 1. Figure (a) shows the values of $\Delta = D_\alpha - D_n$ for each point added from Phase 3 to Phase 4. Figures (b)-(d) show the corresponding changes in the MLE parameter values.

Table 5. *PIRA events across NI.*

		Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
Model 0	μ	0.3020	0.2541	0.2250	0.3028	0.0957
	AIC	3359.6	1784.4	3834.3	5190.5	1079.7
Model 1	μ	0.0543	0.1721	0.0800	0.0597	0.0270
	k_0	0.8241	0.3233	0.6529	0.8040	0.7231
	ω	0.0542	0.7685	0.0426	0.0316	0.0901
	KS Test	0.0686	0.0528*	0.0465*	0.0343*	0.0455*
	KS Critical 95%	0.0492	0.0701	0.0490	0.0396	0.1072
	KS Critical 99%	0.0590	-	-	-	-
	AIC	3083.7	1717.6	3750.5	5004	987

* Significant at 95% level.

Table 6. *PIRA events in Belfast and outside.*

		Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
Model 2	μ	0.0375	0.0397	0.0222	0.0175	0.0178
	k_0	0.7364	0.6456	0.7123	0.8766	0.5925
	ω	0.0246	0.0298	0.0135	0.0103	0.0874
	KS Test	0.0715*	0.0996*	0.0477*	0.0679**	0.0693*
	KS Critical 95%	0.0726	0.1096	0.0878	0.0640	0.1626
	KS Critical 99%	-	-	-	0.0767	-
	AIC	2015.1	967	1724.8	2707.8	556
	Model 3	μ	0.0396	0.0299	0.0189	0.0099
k_0		0.6441	0.5783	0.7126	0.7449	0.3842
ω		0.0272	0.0339	0.0135	0.0138	0.1001
s_0		0.0647	0.0949	0.0208	0.1085	0.2336
ν		0.7840	0.3976	1.1336	0.1099	0.1681
KS Test		0.0561*	0.0961*	0.0667*	0.0562*	0.1745**
KS Critical 95%		0.0726	0.1096	0.0878	0.0640	0.1626
KS Critical 99%		-	-	-	-	0.1948
AIC	2013.6	973	1726.1	2716.4	552.3	

* Significant at 95% level. ** Significant at 99% level.

4.4 Model 0

For this model we look at just a Poisson process applied to IED attacks across NI. Even this simple framework manages to capture a difference in rate across the five phases, showing Phase 4 as having the highest probability of a random event in a given time window with $\mu = 0.3028$. This is in contrast to the Hawkes process in Model 1 where Phase 2 is found to have the highest background rate whilst μ is much lower in Phase 4, suggesting that events in Phase 4 were heavily dependent on the past. Historical dependence is also seen for the other phases since the background parameter values found for Model 1 are all significantly lower than those of the Poisson process in Model 0. According to the AIC comparisons Model 1 is also shown to provide a better model fit in each phase.

Table 7. *PIRA and BSF events in NI.*

		Phase 1
Model 4	μ	0.0189
	k_0	0.1138
	ω	0.9716
	s_0	1.0694
	ν	0.0137
	KS Test	0.1122*
	KS Critical 95%	0.1133
	KS Critical 99%	-
	AIC	1057.8
Model 5	μ	0.0212
	k_0	0.2411
	ω	0.2757
	s_0	0.9774
	ν	0.0087
	KS Test	0.1205**
	KS Critical 95%	0.1133
	KS Critical 99%	0.1358
	AIC	1070.4

* Significant at 95% level. ** Significant at 99% level.

4.5 Model 1

The next model studied was a Hawkes process with a single self-exciting term applied to IED events across NI.

4.5.1 Goodness of Fit

Qualitative evidence for the model's ability to represent the data is gained from visualisation of the intensity function. In Figure 2 a plot of the intensity function over time for the Hawkes process in Phase 1 is presented. Here the peaks and troughs of the model are seen to follow closely the patterns of event times observed in the actual data.

The overall goodness of fit of Model 1 in each phase is determined via the KS test results. The critical values for the 95% confidence level have $D_\alpha = \frac{1.36}{\sqrt{N}}$ [41]. For Phase 1 it was found that the KS test statistic exceeded the critical value even when considering the 99% confidence level with $D_\alpha = \frac{1.63}{\sqrt{N}}$ [41]. Hence there is insufficient evidence in this case to conclude that the model is accurately capturing the dynamics of the data.

Although the KS test results for Phase 1 gave a negative result for the model fit there is some evidence for goodness of fit in a KS plot. Following the method outlined in Brown et al. [10] U_k , as defined in (3.5), is plotted against the hypothesized cumulative distribution, evaluated at $\frac{k-0.5}{N}$. If U_k is indeed uniformly distributed, the resulting graph should be a 45° line. The results of performing this procedure are shown in Figure 3 for Model 1 in each phase of the PIRA. Upper and lower bounds are also shown, obtained by plotting the lines $y = x \pm D_\alpha$ [13]. In cases where points do deviate from the best fit




FIGURE 2. Top graphs show event times of IED attacks and the intensity function for Model 1 in Phase 1. The bottom graphs are the same but only for the first 500 days.

line they still remain within or close to the error bounds, suggesting that the Hawkes process generally represents the data well.

A final test is to inspect a plot of U_{k+1} against U_k to ensure that the U_k are independent [4]. If serial correlation occurs it is most likely to be between adjacent time intervals, hence if these plots reveal some patterning it suggests that the transformed times are not independent. The graphs in Figure 4 present this analysis for each phase of the IED data. Although there is some patterning occurring in the plots for Phases 1 and 4 overall these plots appear to show little correlation between the neighbouring points of the sequence $\{U_k\}$. This serves as reassurance that independence exists and the Hawkes process is a good representation of IED attacks.

An AIC comparison to the simple Poisson process in Model 0 shows that Model 1 provides a better fit in each phase.

4.5.2 Interpreting Parameters

The parameter ω^{-1} gives information concerning the average length of time a series of attacks persists. In Phase 1 the average attack window is $\frac{1}{0.0542} = 18.5$ days. The rate of decay then has a large increase in Phase 2 yielding an av-




FIGURE 3. Figures (a) - (e) show KS plots for Hawkes process Model 1 in Phases 1-5 respectively. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

erage time window of 1.3 days. The work of Asal et al. [3] points out that in Phase 1 there was a more militaristic style of operation within the PIRA which could suggest better attack coordination allowing for longer periods of related waves of attacks. However, this organisational structure made the PIRA susceptible to infiltration by Security Forces thus prompting a shift to a cellular based approach in Phase 2 [3]. One of the consequences of infiltration could be reflected in the shorter attack window, which might be representing the fact

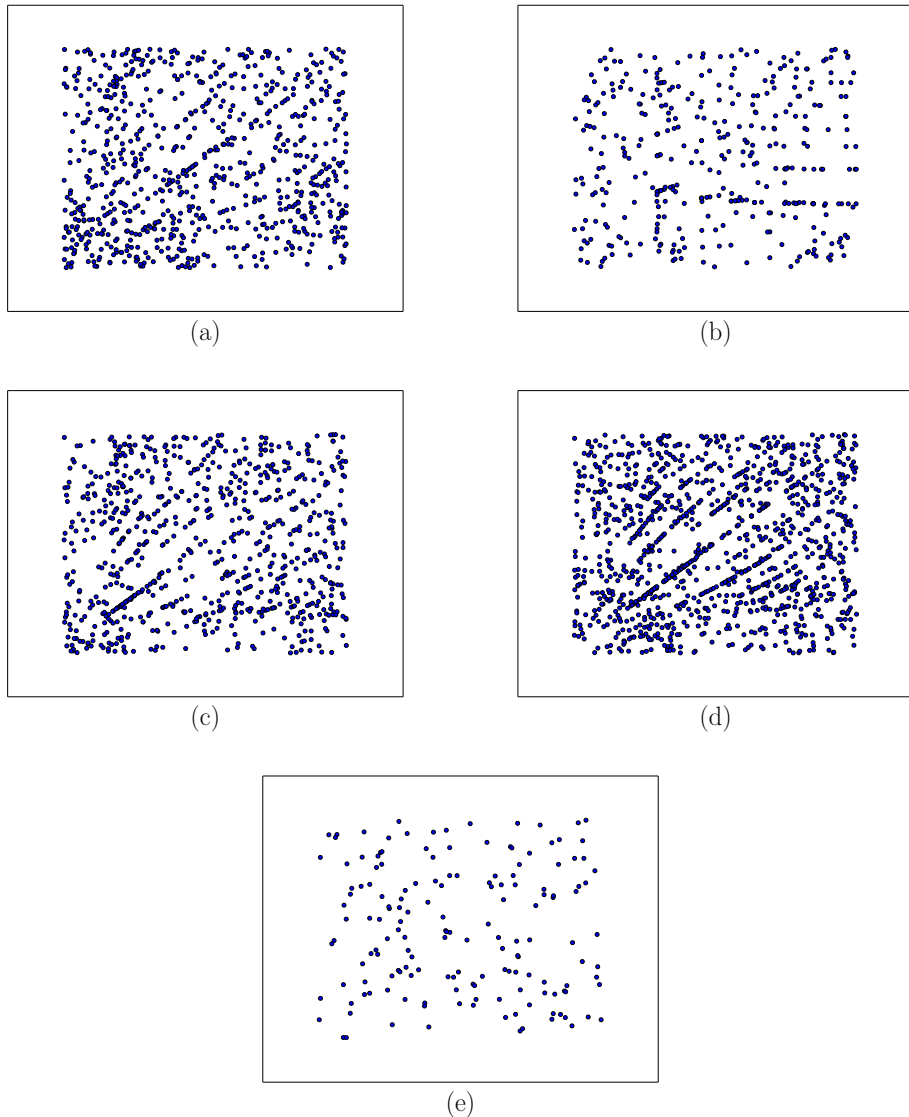


FIGURE 4. Figures (a) - (e) show serial correlation plots for Hawkes process Model 1 in Phases 1-5 respectively. A random dispersion of data points indicates goodness of fit of the model.

that many PIRA members were imprisoned [19], thus, there would have been fewer members to carry out attacks. In Phase 3 the attack window becomes 23.5 days. This phase saw a resurgence of violence by the PIRA with Moloney [36] describing the 1980's as a period of escalating violence similar to the "Tet Offensive" launched by the People's Army of Vietnam in 1968. In particular, the decay rate found could represent the PIRA using sustained attacks to weaken British resolve to remain in NI. This also links to the PIRA objectives described in the organisation's Green Book [39]. One of these objectives was to

use a “bombing campaign aimed at making the enemy’s financial interests in our country unprofitable”. Another objective was to wage a “war of attrition against enemy personnel which is aimed at causing as many casualties and deaths as possible so as to create a demand from their people at home for their withdrawal”.

Phases 4 and 5 then show an increasing trend for the decay rate. In Phase 4 the average time window for attacks is 31.6 days and in Phase 5 it is 11.1 days. These periods were characterised by secret meetings and negotiations that eventually led to the Good Friday Agreement [3]. The PIRA used IED attacks as a bargaining tool with the British Government [3] and as noted by Coogan [15] PIRA had the ability to “turn ... bombing[s] on and off like a tap”. So these shortening periods of IED usage may have been the PIRA using its capabilities as a way to achieve leverage with the British Government during peace negotiations rather than for a war of attrition.

The parameter μ can be interpreted as the background rate at which new events randomly occurred. The trend appears to be for the parameter to increase from Phase 1 to 2 and then fall in the remaining three phases. The rise in Phase 2 may be related to the fact that the organisation of PIRA was shifting and hence with less control members were conducting attacks more randomly. However, as the “Tet Offensive” campaign began in Phase 3 attacks became more systematic. Finally de-escalation of violence over Phases 4 and 5 explain the decreases in the value of μ .

The final parameter k_0 can be interpreted as the jump in the rate of events following an initial event. Phase 1 has the highest value for this parameter. This could again be a result of the military structure of the PIRA leading to more flexibility to escalate events. As before the drop in Phase 2 may be related to imprisonment of PIRA members hindering the extent to which attacks could occur. The rise in Phase 3 may also be interpreted as PIRA adopting a “Tet Offensive” approach and the Green Book objectives both explained above. The final phases have relatively high values of k_0 . This can be linked to the PIRA using IED attacks to demonstrate its capabilities during peace negotiations also as described above.

4.6 Model 2

The work of Fay, Morrissey and Smyth [20] demonstrates that violence during “The Troubles” was highly concentrated in Belfast. Thus in Model 2 we refine the geographical scale of investigation to this region considering a single term self-exciting Hawkes process applied to PIRA IED events.

4.6.1 Goodness of Fit

The results in Table 6 show that there is strong quantitative evidence for Model 2 providing a good fit to the data. The KS plot for the model in Phase 2 is shown in Figure 5(a). It was not felt that the data stayed close enough to the




FIGURE 5. Figures (a)-(b) show KS plots for Hawkes process Model 2 in Phases 2 and 4 respectively. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.




FIGURE 6. Serial correlation plot for Hawkes process Model 2 in Phase 4. A random dispersion of data points indicates goodness of fit of the model.

line of best fit to conclude a good fit of the Hawkes process. For the model in Phase 4 the KS test had to be conducted at the 99% level of confidence to obtain a significant outcome. Despite this the KS plot for Phase 4, as shown in Figure 5(b), gives some qualitative evidence for goodness of fit of the model even at the 95% level. There also appeared to be some patterning in the serial correlation plot for Model 2 in Phase 4 as shown in Figure 6.

4.6.2 Interpreting Parameters

Since Belfast was such a central stage in the NI conflict it is not too surprising that the trends for the parameter values in Model 2 are similar to those seen for Model 1. However, Model 2 in Phase 2 does not share the same significant changes in parameter values as observed for Model 1. This implies the PIRA attacks in Belfast were less susceptible to internal and external changes. One explanation for this observation is the existence of a Northern Command Unit being based in Belfast since 1969 [25]. As such it may have been easier for

PIRA to quickly adapt to internal and external events, such as, the mass imprisonment of PIRA members which led to the organisation becoming cell-based in Phase 2 [19].

4.7 Model 3

The next model considered was a multivariate Hawkes process. With this model we aimed to capture the influence on PIRA IED attacks in Belfast based on self-excitations of past PIRA attacks in Belfast and mutual-excitations of past PIRA attacks in the six counties of NI.

4.7.1 Goodness of Fit

It can be seen from the results table presented earlier that the increased complexity from using a multivariate Hawkes process in Model 3 does not yield a better fitting model compared to Model 2. This may be due to the autonomy of ASUs in the counties of NI and Belfast [25].

Nonetheless, in all but one case there is quantitative evidence for goodness of fit of the models. For the model in Phase 5 goodness of fit was found only after the KS test was conducted at the 99% level. The models in Phases 2, 3 and 5 did not have strong qualitative evidence for the goodness of fit of the model as shown by KS plots in Figure 7. Nonetheless, the transformed time data for each model appeared to be independent as measured by a serial correlation plot.

4.7.2 Interpreting Parameters

From the parameter values presented in Table 6 more evidence is gained for the relative autonomy of the Belfast Brigade from the other PIRA units. In particular, it can be seen that the value of the jump from self-excitations, k_0 , is much higher than that for mutual-excitations, s_0 . Whilst the opposite is true for the decay rate of self-excitations, ω , and those of mutual-excitations, ν . This suggests that the events in the six counties of NI had little impact on IED attacks in Belfast and the impact they did have was short lived. Also it can be seen that the self-excitation part of the model is very similar to that of the Belfast only case in Model 2 suggesting that Model 2 is sufficient for studying the internal dynamics of PIRA attacks in Belfast.

4.8 Model 4

In the final two models examined in this paper the focus will be on the influence of actions by BSFs in NI. These models are only examined in Phase 1 and on spatially aggregated data across NI due to a lack of data concerning BSF attacks. The first of these models is a multivariate Hawkes process considering the influence on PIRA attacks against BSFs based on self-excitations




FIGURE 7. Figures (a) - (c) show KS plots for Hawkes process Model 3 in Phases 2, 3 and 5 respectively. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

of past IED attacks against BSFs and mutual-excitations of past BSF attacks which killed Catholic civilians. The decision to study the impact of the deaths of Catholic civilians resulted from the fact that the PIRA were heavily rooted in the Catholic community [56]. In our datasets the number of IED attacks targeting BSFs is 144 and the number of Catholic civilian deaths resulting from BSF events is 78.

4.8.1 Goodness of Fit

Results of applying goodness of fit tests to Model 4 are shown in Table 7. Quantitatively Model 4 appears to provide a good fit to the IED data. Qualitatively, however, it should be noted that the data points in the KS plot do not lie on the line of best fit, as can be seen in Figure 8. Nonetheless, the serial correlation plot, although not included here, does suggest independence of the data points.




FIGURE 8. KS plot for multidimensional Hawkes process Model 4. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

4.8.2 Interpreting Parameters

Interpreting the MLE parameter values for Model 4 it can be seen that BSF attacks, which lead to the death of Catholic civilians, actually caused a backlash in terms of leading to an increase in IED attacks. In particular, the jump parameter, s_0 , is high compared to the other values seen in this paper implying that following an incident involving the death of a Catholic civilian the PIRA were likely to respond with a large increase in IED attacks targeting BSFs. Also the decay rate of this increase in attacks, ν , is small suggesting a lengthy period of increased violence equivalent to an average of 73 days. This prolonged retaliation by PIRA may be seen as an attempt by the organisation to obtain public legitimacy by acting as defenders of the Catholic population. Such results also give support for the sort of findings made by Braithwaite and Johnson [7] where less discriminatory counter-insurgency operations were found to result in an increase in violence.

4.9 Model 5

The final model studied in this paper is similar to Model 4 except now the influence of BSF attacks which killed PIRA members is considered. In our dataset there are 58 recorded incidents of BSF events resulting in PIRA member deaths.

4.9.1 Goodness of Fit

Goodness of fit test results are shown in Table 7. These results show that, based on the AIC, Model 5 performs worse than Model 4 at modelling PIRA attacks targeting BSFs. This result indicates that Catholic civilian deaths were better predictors of a backlash by the PIRA. Moreover, the KS test is only significant for Model 5 when considered at the 99% level. The KS plot for Model 5, shown in Figure 9, suggests that the data points are not falling on the line of best fit so there is not enough qualitative evidence to declare a




FIGURE 9. KS plot for multidimensional Hawkes process Model 5. Data points falling on the solid goodness of fit line imply a perfect model fit with the dashed lines representing 95% error bounds.

significantly good fit. However, the serial correlation plot did suggest that the data points were independent.

4.9.2 Interpreting Parameters

The parameter values for Model 5 also suggest an increase in violence against BSFs following a BSF event leading to the death of a PIRA member. The increase in the rate of IED attacks, given by s_0 , following such an operation is high compared to other values seen in this paper and the length of time this increase is sustained, given by ν , is quite long at an average of 115 days. These observations may be the result of PIRA trying to revenge the death of its members as well as demonstrating defiance. This is similar to the findings of Lafree, Dugan and Korte [28] which suggested an increase in PIRA attacks following BSF operations which killed PIRA members. When planning CT strategies these observations, and similar ones for Model 4, imply that after a civilian or terrorist death quick interventions are necessary to avoid long time periods of increased violence.

5 Discussion

We began this paper by outlining the importance of mathematical modelling in terrorism and stressing in particular the contributions such modelling could make to improving CT responses. From a firm foundation of criminological theory concerning spatio-temporal patterns of crime and a mathematical model known as a Hawkes process we then proceeded to model IED attacks during the NI conflict.

The first step taken in this paper was to divide the IED dataset we used into five phases corresponding to organisational changes within the PIRA as described by Asal et al. [3]. Although this allowed for a greater depth of analysis it did introduce the issue of edge effects where events in one phase influence those in the next phase. To account for this a novel approach was found using a

moving time frame to incorporate events from the previous phase where an improvement was then judged using the KS test. It was found that in Phases 2 and 3 very few data points were needed to achieve the best fitting model. However in Phases 4 and 5 more significant edge effects were found. For each phase new mathematical boundaries were fixed for the analyses in this paper. Such findings raise interesting questions relating to the timing of tactical and organisational shifts within the PIRA that may have previously been missed.

Having defined the phase boundaries of PIRA we then moved to analyse six models aimed at capturing different temporal patterns of IED usage by the organisation at different geographical scales.

Initially in Model 1 we examined the self-exciting nature of IED attacks across the whole of NI. For this case it was found that the model outperforms a simple Poisson process, defined by Model 0, as measured by the AIC. Moreover, quantitative and qualitative evidence suggested that this model was capturing the temporal patterns of IED attacks. With the model specified we were then able to compare its predictions to historical accounts of the PIRA illustrating how such models could be used in practise to determine how terrorist groups respond to past events.

Being the center of much of the violence seen during “The Troubles” [20] we then decided to refine the geographical scale to focus on Belfast. Here it was found that there was quantitative and qualitative evidence to suggest the model is good at capturing past influences on further IED attacks. In addition, a study of the model parameters revealed a similar pattern to Model 1 although with some difference in Phase 2. But again using historical accounts of the PIRA we were able to account for this difference. This shows that by adjusting the spatial scale it is possible to gain more refined information about a terrorist organisation demonstrating the depth of insights that can be gained from Hawkes process models.

Having studied univariate models the next model considered was a multidimensional model which aimed to examine the interplay between IED attacks in Belfast and those in the six counties of NI. It was found that the additional complexity did not yield significant improvements over the self-exciting model in Belfast only. However, it was found that this could be accounted for by examining the autonomous nature of Active Service Units of the PIRA in the counties of NI and Belfast. Hence this gives an example of how Hawkes processes can also be used to uncover simultaneous influences on different components of terrorist organisations over a range of spatial scales. Quantitative and some qualitative evidence also exists confirming these models are capturing some of the IED dynamics studied.

For the final two models our focus was on determining the effectiveness of multidimensional Hawkes processes for testing CT strategies. The first of these models was used to investigate how BSF attacks which lead to the death of Catholic civilians influenced PIRA attacks. On the other hand the second model examined the influence of BSF attacks leading to PIRA member deaths. The former model proved to be a better fit for the IED data both in terms of a

direct comparison using the AIC and using quantitative and qualitative goodness of fit tests. Although with small datasets for this investigation some care should be taken when drawing conclusions from these models. Despite this the parameter values for both models were consistent with previous research showing the retaliatory nature of terrorist groups. This indicates that important lessons can be learned from Hawkes processes concerning how terrorist groups will respond to different events.

In conclusion, it is hoped that this paper has shown the adaptability of Hawkes process models to study a range of areas within CT. For future research it is planned to extend these models further using an explicit spatial component as opposed to the implicit approach taken here. This should allow further details concerning the patterning of PIRAs attacks to be uncovered and also enable more depth to be gained concerning hotspots of terrorism during the NI conflict. It would also be interesting to take a sociological point of view on the results found here and in particular examine again the changes that occurred in the PIRA to align theory with the mathematical phase boundaries that we have found.

Acknowledgement

This project has been funded by the EPSRC through the UCL Security Science Doctoral Training Centre, reference EP/G037264/1.

References

- [1] Akaike, H. (1974) A new look at the statistical model identification. *IEEE T. Automat. Contr.*, **19**(6), 716–723.
- [2] Andresen, M. A. and Malleson, N. (2011) Testing the Stability of Crime Patterns: Implications for Theory and Policy. *J. Res. Crime Delinq.*, **48**(1), 58–82.
- [3] Asal, V., Gill, P., Rethemeyer, R. K. and Horgan, J. (2013) Killing Range: Explaining Lethality Variance within a Terrorist Organization. *J. Conflict Resolut.* **59**(3), 401–427.
- [4] Berman, M. (1983) Comment on “Likelihood Analysis of Point Processes and Its Applications to Seismological Data”, by Y. Ogata. In: *Bulletin of the International Statistical Institute*, **50** (3), 412–418.
- [5] Bowsher, C. G. (2007) Modelling security market events in continuous time: Intensity based, multivariate point process models. *J. Econometrics*, **141**(2), 876–912.
- [6] Braga, A. A. (2001) The Effects of Hot Spots Policing on Crime. *Ann. Am. Acad. Polit. Ss.*, **578**(1), 104–125.
- [7] Braithwaite, A. and Johnson, S. D. (2012) Space–Time Modeling of Insurgency and Counterinsurgency in Iraq. *J. Quant. Criminol.*, **28**(1), 31–48.
- [8] Brantingham, P. J. and Brantingham, P. L. (1984) *Patterns in Crime*. New York: Macmillan.

- [9] Brémaud, P. and Massoulié, L. (1996) Stability of Nonlinear Hawkes Processes. *Ann. Probab.* **24**(3), 1563–1588.
- [10] Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E. and Frank, L. M. (2002) The Time-Rescaling Theorem and Its Application to Neural Spike Train Data Analysis. *Neural Comput.*, **14**(2), 325–346.
- [11] Burnham, K. P. and Anderson, D. R. (2002) *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. 2nd ed. New York: Springer Science & Business Media.
- [12] Campbell, C. and Connolly, I. (2003) A Model for the “War Against Terrorism”? Military Intervention in Northern Ireland and the 1970 Falls Curfew. *J. Law Soc.*, **30**(3), 341–375.
- [13] Chen, Z., Purdon, P., Brown, E. N. and Barbieri, R. (2012) A unified point process probabilistic framework to assess heartbeat dynamics and autonomic cardiovascular control. *Front. Physio.*, **3**.
- [14] Clauset, A. and Woodard, R. (2013) Estimating the historical and future probabilities of large terrorist events. *J. Appl. Stat.*, **7**(4), 1838–1865.
- [15] Coogan, T. P. (2002) *The IRA*. New York: Palgrave.
- [16] Eck, J. E., Chainey, S., Cameron, J. G., Leitner, M. and Wilson, R. E. (2005) *Mapping Crime: Understanding Hot Spots*. Washington: U.S. Department of Justice.
- [17] Eck, J. E. and Weisburd, D. (1995) *Crime Prevention Series, Volume 4: Crime and Place*. New York: Criminal Justice Press.
- [18] Egesdal, M., Fathauer, C., Louie, K. and Neuman, J. (2010) Statistical and Stochastic Modeling of Gang Rivalries in Los Angeles. *SIURO*, **3**, 72–94.
- [19] English, R. (2004) *Armed Struggle: The History of the IRA*. London: Pan Macmillan.
- [20] Fay, M.-T., Morrissey, M. and Smyth, M. (1999) *Northern Ireland’s Troubles: The Human Costs*. London: Pluto Press.
- [21] Fox, E. W., Short, M. B., Schoenberg, F. P., Coronges, K. D. and Bertozzi, A. L. (2014) Modeling E-mail Networks and Inferring Leadership Using Self-Exciting Point Processes. [online] Available at: <<http://people.math.gatech.edu/~mshort9/papers/IkeNet.pdf>> [Accessed 22 May 2015].
- [22] Gao, P., Guo, D., Liao, K., Webb, J. J. and Cutter, S. L. (2013) Early Detection of Terrorism Outbreaks Using Prospective Space–Time Scan Statistics. *Prof. Geogr.*, **65**(4), 676–691.
- [23] Hawkes, A. G. (1971) Spectra of Some Self-Exciting and Mutually Exciting Point Processes. *Biometrika*, **58**(1), 83–90.
- [24] Hegemann, R. A., Lewis, E. A. and Bertozzi, A. L. (2013) An “Estimate & Score Algorithm” for simultaneous parameter estimation and reconstruction of incomplete data on social networks. *Security Informatics*, **2**(1).
- [25] Horgan, J. and Taylor, M. (1997) The Provisional Irish Republican Army: Command and Functional Structure. *Terror Polit. Violenc.*, **9**(3), 1–32.
- [26] Johnson, N. F., Medina, P., Zhao, G., Messinger, D. S., Horgan, J., Gill, P., Bohorquez, J. C., Mattson, W., Gangi, D., Qi, H., Manrique, P., Velasquez, N., Morgenstern, A., Restrepo, E., Johnson, N., Spagat, M. and Zarama, R.

- (2013) Simple mathematical law benchmarks human confrontations. *Sci. Rep.*, **3**.
- [27] Johnson, S. D. and Bowers, K. J. (2004) The Burglary as Clue to the Future The Beginnings of Prospective Hot-Spotting. *Eur. J. Criminol.*, **1**(2), 237–255.
- [28] Lafree, G., Dugan, L. and Korte, R. (2009) The Impact of British Counterterrorist Strategies on Political Violence in Northern Ireland: Comparing Deterrence and Backlash Models. *Criminology*, **47**(1), 17–45.
- [29] Lewis, E., Mohler, G., Brantingham, P. J. and Bertozzi, A. L. (2012) Self-exciting point process models of civilian deaths in Iraq. *Security Journal*, **25**(3), 244–264.
- [30] Liniger, T. (2009) “Multivariate Hawkes Processes”. PhD thesis. Zurich: Swiss Federal Institute of Technology.
- [31] Lum, C. and Kennedy, L. W. (2012) *Evidence-Based Counterterrorism Policy*. New York: Springer Science + Business Media.
- [32] Massey Jr., F. J. (1951) The Kolmogorov-Smirnov Test for Goodness of Fit. *J. Am. Stat. Assoc.*, **46**(253), 68–78.
- [33] Memon, N., Farley, J. D., Hicks, D. L. and Rosenorn, T. (2009) *Mathematical Methods in Counterterrorism*. New York: Springer Science + Business Media.
- [34] Mohler, G. O., Short, M. B., Brantingham, P. J., Schoenberg, F. P. and Tita, G. E. (2011) Self-Exciting Point Process Modeling of Crime. *J. Am. Stat. Assoc.*, **106**(493), 100–108.
- [35] Mohler, G. (2013) Modeling and estimation of multi-source clustering in crime and security data. *Ann. Appl. Stat.* **7**(3), 1525–1539.
- [36] Moloney, E. (2003) *A Secret History of the IRA*. New York: W. W. Norton.
- [37] Nelder, J. A. and Mead, R. (1965) A Simplex Method for Function Minimization. *Comput. J.*, **7**(4), 308–313.
- [38] Nichols, K. and Schoenberg, F. P. (2014) Assessing the dependency between the magnitudes of earthquakes and the magnitudes of their aftershocks. *Environmetrics*, **25**(3), 143–151.
- [39] O’Brien, B. (1999) *The Long War: The IRA and Sinn Féin*. 2nd ed. New York: Syracuse University Press.
- [40] O’Connor, M. P. and Rumann, C. (2003) Into the Fire: How to Avoid Getting Burned by the Same Mistakes Made Fighting Terrorism in Northern Ireland. *Cardozo Law Review*, **24**(4).
- [41] O’Connor, P. D. T. and Kleyner, A. (2012) *Practical Reliability Engineering*. 5th ed. West Sussex: Wiley.
- [42] Ozaki, T. (1979) Maximum likelihood estimation of Hawkes’ self-exciting point processes. *Ann. I. Stat. Math.*, **31**(1), 145–155.
- [43] Papangelou, F. (1972) Integrability of Expected Increments of Point Processes and a Related Random Change of Scale. *T. Am. Math. Soc.*, **165**, 483–506.
- [44] Pease, K. (1998) *Repeat Victimisation: Taking Stock*. London: Home Office Police Research Group.
- [45] Peng, R. (2003) Multi-dimensional Point Process Models in R. *J. Stat. Softw.*, **8**(16).

- [46] Rasmussen, J. G. (2013) Bayesian Inference for Hawkes Processes. *Methodol. Comput. Appl.*, **15**(3), 623–642.
- [47] Ross, S. (2010) *A First Course in Probability*. 8th ed. London: Pearson.
- [48] Rubin, I. (1972) Regular Point Processes and Their Detection. *IEEE T. Inform. Theory*, **IT-18**(5), 547–557.
- [49] Scipy Optimize Minimize (2014) [online] Available at: <<http://docs.scipy.org/doc/scipy-0.14.0/reference/generated/scipy.optimize.minimize.html#scipy.optimize.minimize>> [Accessed 17 Dec. 2014].
- [50] Short, M. B., Mohler, G. O., Brantingham, P. J. and Tita, G. E. (2014) Gang rivalry dynamics via coupled point process networks. *Discrete Cont. Dyn.-B*, **19**(5), 1459–1477.
- [51] Stevenson, R. and Crossley, N. (2014) Change in Covert Social Movement Networks: The “Inner Circle” of the Provisional Irish Republican Army. *Social Movement Studies*, **13**(1), 70–91.
- [52] Stomakhin, A., Short, M. B. and Bertozzi, A. L. (2011) Reconstruction of missing data in social networks based on temporal patterns of interactions. *Inverse Probl.* **27**(11).
- [53] Sutton, M. (1994) *An Index of Deaths from the Conflict in Ireland 1969-1993*. Belfast: Beyond the Pale.
- [54] Varadhan, S. R. S. (2007) *Stochastic Processes (Courant Lecture Notes)*. Rhode Island: American Mathematical Soc.
- [55] White, G., Porter, M. D. and Mazerolle, L. (2012) Terrorism Risk, Resilience and Volatility: A Comparison of Terrorism Patterns in Three South-east Asian Countries. *J. Quant. Criminol.*, **29**(2), 295–320.
- [56] White, R. W. (1993) *Provisional Irish Republicans: An Oral and Interpretive History*. Westport, CT: Greenwood Press.
- [57] Zar, J. H. (2014) *Biostatistical Analysis*. 5th ed. Essex: Pearson Education Limited.