ORIGINAL ARTICLE
Journal Section

(No) Spillovers in reporting domestic abuse to police

Lara Vomfell[1] | Jan Povala[2]

[1]Warwick Business School, University of Warwick, Coventry, UK
[2]Department of Mathematics, Imperial College London, London, UK
[2]The Alan Turing Institute, London, UK

Correspondence
Warwick Business School, Scarman Road, Coventry, CV4 7AL
Email: L.Vomfell@warwick.ac.uk

Funding information
JP was supported by EPSRC (EP/L015129/1)

Spillover effects in crime are typically studied as a result of offender behaviour. This study investigates whether spillover effects can occur in the reporting of domestic abuse by victims. Domestic abuse is a particularly interesting context because of its high prevalence but low reporting rate. Extending existing spatio-temporal Hawkes process specifications, we test for the presence of two spillover channels in all domestic abuse calls in a major English city. We find no evidence to support such effects in the reporting of domestic abuse.

KEYWORDS
spatio-temporal point process, crime, domestic abuse, self-exciting process, spillover effect, non-parametric estimation

1 | INTRODUCTION

Crimes like burglaries, homicides and robberies exhibit triggering behaviour: immediately after their occurrence, the risk of another crime in the same neighbourhood is increased (Johnson and Bowers, 2004; Mohler et al., 2011; Reinhart and Greenhouse, 2018; Flaxman et al., 2019). The triggering of these crime types is typically explained with the same offender repeating their crime or similar offenders learning about the promising criminal opportunity (Bernasco, 2008).

Victim behaviour, rather than offender behaviour, is yet to be examined as a source of triggering. Crime victims share their experiences with the crime itself and with reporting it to the police. This could influence other crime victims in their reporting decisions and imply triggering of crime reporting.

In this paper, we investigate if the reporting of domestic abuse to police exhibits triggering behaviour. Domestic abuse is a particularly interesting context to study spillovers in victim behaviours due to four facts:

1. It is highly prevalent. In the United Kingdom, approximately one in four adults experience domestic abuse in their adult life and an estimated 2.4 million adults have experienced some form of it in 2018 (Office for National Statistics, 2019a).

2. It is significantly under-reported. Domestic abuse is largely hidden from public view. It is one of the most under-reported serious crimes and victims/survivors often endure abuse for many years before seeking outside institutional support with police or a support provider (SafeLives, 2015). Reporting domestic abuse to police is a complex decision, as we detail in Section 2.3.

3. It has a high degree of social disclosure, that is many victims/survivors tell a friend, neighbour or family member about the abuse (Osborne et al., 2012).

4. Typically, perpetrator and victim(s) constellations remain stable over time such that one perpetrator acts within only one household. That means that if we were to document any spillovers in reports of domestic abuse, they must be due to spillovers in victim behaviour rather than in offender behaviour.

To clarify terminology, we use the terms 'victim' and 'survivor' interchangeably. While 'victim' is commonly used in the criminal justice context to describe someone who experienced domestic abuse, it has also been criticised for implying a passive role with 'survivor' as the proposed alternative. We use both terms to reflect the tension between the victimisation that people experiencing abuse go through and the agency they possess to define the experience on their own terms.

High rates of social disclosure paired with the high prevalence of domestic abuse provide a context in which the disclosure of domestic abuse to police by *others* could affect an individual's decision to report. However, identifying spillovers in domestic abuse reporting is challenging because the number of calls to police about domestic abuse varies, both due to the variation in the underlying abuse and due to variation in reporting behaviour (Cohn, 1993). Furthermore, since neighbourhood characteristics have been shown to influence both the incidence of domestic abuse and the reporting of crime, we would expect both to cluster in neighbourhoods (Goudriaan et al., 2006; Beyer et al., 2015).

To separate spillover effects from spatio-temporal clustering, we model a data set of domestic abuse reports to police using a Hawkes-type triggering spatio-temporal point process, called Hawkes process for short (Hawkes, 1971; Ogata, 1988; Reinhart, 2018). Hawkes processes have been an invaluable tool in identifying spillovers in criminal behaviour and disentangling them from clusters in space and time (Mohler, 2014; Reinhart and Greenhouse, 2018; Flaxman et al., 2019; Park et al., 2021). The Hawkes process predicts a series of discrete events (crime reports) that arise from one of two components: the background component which corresponds to the "typical" incidence of domestic abuse reports or the triggering component where past reports can trigger the occurrence of future reports. Hawkes processes do not identify a causal link between events, as in "event i caused event j". Instead, past events increase the likelihood of future events in their spatio-temporal vicinity. Triggering is therefore the statistical quantification of said increase or spillover, not the identification of a causal relationship.

Specifying the forms of the background and triggering components is a crucial modelling choice since their shapes have important implications for inference and detection of triggering behaviour (Reinhart and Greenhouse, 2018; Park et al., 2021). We employ the specification of the Hawkes process proposed by Zhuang and Mateu (2019) because it includes periodic components that account for daily and weekly periodicity of reporting behaviour. Periodicity is predicted by criminological theories such as routine activity as well as empirical data (Cohen and Felson, 1979; Rotton and Cohn, 2001; Johnson and Bowers, 2004).

We also develop an extension of Zhuang and Mateu (2019)'s model such that reports of domestic abuse can be triggered by past reports, but also by the police response to past events. A meta-study by Davis et al. (2008) analysed the effect of police returning to the incident household to follow-up with the victim(s). They find that such follow-up visits do not reduce future violence occurring but can increase victim reporting of the violence. It stands to reason that if follow-up visits by police amplify the effects of a report to police, this effect could potentially extend beyond the reporting household. Our model extension adding an additional spillover channel tests this explicitly. With this extension, we can statistically distinguish two channels of spillover effects: 1. report-to-report and 2. follow-up-to-report spillovers.

We test the presence of triggering effects in domestic abuse reporting on a data set of 11,691 calls for service to police. The data set contains all calls concerning domestic abuse in a major English city between January 2017 and December 2018. We use the 2017 data to select a model for the 2018 data. We find extremely limited evidence of spillover effects of domestic abuse reporting. This is true for both types of spillovers tested in this study, report-to-report and follow-up-to-report.

2 | MOTIVATION

The focus of our study is on potential spillover effects of domestic abuse reporting. Section 2.1 gives an overview of the literature on clustering and triggering of crimes. Section 2.2 introduces the Hawkes process as an explicit test to separate clustering from spillover effects. Section 2.3 introduces domestic abuse as a particularly interesting crime type for potentially triggering behaviour, and in Section 2.4 we discuss potential channels of such behaviours.

2.1 | Concentration of crime and spillover effects

Crime events concentrate in space and time, leading Weisburd (2015) to argue for a ‘law of crime concentration’: Just a few street segments account for the vast majority of crimes in a city. However, it is important to distinguish whether the increased level of criminal activity is due to exogenous environmental factors (clustering), or due to past crime events themselves (triggering). Different criminological theories have been put forward to understand the factors contributing to these spatio-temporal patterns (Johnson and Bowers, 2004; Gorr and Lee, 2015).

Multiple theories consider the effects of time and spatial locations, giving rise to clustering. *Routine activity theory* analyses crimes as the intersection of a suitable target, a motivated offender and lack of supervision (Cohen and Felson, 1979). *Rational choice frameworks* suggest that rational offenders choose location and time with a high number of potential targets, while considering potential costs and payoffs of a criminal opportunity (Clarke and Cornish, 1985; Sanders et al., 2017). *Social disorganisation theory* proposes that areas with low levels of social cohesion – indicated by high social and economic deprivation levels, family disruption and ethnic heterogeneity – tend to have higher density of crime perpetrators (Shaw and McKay, 1942; Sampson and Groves, 1989). Focusing on the spatial dimension, *crime pattern theory* posits that spatial topologies such as street layouts play an important role in the offender’s decision-making process (Brantingham and Brantingham, 1993a,b).

A second strand of literature focuses on triggering, conceptualised as repeat victimisation or near-repeat victimisation, where certain properties, places and people are more likely to become victimised multiple times (Farrell and Pease, 1993). The near-repeat victims are targets that are situated in close spatio-temporal proximity to an original target. The reasoning is that a successful offender will seek to repeat their success in addition to other offenders picking up similar cues (Bernasco, 2008). For example, a house that has been burgled may be burgled again to target replaced goods, or the house may be re-victimised because it is now familiar to the offender (Short et al., 2009; Bernasco et al., 2015). Retaliation can also lead to new crimes following a single incident, as has been shown for gang violence and shootings (Loeffler and Flaxman, 2018; Ratcliffe and Rengert, 2008; Brantingham et al., 2018).

The phenomena mentioned in the previous paragraph imply a dependence between crime events, called triggering or spillover. This dependency is the result of the offender's decision-making process of target selection. A phenomenon that has not yet been explored in depth is whether there are any spillovers in victim behaviour (other than retaliatory action). The decision to report a crime to police is an important one. One could therefore ask whether this decision has a spillover effect on the victim's surroundings, such as neighbours or their social network.

Spillovers resulting from a person's or household's behaviour to the surroundings has been shown to take place in various contexts and provides us with a reference frame on how crime victim behaviour might spill over. For example, sending letters about TV licences to households increases compliance in non-treated households in the neighbourhood (Rincke and Traxler, 2011; Drago et al., 2020). Beyond licence compliance, such effects have been documented in, e.g., voting, insurance, or school performance (Nickerson, 2008; Hong and Raudenbush, 2006; Sobel, 2006; Cai et al., 2015; Halloran and Hudgens, 2016).

2.2 | Hawkes inference as a test of spillover effects in observational studies

Identifying spillover effects amounts to answering the following question: 'Is the increased manifestation of a specific behaviour the result of spillovers or simply concentration due to exogenous factors such as the environment?' This task is relatively feasible when evaluating a specific, randomised intervention such as sending TV licence reminders (Aronow et al., 2021), but becomes much harder in observational studies. For example, Bertrand et al. (2000) demonstrate that women are more likely to use welfare when there is a local network of the same ethnic group (and the group also has a high level of welfare use). Aizer and Currie (2004) try to separate neighbourhood from network effects in the similar use of publicly funded prenatal care within ethnic groups. In contrast, they find no evidence for information sharing through networks and instead show that behaviour is highly similar in ethnic groups because of local hospital policies. Without assumptions or explicit knowledge about the underlying structure of networks (e.g., Fadlon and Nielsen, 2019; Nicoletti et al., 2018), it is challenging to separate concentration clusters and spillovers. It is in this precise context that the Hawkes process becomes a particularly valuable statistical model. Its initial applications focused on the triggering component of the process, e.g., earthquakes (Ogata, 1988). Later interpretations of Hawkes processes present them as a principled, model-based way of separating space-time clusters from the spillover of a behaviour (Meyer et al., 2016). Hawkes processes model discrete events happening in a spatio-temporal domain, such as occurrences of crime. The process is characterised by its conditional intensity – the intensity or rate of events at time t, conditioned on the events that occurred before time t. The intensity has two components: the *background* component representing the variation in time and space due to exogenous factors, and the *triggering* component which accounts for the events that come about as a result of past events (i.e., that are triggered). Inferring parameters of the model enables us to disentangle the triggering behaviour from the background variation such as geographical differences and temporal periodicity.

As we will argue in the next section, domestic abuse is an ideal context in which to explore the potential for reporting spillovers: It is under-reported but highly prevalent and exhibits a high degree of social disclosure while typically only having one offender/victim constellation per household. In particular, the latter property implies that if there is a spillover effect, it is not due to one offender choosing multiple targets. The use of a Hawkes process allows us to identify if a single report of domestic abuse to police increases the likelihood that a victim of domestic abuse in another household in the neighbourhood will also report the crime, without having to make any assumptions about the nature or structure of local networks.

2.3 | Reporting Domestic Abuse

A persistent challenge to understanding domestic abuse is its hidden nature: the abuse is hidden away from the public eye in private homes, but domestic abuse is also extremely under-reported. Most victims endure domestic abuse for a long period before disclosing their abuse to a formal institution, if at all (SafeLives, 2015). As a result, police-reported numbers are significant undercounts of the actual prevalence. Even numbers from large-scale surveys of the general population such as the Crime Survey for England and Wales are incomplete because they are highly sensitive to methodological details and exclude people outside of stable households (Ellsberg et al., 2001; Walby and Allen, 2004; Emery, 2010; Agüero and Frisancho, 2021; Office for National Statistics, 2017).

It is crucial to understand why victims/survivors of domestic abuse do not report their abuse. For the purposes of our study, the following section will focus specifically on the decision to report to the police. The first crucial step is recognising the abuse as a serious offence worth reporting, rather than a "private issue" (Bachman and Coker, 1995; Tjaden and Thoennes, 2000; Rogers et al., 2016). Perceiving a situation as "legitimate abuse" depends on a variety of factors of the situation: the seriousness of the physical abuse, whether alcohol was involved, whether the victim belongs to a marginalised group, and others (Trute et al., 1992; Robinson et al., 2018; O'Neal, 2019). However, these factors not only affect victims of domestic abuse, but also police officers, with important implications for how seriously police take individual cases (Miller and Segal, 2019; Kavanaugh et al., 2019). This arbitration of legitimacy can turn police officers—who are often the first point of contact to the criminal justice system—into gatekeepers of access to said systems (Taylor and Gassner, 2010).

Surveys have shown that victims fear not being believed or taken seriously by police (Tjaden and Thoennes, 2000; Hawkins and Laxton, 2014; Her Majesty's Inspectorate of Constabulary, 2014). Similarly, victims/survivors who decide to report to police report mixed experiences. In the Crime Survey for England and Wales, 72% of victims of domestic abuse stated that they found the police fairly or very helpful, while only 55% reported feeling safer after contacting police. Approximately 14% reported feeling less safe (Osborne et al., 2012).

For some victims, involving the police can also set into motion a series of consequences which they do not want. For example, victims do not always want to leave the relationship or family environment due to love or family bonds (Strube, 1988). More importantly, many victims are potentially isolated or without means to leave the abusive environment. Combined with underfunded domestic abuse services, the decision to disclose domestic abuse can result in homelessness (Cramer and Carter, 2002; Walby and Towers, 2012; Sanders-McDonagh et al., 2016; Office for National Statistics, 2019b). The upshot of analysing the factors influencing the decision to report domestic abuse to police is that many victims/survivors view calling the police as a last resort (Fitzgerald et al., 1995; Women's Aid, 2009).

The low level of reporting would leave us with a scant premise for investigating spillovers of police reports of domestic abuse. However, surveys of survivors reveal that they do disclose their abuse, just not necessarily to police: In England, more than 73% of victims of domestic abuse told a friend or relative about the abuse, compared with only 23% reporting to police (Osborne et al., 2012). While estimates of reporting rates to police vary (from as low as 2% to almost 50%), other work similarly finds that more than half of victims disclose their abuse to someone close to them (Greenberg and Ruback, 1992; Fisher et al., 2003; Coy and Kelly, 2011; Stark et al., 2013).

High rates of social disclosure paired with the high incidence of domestic abuse provide a context in which the disclosure of domestic abuse to police by others could affect the decision to report. This establishes the basis for our core hypothesis: Someone close to a victim of domestic abuse reporting their own abuse to police might induce the victim to contact police themselves.

2.4 | Spillover Channels

In our model, we distinguish between two channels of spillovers: Report-to-report and follow-up-to-report.

The first channel, report-to-report, accounts for spillovers due to information passing through social peer networks. Knowing that others in their vicinity reported domestic abuse to the police might affect victims' decision to report abuse. While this has not yet been studied for domestic abuse, some work demonstrates that information transmissions matters for the reporting of sexual violence: Cheng and Hsiaw (2020) develop a formal model of reporting sexual misconduct in the workplace. The context of workplace harassment differs from domestic abuse: Mainly, their work centres on corroboration, that is, multiple individuals need to report misconduct before action against an offender is taken. Corroboration is difficult to translate to the context of domestic abuse (where a perpetrator typically only abuses within their immediate household). Still, the Cheng and Hsiaw (2020) model illustrates that individuals face information frictions about how wide-spread a behaviour is, which affects their propensity to report.

Levy and Mattsson (2020) study the effect of the #MeToo movement on reporting behaviour and find that it resulted in a persistent increase of reports of sexual violence. They argue that their results are plausibly explained by a rapid change in social norms and information. Similarly, Iyer et al. (2012) and McDougal et al. (2018) find that visible social changes (the election of female politicians and a highly publicised case of sexual violence, respectively) lead to a large increase of reports of sexual violence.

The second channel of spillovers, follow-up-to-report, accounts for the effects of police intervention. Some studies have investigated the effects of police intervention on future violence within the same household, with mixed results (Hoppe et al., 2020; Hanmer et al., 1999). A meta analysis by Davis et al. (2008) of ten studies investigates follow-up visits by police officers and/or victim advocates as part of a violence interruption programme. Results indicate that the follow-up visits had no significant effect on the occurrence of repeat violence as measured by victim surveys (standardised difference in group means: -0.01, $p = 0.82$). Instead, there is a modest positive increase in reports to police (standardised difference in group means: 0.12, $p = 0.01$). The results suggest that while the follow-up visits do not reduce the likelihood of abuse, victims seem more confident about reporting the violence to police.

Our study tests an extension of this effect: Do police visits following an incident have an effect outside the directly affected household? The mechanism of this effect links back to the notion that victims of domestic abuse often do not recognize their abuse as such. Repeat visits by police to another call can then serve as validation and amplification: Police are taking the incident seriously and paying attention.

3 | DATA

Our data set covers all 11,691 calls to police about an incident of domestic abuse between 01/01/2017 to 31/12/2018 in an English city with over 300,000 inhabitants. To preserve the anonymity of the city, all figures of the spatial domain show only the maximum inscribed circle of the city area, that is the largest possible circle completely contained by the domain. However, all data points and results reported in the paper refer to the entirety of the available data, not just the circular area presented.

When someone calls the police for service, the call will be picked up by a call handler in the police contact centre. The handler will ask a series of questions to evaluate the situation and decide on an appropriate police response. If the call handler at this point assesses the situation to take place in the context of domestic abuse, they raise a flag in the system which notifies the responding officer of that context.

In the United Kingdom, there is no statutory crime of domestic abuse. But many forms of domestic violence constitute criminal offences such as assault, sexual offences, stalking or criminal damage. Police forces in the UK classify such incidents as domestic abuse if they meet the cross-government definition: "Any incident or pattern of incidents of controlling, coercive or threatening behaviour, violence or abuse between those aged 16 or over who are or have been intimate partners or family members regardless of gender or sexuality. This can encompass but is not limited to the following types of abuse: psychological, physical, sexual, financial, emotional." (Home Office, 2013).

In the past, English police forces have been criticized for not supplying responding officers with sufficient information. For example, officers may often not have any information on the perpetrator or know that the victim/survivor may be a repeat victim (Her Majesty's Inspectorate of Constabulary, 2014). Similarly, the initial call handler may not identify a situation as domestic abuse, but the police responders at the scene may do so.

A call enters our data set when either the handler or the responders classify the call as domestic abuse. Research conducted in the UK finds evidence that there is variation between call handlers and officers in their handling of domestic abuse: Female call handlers result in faster police response and cases handled by response teams with more female officers have lower legal attrition (Hawkins and Laxton, 2014).

For each call in our data set, we know the time when the call was placed and the location of the incident. A key feature of our analysis is the question of police follow-ups where officers return to places of domestic abuse. A key challenge here, however, is that we cannot consistently check *why* police officers return: Is it for a scheduled routine visit or is it because the domestic situation escalated and requires intervention? This reason for this inconsistency is inconsistent police record keeping. Some officers who return to the scene will link their new visit to the old case identifier which means that we can follow this link. However, some officers will also create a new case identifier unconnected to the original case. Is this because there is a new incident of domestic abuse at the house, therefore necessitating the creation of a new case file? Or is this simply an oversight on the officer's part?

While more than two thirds of return visits happen between business hours, suggestive of scheduled visits, our approach to this issue is conservative: Every return visit by police officers to the same address within two weeks of the initial call is classified as a follow-up visit, without distinction as to why the officers might have returned. If officers return after more than 15 days of the initial call, we consider it a new incident of domestic abuse due to escalated violence. This concerns only 29 calls in our data set. Choosing 15 days as the cut-off is motivated by police response: Police aim to respond to non-urgent incidents within 5 days (and within 15-60min to urgent incidents, depending on the urgency) which means that there is a reasonable range of days after an initial incident during which officers might follow up.

Taken together, this results in 6,084 initial calls for service and 2,286 follow-up visits by police in 2018. Figure 1a and Figure 1b show the temporal and spatial dimension of the raw data for that year. In 2017, the data contain 5,607 initial calls for service and 2,551 follow-up visits.

4 | METHOD

We formally introduce our specification of the Hawkes process in Section 4.1. In Section 4.2 we propose an extension to Zhuang and Mateu (2019), and in Section 4.3, we provide the details of the inference of the model, including the extension.

4.1 | Model

We treat reports of domestic abuse to police as events occurring at a particular point in space and time. Each event consists of a timestamp and a geographic coordinate. We can then treat these events as a realisation of a point process

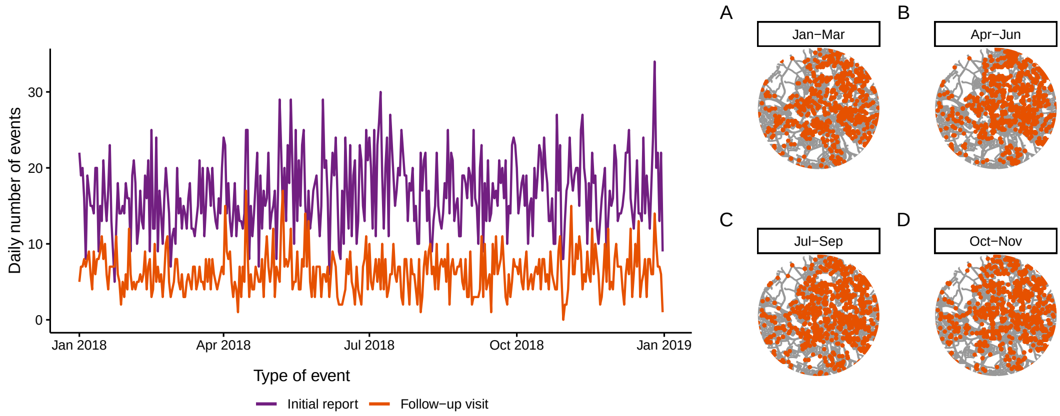

(a) Time series of events.

(b) Event locations by quarter.

FIGURE 1 Time and location of events of domestic abuse and follow-up visits by police. The plots show the data for the whole of 2018, as described in Section 3.

and model the process. This point process has n points where $\{t_1, t_2, \ldots, t_n\}$ denotes the time-ordered sequence of event times and $\{s_1, s_2, \ldots, s_n\}$ denotes the time-ordered sequence of event locations (Daley and Vere-Jones, 2003; van Lieshout, 2010). A point process is completely defined by its conditional intensity, which gives the expected rate of events, conditional on the history of the process of all events up to time t. More formally, the conditional intensity at time $t \in [0, T)$ and location $s \in X \subseteq \mathbb{R}^d$ is given by:

$$\lambda(t, s|\mathcal{H}_t) = \lim_{\Delta_t, \Delta_s \downarrow 0} \frac{\mathrm{E}\{N((t, t+\Delta_t) \times (s, s+\Delta_s)) \mid \mathcal{H}_t\}}{\Delta_t \Delta_s}, \tag{1}$$

where $\mathcal{H}_t$ is the history of the process up to and including t, $N(\cdot)$ is a counting measure for the number of events in a space-time volume specified by $(t, t+\Delta_t) \times (s, s+\Delta_s)$.

The conditional intensity for a Hawkes process is given by

$$\begin{aligned} \lambda(t, s|\mathcal{H}_t) &= \mu(t, s) + \int_0^t \int_X f(t-u, s-v) \mathrm{d}N(u \times v) \\ &= \mu(t, s) + \sum_{j: t_j < t} f(t - t_j, s - s_j), \end{aligned} \tag{2}$$

where $\mu(t, s)$ describes the background rate, and term $f(\cdot)$ is the self-exciting component of the process. The integral form counts the contribution of past events towards the intensity. The specification of $f(\cdot)$ determines for how long (in time) and how far (in geographical distance) past events can trigger events, in addition to the base background rate, $\mu(\cdot)$. The time component of triggering is uni-directional, i.e., forward in time, whereas the spatial component considers all points in the neighbourhood around s. Intensity $\lambda(t, s|\mathcal{H}_t) \geq 0$ must be greater than or equal to zero, so we set $\mu(t, s) \geq 0$ and $f(t, s) \geq 0$ for every t and s. For ease of notation, we will omit the explicit conditioning on $\mathcal{H}_t$, but $\lambda(t, s)$ depends on all past events $\mathcal{H}_t$, for all spatial locations s. Note that the process specified above is not a Poisson process (due to triggering, for two disjoint Borel sets A and B, random variables $N(A)$ and $N(B)$ are not independent)

The specification of the background and triggering components depends on the application context (Reinhart, 2018). Our specification of the background component follows Zhuang and Mateu (2019), who consider a decomposition of the background into periodic components as follows:

$$\mu(t,s) = m_0 \mu_{\text{trend}}(t)\mu_{\text{weekly}}(t)\mu_{\text{daily}}(t)\mu_{\text{area}}(s), \tag{3}$$

where $\mu_{\text{trend}}(t), \mu_{\text{weekly}}(t)$ and $\mu_{\text{daily}}(t)$ represent the trend term over the whole study window, the weekly and daily periodicity in the time dimension of the background rate, respectively. The term $\mu_{\text{area}}(s)$ is an estimate of the background spatial intensity in the study area. These terms are normalized to have mean 1. As a consequence, the non-negative term m_0 weights the overall contribution of the background component to the overall intensity (Loeffler and Flaxman, 2018; Zhuang and Mateu, 2019). Alternatively, one could incorporate covariates into the specification of the background rate $\mu(t,s)$. However, this approach requires that the covariates approximate the background rate well, as otherwise the parameter inference is considerably affected (Reinhart and Greenhouse, 2018). Park et al. (2021) offer a non-parametric estimation of a mapping from covariates to the background rate using generalised additive models. As we do not have access to relevant and up-to-date spatial covariates, we use a non-parametric approach.

For the triggering component, $f(\cdot)$, we follow a common practice in the literature and assume that it is separable in time and space, such that $f(t,s) = \theta g(t)h(s)$, where θ is an unknown constant (Meyer et al., 2012; Reinhart and Greenhouse, 2018; Zhuang and Mateu, 2019; Loeffler and Flaxman, 2018). We normalize g and h to integrate to 1 such that θ gives the average number of events coming from the trigger component. As explained in the following section, functions $g(\cdot)$ and $h(\cdot)$ are estimated non-parametrically from the data. We measure distance in kilometres and time in days, such that $t_i = 1.5$ denotes the time of an event i that took place 1.5 days (= 36 hours) since the beginning of the study window. In contrast with e.g., Kalair et al. (2021), we do not enforce monotonicity of $g(t)$ and $h(s)$ since we might expect the social dynamics of reporting to be non-monotonic.

We choose the maximum number of days an event (a first-time report or a follow-up visit) can trigger a first-time report at 40 days. This choice is driven by the need to allow for potentially delayed triggering (for example, Levy and Mattsson (2020) found it was around 30 days for #MeToo reporting) as well as to maintain computational tractability.

4.2 | Model extension: additional event types

We explicitly consider the effect of follow-up visits by police, in addition to the self-exciting effect of a new instance of reporting domestic abuse itself. Doing so requires us to consider additional event times and locations, those of police follow-ups. This results in an additional sequence of event times $\{t'_1, \dots, t'_k\}$ and of locations $\{s'_1, \dots, s'_k\}$, where (t'_j, s'_j) gives the time and location of a follow-up event j. These different event types are distinguished by a sequence of labels, $\{M_j\}_{j=1}^{n+k}$, where $M_j = 0$ if event j is a report of domestic abuse to police and $M_j = 1$ if event j is a follow-up police visit.

Note that despite the introduction of additional events, we are still only modelling the intensity of first-time domestic abuse reports (events for which $M_j = 0$), and not the follow-up visits ($M_j = 1$). The follow-ups ($M_j = 1$) serve as events that may trigger additional first-time reports of domestic abuse. As a result of this extension, Equation (2) takes the form

$$\lambda(t,s|\mathcal{H}_t) = \mu(t,s) + \sum_{\substack{j:t_j<t_i \\ t_j \in t^{\text{all}}}} f(t-t_j, s-s_j), \tag{4}$$

where t^{all} is the time-ordered sequence of *all* events in $t^{\text{reports}} \cup t^{\text{followups}}$. To avoid such cumbersome indexing, we abuse notation and assume that any time we index over j, we are implicitly indexing over all events in t^{all}. Using this shortcut, the full form of the conditional intensity function is:

$$\lambda(t, s) = m_0 \mu_{\text{trend}}(t) \mu_{\text{weekly}}(t) \mu_{\text{daily}}(t) \mu_{\text{area}}(s) + \sum_{j: t_j < t} \theta_{M_j} g(t - t_j) h(s - s_j), \tag{5}$$

where θ_{M_j} is the expected number of events triggered by events of type $M_j \in \{0, 1\}$.

4.3 | Inference

Our inference approach follows the inference procedure first proposed in Zhuang et al. (2002) and applied in the context of crime in Zhuang and Mateu (2019). We summarise the key elements of this procedure below. For full details, we refer the reader to Zhuang et al. (2002); Zhuang and Mateu (2019).

We specify $g(\cdot)$, $h(\cdot)$, and all the background components ($\mu_{\text{daily}}(\cdot)$, $\mu_{\text{weekly}}(\cdot)$, $\mu_{\text{trend}}(\cdot)$, $\mu_{\text{area}}(\cdot)$) in a non-parametric manner. This implies that both the background and the self-exciting components are estimated from the observed events alone, without assuming a parametric form. This data-driven estimation necessitates distinguishing events arising from the background intensity, $\mu(t, s)$, from those triggered by past events. This is known as *stochastic declustering* (Zhuang et al., 2002). Once the events have been classified as either coming from the background or due to self-excitation, individual components can be estimated.

4.3.1 | Stochastic declustering

We introduce two quantities which are essential for deriving the estimation procedure for the model components:

1. the probability that an event came from the background, rather than the trigger component,

$$\varphi_i = \text{P(event } i \text{ came from background)} = \frac{\mu(t_i, s_i)}{\lambda(t_i, s_i)} \tag{6}$$

2. and the probability that an event was triggered by a past event (either a first-time report or a follow-up visit),

$$\rho_{ij} = \text{P(event } i \text{ was triggered by j, } t_j < t_i) = \frac{f(t_i - t_j, s_i - s_j)}{\lambda(t_i, s_i)}. \tag{7}$$

By the law of total probability, it is clear that

$$\varphi_i + \sum_{j=1}^{t_j < t_i} \rho_{ij} = 1. \tag{8}$$

One way to interpret Equation (8) is that all events j preceding event i added probability mass ρ_{ij} onto the event i, which allows us to decompose event i into background and trigger.

4.3.2 | Estimation of the non-parametric model components

The estimation procedure by Zhuang et al. (2002) relies on the Georgii-Nguyen-Zessin formula (Georgii, 1976; Nguyen and Zessin, 1979; Daley and Vere-Jones, 2008). For a time interval $[T_1, T_2]$, area X and a non-negative function γ, the following holds:

$$\mathbb{E}\left[\int_{[T_1,T_2]\times X} \gamma(t,s)\mathrm{d}N(t\times s)\right] = \mathbb{E}\left[\int_{T_1}^{T_2}\int_X \gamma(t,s)\lambda(t,s)\mathrm{d}t\mathrm{d}s\right]. \tag{9}$$

Loosely speaking, we can evaluate $\gamma(\cdot)$ over all observed points and the expectation of this (i.e., the left-hand side of Equation (9)) is equivalent to the expectation of that function over the entire space weighted by the conditional intensity, $\lambda(\cdot)$. This property is viewed as a self-consistency equation for the process. This then allows us to rearrange terms and obtain the expectation of the function over the entire space.

For each of the model components, we can construct a non-negative function w which can be plugged into (9) and used to derive the estimates of the component. The function w represents the component's contribution to the overall intensity. We show the full derivation for μ_{trend}, others follow analogously.

We define

$$w^{\text{trend}}(t,s) = \frac{\mu_{\text{trend}}(t)\mu_{\text{area}}(s)}{\lambda(t,s)}., \tag{10}$$

and substitute it for $\gamma(\cdot)$ in (9). Considering the time interval $[t-\Delta_t, t+\Delta_t]$, where Δ_t is a small positive number, and the whole of domain X we obtain:

$$\begin{aligned}\sum_i w^{\text{trend}}(t_i,s_i)\mathbb{I}(t_i \in [t-\Delta_t, t+\Delta_t]) &\approx \int_{T_1}^{T_2}\int_X w^{\text{trend}}(u,v)\lambda(u,v)\mathbb{I}(u \in [t-\Delta_t, t+\Delta_t])\mathrm{d}u\mathrm{d}v \\ &= \int_{t-\Delta_t}^{t+\Delta_t} \mu_{\text{trend}}(u)\mathrm{d}u \int_X \mu_{\text{area}}(v)\mathrm{d}v \\ &\propto \int_{t-\Delta_t}^{t+\Delta_t} \mu_{\text{trend}}(u)\mathrm{d}u \\ &\approx \mu_{\text{trend}}(t)2\Delta_t. \end{aligned} \tag{11}$$

We rearrange the last expression to

$$\hat{\mu}_{\text{trend}}(t) \propto \sum_i \underbrace{\frac{\mu_{\text{trend}}(t_i)\mu_{\text{area}}(s_i)}{\lambda(t_i,s_i)}}_{:=w_i^{\text{trend}}} \mathbb{I}(t_i \in [t-\Delta_t, t+\Delta_t]). \tag{12}$$

Finally, we smooth this histogram-like estimate by using a kernel density estimate – we replace the indicator function in Equation (12) by a Gaussian kernel $Z(\cdot)$ with bandwidth b_{trend} to obtain:

$$\hat{\mu}_{\text{trend}}(t) \propto \sum_i w_i^{\text{trend}} Z(t-t_i; b_{\text{trend}}). \tag{13}$$

Similarly to Equation (10), we can define functions w and then estimators for all background components:

$$\hat{\mu}_{\text{daily}}(t) \propto \sum_i w_i^{\text{daily}} \sum_{k=0}^{T} Z(t - (t_i - \lfloor t_i \rfloor + k); b_{\text{daily}}) \tag{14}$$

$$\hat{\mu}_{\text{weekly}}(t) \propto \sum_i w_i^{\text{weekly}} \sum_{k=0}^{\lfloor T/7 \rfloor} Z(t - (t_i - 7\lfloor t_i/7 \rfloor + 7k); b_{\text{weekly}}) \tag{15}$$

$$\hat{\mu}_{\text{area}}(s) \propto \sum_i \varphi_i Z(s - s_i; b_{\text{area}}), \tag{16}$$

where $\lfloor x \rfloor$ denotes the largest integer smaller than or equal to x. The periodicity in μ_{daily} and μ_{weekly} comes from mapping the input event time t into the periodic domain. For the daily component, we simply subtract the day on which the event took place and are left with the time of day on which the event took place ($t_i - \lfloor t_i \rfloor$). For the weekly component, we subtract the week from the event time ($t_i - 7\lfloor t_i/7 \rfloor$). The bandwidth for the spatial component, b_{area}, is set to be event-specific. This adaptive bandwidth accounts for the fact that a single bandwidth is often a poor choice with clustered point processes because it oversmooths some areas while being too noisy in other areas (Reinhart, 2018). Instead, b_{area} is set such that a spatial disk centred on event i with radius b_i contains n_p other events (Zhuang et al., 2002; Zhuang, 2011).

For the triggering components, we obtain similar expressions. As before, we define a function

$$w^f(t, s, u, v) = \begin{cases} g(u-t)h(v-s)/\lambda(u, v) & \text{if } u > t, \\ 0 & \text{otherwise,} \end{cases}$$

and substitute it for γ into Equation (9). For a fixed t_j and s_j, and considering the time interval $[t - \Delta_t, t + \Delta_t]$, we obtain

$$\begin{aligned} \sum_i w^f(t_j, s_j, t_i, s_i)\mathbb{I}(t_i - t_j \in [t - \Delta_t, t + \Delta_t]) &\approx \int_0^T \int_X w^f(t_j, s_j, u, v)\mathbb{I}(u - t_j \in [t - \Delta_t, t + \Delta_t])\lambda(u, v)\mathrm{d}u\mathrm{d}v \\ &\approx \int_{t-\Delta_t}^{t+\Delta_t} g(u - t_j)\mathrm{d}u \int_X h(v - s_j)\mathrm{d}v \\ &\propto g(t). \end{aligned}$$

To now obtain a more stable estimate for $g(t)$, we can additionally sum over the left-hand side for all j and obtain

$$g(t) \propto \sum_{i,j} w^f(t_j, s_j, t_i, s_i)\mathbb{I}(t_i - t_j \in [t - \Delta_t, t + \Delta_t]).$$

Observing that $w^f(t_j, s_j, t_i, s_i)$ is just ρ_{ij}, the estimate of $g(t)$ simplifies to:

$$\hat{g}(t) \propto \sum_{i,j} \rho_{ij}\mathbb{I}(t_i - t_j \in [t - \Delta_t, t + \Delta_t]). \tag{17}$$

Similarly, we can obtain an estimate for $h(s)$:

$$\hat{h}(s) \propto \sum_{i,j} \rho_{ij}\mathbb{I}(s_i - s_j \in [s - \Delta_s, s + \Delta_s]), \tag{18}$$

where $\Delta_s = [\delta_s, \delta_s]^\top$ and δ_s is a small positive number. The membership in the interval is applied coordinate-wise as $s, s_i, s_j \in \mathbb{R}^2$. In addition, we apply kernel smoothing and the necessary repetition correction which counts how many times the triggering effect could be potentially observed at a specific time lag or spatial distance (Zhuang and Mateu, 2019):

$$\hat{g}(t) \propto \frac{\sum_{i,j} \rho_{ij} Z(t - (t_i - t_j); b_g)}{\sum_j \mathbb{I}(t_j + t \leq T)} \tag{19}$$

$$\hat{h}(s) \propto \frac{\sum_{i,j} \rho_{ij} Z(s - (s_i - s_j); b_h)}{\sum_j \mathbb{I}((s_j + s) \in X)}. \tag{20}$$

4.3.3 | Estimation of the weighting terms and including additional event types

Having derived estimates of all background and triggering components, we detail the procedure for estimating the weighting terms, m_0 and θ_M from Equation (5). The original paper by Zhuang and Mateu (2019) does not consider additional event types, which is why we extend their model to accommodate the effect of follow-up visits by police.

We use the maximum likelihood approach for estimating the weighting terms, $\Theta = \{m_0, \theta_0, \theta_1\}$. The log-likelihood for a point process specified by its conditional density is given by

$$\ell(\Theta) = \sum_i \log \lambda(t_i, s_i) - \int_0^T \int_X \lambda(t, s) \mathrm{d}t \mathrm{d}s. \tag{21}$$

For more details, we refer the reader to Daley and Vere-Jones (2003, Ch. 7). Differentiating the log-likelihoood with respect to each of the parameters Θ and equating to zero, we obtain

$$\begin{aligned} \partial \ell(\Theta) / \partial m_0 &= 0 \\ &= \sum_i \frac{\mu_{\text{trend}}(t_i) \mu_{\text{weekly}}(t_i) \mu_{\text{daily}}(t_i) \mu_{\text{area}}(s_i)}{\lambda(t_i, s_i)} \\ &\quad - \int_0^T \int_X \mu_{\text{trend}}(t) \mu_{\text{weekly}}(t) \mu_{\text{daily}}(t) \mu_{\text{area}}(s) \mathrm{d}t \mathrm{d}s \end{aligned} \tag{22}$$

and

$$\begin{aligned} \partial \ell(\Theta) / \partial \theta_0 &= 0 \\ &= \sum_i \frac{\sum_{j: t_j < t_i} \mathbb{I}(M_j = 0) g(t_i - t_j) h(s_i - s_j)}{\lambda(t_i, s_i)} \\ &\quad - \int_0^T \int_X \sum_{j: t_j < t} \mathbb{I}(M_j = 0) g(t - t_j) h(s - s_j) \mathrm{d}s \mathrm{d}t. \end{aligned} \tag{23}$$

Differentiation with respect to θ_1 is analogous to θ_0. For full details, see Section 1 of the Supplementary Material. Similar to Zhuang and Mateu (2019), this system of equations can be solved by basing the estimates for m_0, θ_0 and θ_1 in inference round $(k + 1)$ on estimated quantities from round (k).

4.3.4 | Kernel edge correction

Kernel density estimates are well-known to behave poorly around the edges. To address this problem, edge correction techniques are necessary. For the periodic and area kernels and the kernel smoothing $g(t)$ and $h(s)$, we use a truncated kernel which normalizes the kernel density estimator by its integral over the support (Hall and Turlach, 1999). For example, the estimator for μ_{daily} from Equation (14) is modified to:

$$\hat{\mu}_{\text{daily}}(t) \propto \sum_i w_i^{\text{daily}} \frac{\sum_{k=0}^{T} Z(t - t_i + \lfloor t_i \rfloor - k; b_{\text{daily}})}{\int_0^T Z(u - t_i; b_{\text{daily}}) \mathrm{d}u}. \tag{24}$$

This modification was not sufficient to ensure sensible edge behaviour for the trend kernel. That is because the support for the trend kernel is bounded between $[0, T]$ and the density of points used for estimation is lower than it is for estimation of b_{daily} and b_{weekly}, where all events are projected onto smaller intervals. Instead, we apply an edge correction proposed by Schuster (1985) called boundary folding, where the density "leaking" outside the support is mirrored or folded back onto the support. For a kernel with support in $[a, b]$, we correct the standard kernel density estimator $f_h(x) = \frac{1}{nh} \sum_i K(\frac{x - x_i}{h})$ to

$$f_h(x) = \frac{1}{nh} \sum_i K\left(\frac{x - 2a + x_i}{h}\right) + K\left(\frac{x - x_i}{h}\right) + K\left(\frac{x + 2b - x_i}{h}\right). \tag{25}$$

5 | RESULTS

Firstly, we discuss the selection of model parameters, before we provide interpretation of the inferred quantities. Afterwards, we discuss different model validation checks and the consistency of inferred quantities with other results from the relevant criminological literature.

5.1 | Choice of Model Parameters

The inference procedure described in the previous section requires setting various parameters: the kernel bandwidths (b_{daily}, b_{trend}, b_{weekly}, b_g and b_h) and the number of neighbours, n_p, required to define a disc radius around each event which is eventually used as b_{area} for each individual event (see Equation (16)). To proceed, we perform a grid search over the space of parameters, minimising the out-of-sample prediction error. To guide the search, i.e., limit the search space, we also followed advice from the literature. We compute the prediction error by comparing the realised number of events from 2018, N_{realised}, to the expected number of events obtained by fitting the model to the data from 2017. More specifically, for each 1km $\times$ 1km cell c of the domain, we compute the expected number of events over the period of one year, $N_{\text{predicted}}$. The prediction error is then given by

$$\text{Prediction Error} = \sqrt{\frac{1}{n} \sum_c \left(N_{\text{predicted}}^{(c)} - N_{\text{realised}}^{(c)}\right)^2}. \tag{26}$$

As already discussed in Section 4.3.1, b_{area} is set to be adaptive, so it does not require an explicit choice. However, it does depend on choice of n_p, the number of neighbours to the event. In an earthquake modelling setting, Zhuang (2011) proposes setting n_p between 3 and 6. This choice has led to overfitting in our context – assessed by the out-of-sample prediction error. We set $n_p = 10$.

TABLE 1 The inferred values for the model parameters. The confidence intervals for the parameters are not available analytically, and simulation-based methods would be computationally too expensive as each simulation of the model takes several hours.

Parameter	Estimate
m_0	0.170
θ_0	1.05×10^{-9}
θ_1	3.08×10^{-9}

Based on the grid search (see Section 4 of the Supplementary Material for full output), we set $b_{\text{daily}} = 3/24 \approx 0.13$, $b_{\text{weekly}} = 16/24 \approx 0.67$, $b_{\text{trend}} = 30$. Bandwidths for Gaussian kernels provide an intuitive interpretation for the temporal components of the background term: the bandwidth corresponds to the temporal range over which we want the kernel to smooth over. For example, by setting $b_{\text{daily}} = 3/24$, one standard deviation of the corresponding kernel function corresponds to 3 hours. That implies that 99.7% (= 3 standard deviations) of the contributions to our kernel density estimate come from events within 9 hours of the event.

Lastly, we set $b_g = 1.0$ and $b_h = 1.0$. Experimenting with different values suggests that the inference is not sensitive to the choice of b_g and b_h. This could be for several reasons: 1) as we will see in the following sections, the estimates for θ_0 and θ_1 are approximately zero, thus effectively excluding the contribution from the triggering functions $g(\cdot)$ and $h(\cdot)$; 2) due to the high density of points for the kernel density estimate as $O(N^2)$ pair-wise distances (both temporal and spatial) are mapped to a small temporal and spatial domain (we ignore pairs of events that are more than 3km or more than 40 days apart). As discussed in Section 3 of the Supplementary Material, we explored different initializations of $g(\cdot)$. Our experiments have shown that the inference was insensitive to this choice and the same conclusions have been made even if the chosen initialization forced a delayed peak of the triggering in time as proposed by Gilmour (2019).

5.2 | Interpretation of the Inferred Model Components

Having chosen model parameters in the section above, we discuss the inferred weighting terms m_0, θ_0, and θ_1 in Equation (5). Table 1 summarises the inferred values. Because the weights of the triggering component (θ_0, θ_1) are almost zero, the main contribution to the intensity of first-time reports of domestic abuse comes from the background component of the intensity. The triggering weights mean that one report of domestic abuse triggers, on average, 1.05×10^{-9} further reports. Similarly, one follow-up by police triggers, on average, 3.08×10^{-9} reports of domestic abuse. Together, the model implies that of the 6,084 initial reports of domestic abuse, 0.00013429 reports were triggered by other events. To reiterate, there is very little evidence to support the notion of spillovers in domestic abuse reporting. Next, we discuss the individual background and triggering components that have been inferred.

In Figures 2a through 2e we visualize the estimated intensities coming from the background components. As shown in Equation (3), these estimated intensities are multiplied together and then weighted by $m_0 = 0.170$. The trend component in Figure 2a demonstrates that there is no dominant trend in domestic abuse reports over our study period since most of the normalized kernel density lies between 0.9 and 1.1, i.e., close to the mean of 1. We observe an increase in domestic abuse reporting, however, in the summer months of July and August. We document strong time of day and day of week effects. There are remarkably few calls between the hours of midnight and 4am, with the estimated intensity increasing during the day. We observe two peaks of intensity, one between 12:00 and 13:00

and another one between 20:00 and 21:00 before calls drop off at night again. Looking at the weekly periodicity, we observe a strong weekend effect as calls begin to pick up from Friday onwards throughout the weekend. Together, the daily and weekly periodicity visualized in Figure 2d show that Friday evenings, Saturday evenings and Sunday mornings are particularly high-intensity periods for reports of domestic abuse.

These findings mirror those found in other studies: Reports of domestic abuse are lowest during the week and strongly increase on the weekend, with Sunday being the peak day (Rotton and Cohn, 2001; Brimicombe and Cafe, 2012).

Even though the inferred effect of triggering is zero (given by θ_0 and θ_1), the inferences of the forms of $g(\cdot)$ and $h(\cdot)$ can still be interpreted. We visualize the estimated $g(\cdot)$ and $h(\cdot)$ in Figure 3a and 3b. Both figures show that the potential for an event to trigger offshoot events peaks at the time and the location of the event and decreases monotonically, with time for $g(\cdot)$, and with increasing distance from the event for $h(\cdot)$.

Lastly, we find that there is marked variation in space. Specifically, we find that for some locations, the estimated spatial intensity of the background is particularly pronounced. This is clear from the small, dark spots in Figure 2e. This raises a question of why these areas might see such high levels of domestic abuse reporting. We discuss this in the section on contextual model validation.

In summary, our model finds no evidence of spillovers in domestic abuse reporting.

5.3 | Model Validation

Since Hawkes processes are complex statistical objects with challenging inference procedures, it is important to validate the plausibility of model outputs. For example, Reinhart and Greenhouse (2018) show that when the background component does not provide a good fit to the data, the triggering component is inflated. In other words, one typically over-estimates the triggering component. This is not the case for our model, since our estimates of the triggering component are negligibly small. Still, we perform two additional model validation checks. In Section 5.3.1, we perform the transformed time process test, in Section 5.3.2 we show the results of the Voronoi residuals test and we conclude with a contextual consistency check in Section 5.4.

5.3.1 | Transformed Time Process

A common way of checking if the model is well-calibrated is a temporal residual plot. To do so, we calculate the following function of event times t_i

$$t_i \rightarrow \tau_i = \int_0^{t_i} \int_X \lambda(u, v) \mathrm{d}u \mathrm{d}v, \tag{27}$$

such that τ_i is the expected number of events in the time interval $[0, t_i)$ or the cumulative number of events by event time t_i. This simple transformation uses the time-rescaling theorem: If the model is correct, the sequence of τ_i is a stationary Poisson process with unit rate (Ogata, 1988; Brown et al., 2002). Accordingly, a plot of the event index i against τ_i should form a 45° diagonal line. This property can be used to assess model fit: If $\hat{\lambda}$ is a good approximation of the true model, then its sequence of $\hat{\tau}_i$ will behave similarly to the sequence of theoretical τ_i (Schoenberg, 2002). For the theoretical τ_i we can construct confidence intervals for each τ_i by taking the $\frac{\alpha}{2}$ and $1 - \frac{\alpha}{2}$ quantiles of a Beta distribution with parameters $(i + 1, n - i + 1)$ and then multiply by n (Zhuang and Mateu, 2019). We set α at the usual 0.05 level.

In Figure 4 we plot the deviation of $\hat{\tau}_i$ from the diagonal against the event index i to verify how far away from

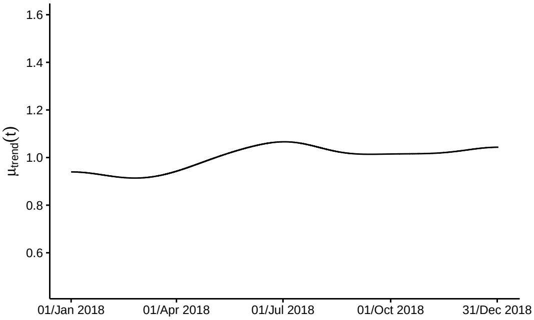

(a) Temporal trend, $\mu_{\text{trend}}(t)$.

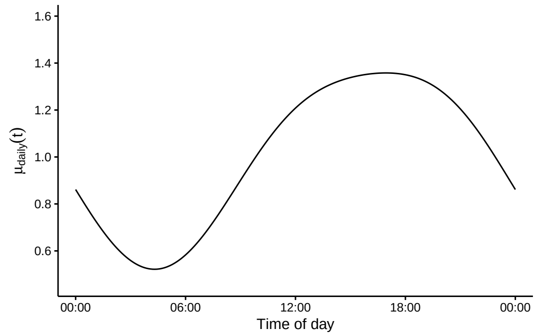

(b) Daily periodicity, $\mu_{\text{daily}}(t)$.

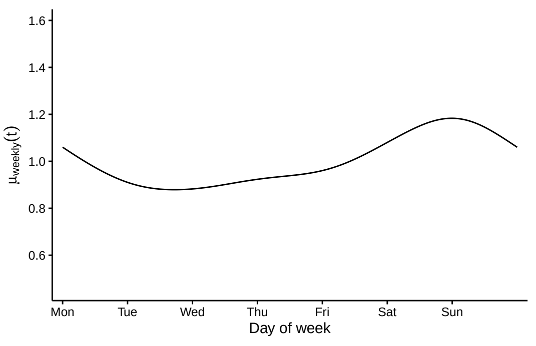

(c) Weekly periodicity, $\mu_{\text{weekly}}(t)$.

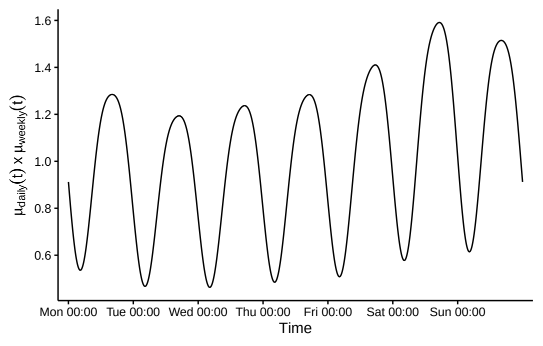

(d) Daily and weekly periodicity combined, $\mu_{\text{weekly}}(t)\mu_{\text{daily}}(t)$.

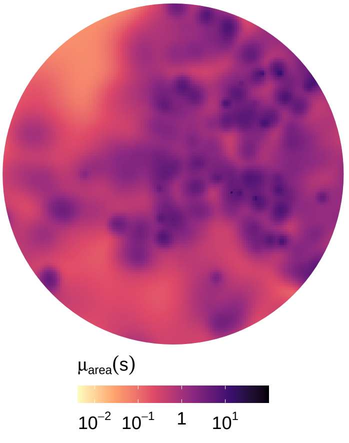

(e) Spatial background intensity, $\mu_{\text{area}}(s)$.

FIGURE 2 Estimated components of the background intensity, $\mu(t, s)$, as specified in Equation (3).

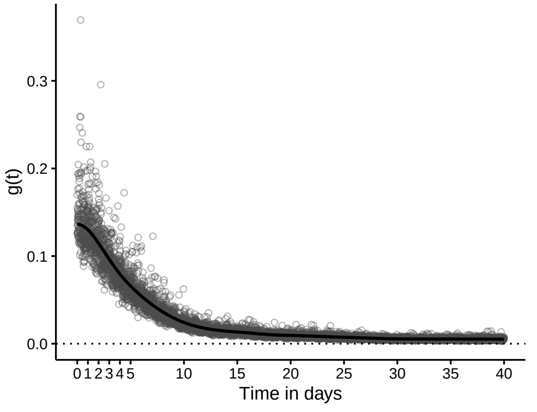

(a) Temporal excitation, $g(t)$. The points on the plot show evaluations of a version of $g(t_i - t_j)$ before a smoother is applied, for all relevant pairs of events. The line shows the smoother version, $g(t)$.

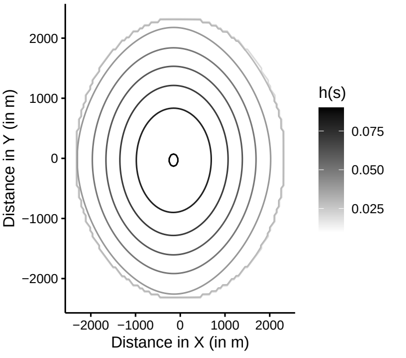

(b) Spatial excitation, $h(s)$. The contour plot shows the surface representing function $h(s)$.

FIGURE 3 Estimates of the triggering component of the intensity, $f(t, s) = \theta_M g(t)h(s)$.

the diagonal our model deviates. Indeed, we find that $\hat{\tau}_i$ is within the 95% confidence bounds of the true model and that our model is therefore a reasonable approximation.

5.3.2 | Voronoi Residuals

Similar to the temporal residual plot, one can examine the spatial residuals. To do so, one chooses a time window $[T_1, T_2]$ and computes the difference between the predicted number of events over some division of spatial domain in the time window and the actual number. As discussed in Bray et al. (2014), a regular grid is a poor choice to divide the spatial domain, as the distribution of the residuals becomes sensitive to the choice of the size of cells. Instead, the paper suggests the use of a Voronoi tessellation of the domain, which produces a set of convex non-overlapping polygons such that each polygon contains a single event. For a polygon c of the tessellation, the residuals for that polygon, R_c, are given by

$$R_c = 1 - \int_{T_1}^{T_2} \int_{S_c} \lambda(t, s) \mathrm{d}s\mathrm{d}t, \tag{28}$$

where S_c is the set containing all locations in the cell c. It has been shown that for a homogeneous Poisson process, the residuals follow

$$R_c \sim 1 - \Gamma(3.569, 3.569), \tag{29}$$

where $\Gamma(\alpha, \beta)$ is the Gamma distribution with location parameter α and rate parameter β. Bray et al. (2014) argue that this holds for non-homogeneous point processes as well, provided the conditional intensity is approximately constant near the location in question.

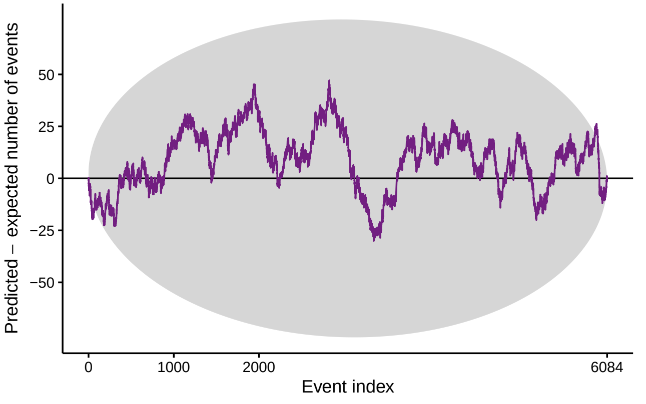

FIGURE 4 Deviation of the transformed time sequence (purple) from the theoretical (black) sequence, with 95% confidence bands (grey). The transformed time sequence is defined in Equation (27).

We perform the integration over the time domain in (28) using a simple Riemann sum, and we estimate the integral over space using Monte Carlo sampling. For more details, see Section 5 of the Supplementary Material. We show the residuals obtained using a Voronoi tessellation for the time window $[0, T]$ in Figure 5. We discuss some statistical properties of these residuals in Section 5 of the Supplementary Material. Overall, we find that our model slightly under-predicts the number of events in a cell, as indicated by the negative residuals, while still providing a good fit overall.

5.4 | Contextual Consistency Checks

A number of studies document a relationship between levels of domestic abuse and deprivation (e.g., Gracia et al., 2015). In Figure 6, we show the spatial background intensity at event locations against local deprivation. Deprivation is measured using the 2015 index of multiple deprivation, a composite index combining measures of deprivation from seven domains such as income, employment, education and health (Office for National Statistics, 2015). Overall, the relationship between abuse reporting and deprivation is not straightforward: We see high levels of reporting in areas with high and low levels of deprivation. However, event locations with very high spatial background intensity (points in darker colours) are consistently in areas with high levels of deprivation.

6 | CONCLUSIONS

This paper studies if the reporting of domestic abuse by victims exhibits triggering behaviour, similar to the criminal behaviour of offenders.

Identifying whether events happen close to each other in time and space as a result of the spatio-temporal environment or because they were triggered by a past event is a challenging task. Hawkes processes are a conducive tool to disentangling triggering from clustering. They decompose the intensity of events into a background component modelling the periodical and geographic occurrence of domestic abuse, and into a triggering component modelling reports coming about as a result of past reports. We employed a non-parametric specification of the Hawkes process proposed by Zhuang and Mateu (2019) and developed an extension which allows for considering additional event

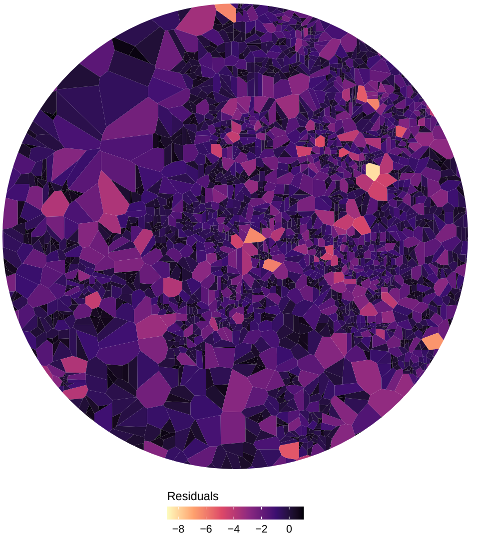

FIGURE 5 Voronoi residuals for the chosen model. The residuals are computed by considering the predicted number of crimes for each cell, see Equation (28). The cells are formed by a Voronoi tessellation, which partitions the domain so that each cell contains exactly one crime event.

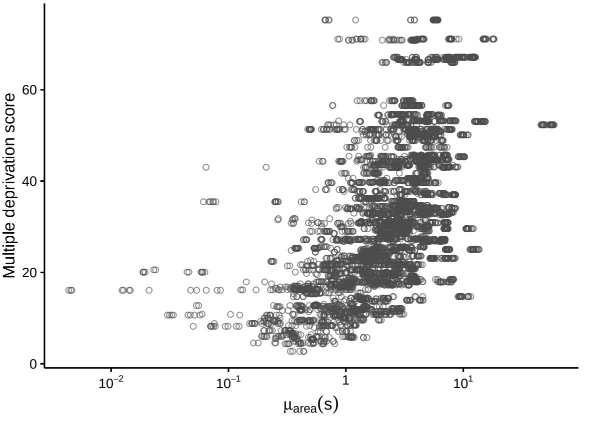

FIGURE 6 Spatial intensity, $\mu(s_i)$, evaluated at all events of first-time reports against deprivation.

types. The extension allows two types of events to trigger reports of domestic abuse: the reports themselves, and the follow-up visits by police.

Analysing data from one year of calls for service concerning domestic abuse in a large English city, we find no convincing evidence for spillover effects. These effects do not plausibly account for spillovers due to information sharing by victims in a social or neighbourhood network. Overall, reporting of domestic abuse by victims does not appear to exhibit any triggering behaviour since only 0.13×10^{-3} % of the reports in our sample are predicted to have been triggered.

The estimation of the background intensity of domestic abuse reporting shows that reporting follows highly periodic patterns. Calls for service of domestic abuse increase on the weekend and particularly in the evening. The spatial component of the inferred background intensity shows several hotspot areas with high levels of reported domestic abuse.

There are a few reasons why we did not find an effect. Our work assumed that the spillover effects spread purely by spatial proximity, ignoring the exact structure of social networks, which we expect to be a more realistic representation of the information flows needed for any spillovers to take place. For several reasons, such as privacy, it is not realistic to expect to have access to the exact structure of the social networks. One possible approach is to infer these latent structures from the data. Future work could consider approaches similar to Linderman and Adams (2014). A more immediate suggestion is to improve the data collection so that each report of domestic abuse indicates whether the victim has been encouraged to report by external factors, such as talking to a person from their social circles.

The functional form of the triggering component of the Hawkes process might not be able to accommodate the shape of spillover effects. In their study of the effect of the #MeToo movement, Levy and Mattsson (2020) find that the effect is largest on crimes reported a month after they took place. Therefore, it is possible that assuming temporal spillovers in the limited time frame of 40 days is an ill fit to the nature of spillovers in reporting.

The non-discrete temporal nature of domestic abuse poses a challenge to testing for spillovers. Domestic abuse consists of both discrete events such as assaults but more so of patterns of abusive behaviour. With this combination, domestic abuse cannot be thought of as a sequence of individual criminal offences (Hawkins and Laxton, 2014). The discretization of a latent, ongoing phenomenon such as domestic abuse into reports may not capture if and when victims share information about domestic abuse and the police response to it. In contrast to domestic abuse, burglaries, homicides and shootings are all discrete events with a distinct time and place of offence which gives crime victims a concrete event to report.

It may also be that expecting reporting to police to affect further domestic abuse reporting is overly optimistic. Davis et al. (2008)' meta analysis showed that follow-up visits by police officers did not have a significant effect on the reporting of future violence (e.g., Davis et al., 2010; Pate et al., 1992). If such visits already do not encourage victims in the treated households to turn to police again, the likelihood that such visits would have any effect outside those households is low.

We believe that our work could be adapted to other contexts. Other crime types could be investigated from the reporting standpoint that we have advocated for in this paper, instead of the offender's perspective. *Sexual offences* are an obvious candidate because they share characteristics of domestic abuse, especially regarding the decision to report. Other areas of science that employ Hawkes processes could benefit from our extension of the non-parametric framework with additional event types.

Implementation

Code which performs the inference method detailed in Section 4 is available at `https://github.com/laravomfell/reporting_spillovers`. It includes a routine to generate synthetic data, as we are not allowed to share the original dataset.

Acknowledgements

An anonymous English police force provided access to the data. They provided comments on the manuscript, but had no role in study design and analysis, decision to publish or preparation of the manuscript.

Conflict of interest

The authors declare no competing interests.

Data availablity

This research is based on data resources provided by an anonymous English police force. Data were originally collected as part of routine police record-keeping. The data are not publicly available and were provided to the authors under an Information Sharing Agreement with the police force. Under the terms of this agreement, the authors are not at liberty to share the data. Other researchers can contact the authors to inquire about obtaining a data sharing agreement with the police force.

Funding statement

Jan Povala was supported by EPSRC grant EP/L015129/1.

references

Agüero, J. M. and Frisancho, V. (2021) Measuring violence against women with experimental methods. *Economic Development and Cultural Change*. Forthcoming.

Aizer, A. and Currie, J. (2004) Networks or neighborhoods? correlations in the use of publicly-funded maternity care in california. *Journal of Public Economics*, **88**, 2573–2585.

Aronow, P. M., Eckles, D., Samii, C. and Zonszein, S. (2021) Spillover effects in experimental data. In *Advances in Experimental Political Science* (eds. J. N. Druckman and D. P. Green), 289–319. Cambridge University Press.

Bachman, R. and Coker, A. L. (1995) Police involvement in domestic violence: The interactive effects of victim injury, offender's history of violence, and race. *Violence and Victims*, **10**, 91–106.

Bernasco, W. (2008) Them again? Same-offender involvement in repeat and near repeat burglaries. *European Journal of Criminology*, **5**, 411–431.

Bernasco, W., Johnson, S. D. and Ruiter, S. (2015) Learning where to offend: Effects of past on future burglary locations. *Applied Geography*, **60**, 120–129.

Bertrand, M., Luttmer, E. F. and Mullainathan, S. (2000) Network effects and welfare cultures. *The Quarterly Journal of Economics*, **115**, 1019–1055.

Beyer, K., Wallis, A. B. and Hamberger, L. K. (2015) Neighborhood environment and intimate partner violence: A systematic review. *Trauma, Violence, & Abuse*, **16**, 16–47.

Brantingham, P. J., Yuan, B., Sundback, N., Schoenberg, F. P., Bertozzi, A. L., Gordon, J. and Malinowski, S. (2018) Does violence interruption work? *Proceedings of the National Academy of Sciences*, **8**, 1–6.

Brantingham, P. L. and Brantingham, P. J. (1993a) Environment, routine and situation: Toward a pattern theory of crime. *Advances in criminological theory*, **5**, 259–94.

— (1993b) Nodes, paths and edges: Considerations on the complexity of crime and the physical environment. *Journal of Environmental Psychology*, **13**, 3–28.

Bray, A., Wong, K., Barr, C. D. and Schoenberg, F. P. (2014) Voronoi residual analysis of spatial point process models with applications to California earthquake forecasts. *The Annals of Applied Statistics*, **8**, 2247–2267.

Brimicombe, A. and Cafe, R. (2012) Beware, win or lose: Domestic violence and the world cup. *Significance*, **9**, 32–35.

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 Computation*, **14**, 325–346.

Cai, J., De Janvry, A. and Sadoulet, E. (2015) Social networks and the decision to insure. *American Economic Journal: Applied Economics*, **7**, 81–108.

Cheng, I.-H. and Hsiaw, A. (2020) Reporting sexual misconduct in the #metoo era. *American Economic Journal: Microeconomics*. [Forthcoming] Available from: `http://people.brandeis.edu/~ahsiaw/chenghsiaw_reportingmisconduct.pdf`. [Accessed 2020-12-10].

Clarke, R. V. and Cornish, D. B. (1985) Modeling offenders' decisions: A framework for research and policy. *Crime and justice*, **6**, 147–185.

Cohen, L. E. and Felson, M. (1979) Social change and crime rate trends: A routine activity approach. *American Sociological Review*, **44**, 588–608.

Cohn, E. G. (1993) The prediction of police calls for service: The influence of weather and temporal variables on rape and domestic violence. *Journal of Environmental Psychology*, **13**, 71–83.

Coy, M. and Kelly, L. (2011) *Islands in the Stream: An Evaluation of Four Independent Domestic Violence Advocacy Schemes*. London: Henry Smith Charity, London Metropolitan University and Trust for London.

Cramer, H. and Carter, M. (2002) *Homelessness: what's gender got to do with it*. London: Shelter.

Daley, D. J. and Vere-Jones, D. (2003) *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods*. New York: Springer.

Daley, D. J. and Vere-Jones, D. (2008) *An Introduction to the Theory of Point Processes, Volume II: General Theory and Structure*. Probability and Its Applications. New York, NY: Springer-Verlag New York.

Davis, R. C., Weisburd, D. and Hamilton, E. E. (2010) Preventing repeat incidents of family violence: A randomized field test of a second responder program. *Journal of Experimental Criminology*, **6**, 397–418.

Davis, R. C., Weisburd, D. and Taylor, B. (2008) Effects of second responder programs on repeat incidents of family abuse: A systematic review. *Campbell Systematic Reviews*, **4**, 1–38.

Drago, F., Mengel, F. and Traxler, C. (2020) Compliance behavior in networks: Evidence from a field experiment. *American Economic Journal: Applied Economics*, **12**, 96–133.

Ellsberg, M., Heise, L., Pena, R., Agurto, S. and Winkvist, A. (2001) Researching domestic violence against women: methodological and ethical considerations. *Studies in Family Planning*, **32**, 1–16.

Emery, C. R. (2010) Examining an extension of johnson's hypothesis: Is male perpetrated intimate partner violence more underreported than female violence? *Journal of Family Violence*, **25**, 173–181.

Fadlon, I. and Nielsen, T. H. (2019) Family health behaviors. *American Economic Review*, **109**, 3162–3191.

Farrell, G. and Pease, K. (1993) Once bitten, twice bitten: Repeat victimisation and its implications for crime prevention. *Crime Prevention Unit Paper 46*, Home Office Police Research Group London.

Fisher, B. S., Daigle, L. E., Cullen, F. T. and Turner, M. G. (2003) Reporting sexual victimization to the police and others: Results from a national-level study of college women. *Criminal Justice and Behavior*, **30**, 6–38.

Fitzgerald, L. F., Swan, S. and Fischer, K. (1995) Why didn't she just report him? the psychological and legal implications of women's responses to sexual harassment. *Journal of Social Issues*, **51**, 117–138.

Flaxman, S., Chirico, M., Pereira, P., Loeffler, C. et al. (2019) Scalable high-resolution forecasting of sparse spatiotemporal events with kernel methods: A winning solution to the nij "real-time crime forecasting challenge". *The Annals of Applied Statistics*, **13**, 2564–2585.

Georgii, H.-O. (1976) Canonical and grand canonical gibbs states for continuum systems. *Communications in Mathematical Physics*, **48**, 31–51.

Gilmour, C. (2019) *Self-Exciting Point Processes and Their Applications to Crime Data*. Ph.D. thesis, University of Strathclyde, Glasgow.

Gorr, W. L. and Lee, Y. (2015) Early warning system for temporary crime hot spots. *Journal of Quantitative Criminology*, **31**, 25–47.

Goudriaan, H., Wittebrood, K. and Nieuwbeerta, P. (2006) Neighbourhood characteristics and reporting crime: Effects of social cohesion, confidence in police effectiveness and socio-economic disadvantage. *British Journal of Criminology*, **46**, 719–742.

Gracia, E., López-Quílez, A., Marco, M., Lladosa, S. and Lila, M. (2015) The spatial epidemiology of intimate partner violence: do neighborhoods matter? *American Journal of Epidemiology*, **182**, 58–66.

Greenberg, M. S. and Ruback, R. B. (1992) *After the crime: Victim decision making*. New York: Springer Science & Business Media.

Hall, P. and Turlach, B. A. (1999) Reducing bias in curve estimation by use of weights. *Computational Statistics & Data Analysis*, **30**, 67–86.

Halloran, M. E. and Hudgens, M. G. (2016) Dependent happenings: a recent methodological review. *Current Epidemiology Reports*, **3**, 297–305.

Hanmer, J., Griffiths, S. and Jerwood, D. (1999) Arresting evidence: domestic violence and repeat victimisation. *Police Research Series Paper 104*, Home Office, Policing & Reducing Crime Unit.

Hawkes, A. G. (1971) Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, **58**, 83–90.

Hawkins, S. and Laxton, C. (2014) *Women's access to justice: from reporting to sentencing*. London: All-Party Parliamentary Group on Domestic and Sexual Violence.

Her Majesty's Inspectorate of Constabulary (2014) *Everyone's business: Improving the police response to domestic abuse*. London: Her Majesty's Inspectorate of Constabulary.

Home Office (2013) New government domestic violence and abuse definition. *Correspondence: Circular 003/2013*, Home Office.

Hong, G. and Raudenbush, S. W. (2006) Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*, **101**, 901–910.

Hoppe, S. J., Zhang, Y., Hayes, B. E. and Bills, M. A. (2020) Mandatory arrest for domestic violence and repeat offending: A meta-analysis. *Aggression and Violent Behavior*, **53**, 101430.

Iyer, L., Mani, A., Mishra, P. and Topalova, P. (2012) The power of political voice: women's political representation and crime in india. *American Economic Journal: Applied Economics*, **4**, 165–93.

Johnson, S. D. and Bowers, K. J. (2004) The burglary as clue to the future: The beginnings of prospective hot-spotting. *European Journal of Criminology*, **1**, 237–255.

Kalair, K., Connaughton, C. and Alaimo Di Loro, P. (2021) A non-parametric Hawkes process model of primary and secondary accidents on a uk smart motorway. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **70**, 80–97.

Kavanaugh, G., Sviatschi, M. M. and Trako, I. (2019) Female officers, gender violence and human capital: Evidence from all-women's justice centers in peru. [Preprint] Available from: `https://rpds.princeton.edu/sites/rpds/files/sviatschi-female-officers-gender-violence-and-children-feb2019.pdf`. [Accessed 2021-02-09].

Levy, R. and Mattsson, M. (2020) The effects of social movements: Evidence from #metoo. [Preprint] Available from: `http://dx.doi.org/10.2139/ssrn.3496903`. [Accessed 2020-09-02].

Linderman, S. and Adams, R. (2014) Discovering latent network structure in point process data. In *Proceedings of the 31st International Conference on Machine Learning* (eds. E. P. Xing and T. Jebara), vol. 32 of *Proceedings of Machine Learning Research*, 1413–1421. Bejing, China: PMLR. URL: `https://proceedings.mlr.press/v32/linderman14.html`.

Loeffler, C. and Flaxman, S. (2018) Is gun violence contagious? a spatiotemporal test. *Journal of Quantitative Criminology*, **34**, 999–1017.

McDougal, L., Krumholz, S., Bhan, N., Bharadwaj, P. and Raj, A. (2018) Releasing the tide: how has a shock to the acceptability of gender-based sexual violence affected rape reporting to police in india? *Journal of Interpersonal Violence*. OnlineFirst.

Meyer, S., Elias, J. and Höhle, M. (2012) A space–time conditional intensity model for invasive meningococcal disease occurrence. *Biometrics*, **68**, 607–616.

Meyer, S., Warnke, I., Rössler, W. and Held, L. (2016) Model-based testing for space–time interaction using point processes: An application to psychiatric hospital admissions in an urban area. *Spatial and Spatio-temporal Epidemiology*, **17**, 15–25.

Miller, A. R. and Segal, C. (2019) Do female officers improve law enforcement quality? effects on crime reporting and domestic violence. *The Review of Economic Studies*, **86**, 2220–2247.

Mohler, G. (2014) Marked point process hotspot maps for homicide and gun crime prediction in chicago. *International Journal of Forecasting*, **30**, 491–497.

Mohler, G. O., Short, M. B., Brantingham, P. J., Schoenberg, F. P. and Tita, G. E. (2011) Self-exciting point process modeling of crime. *Journal of the American Statistical Association*, **106**, 100–108.

Nguyen, X. X. and Zessin, H. (1979) Integral and differential characterizations of the gibbs process. *Mathematische Nachrichten*, **88**, 105–115.

Nickerson, D. W. (2008) Is voting contagious? evidence from two field experiments. *The American Political Science Review*, **102**, 49–57.

Nicoletti, C., Salvanes, K. G. and Tominey, E. (2018) The family peer effect on mothers' labor supply. *American Economic Journal: Applied Economics*, **10**, 206–34.

Office for National Statistics (2015) National statistics—English indices of deprivation 2015 file 1. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/467764/File_1_ID_2015_Index_of_Multiple_Deprivation.xlsx. [Accessed 2020-12-15].

— (2017) *Crime Survey for England and Wales Technical Report 2018/19—Volume 1*. London: Kantar. Available from: https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/crimeandjustice/methodologies/crimeandjusticemethodology/201819csewtechnicalreportvolume1.pdf. [Accessed 2021-02-10].

— (2019a) Domestic abuse prevalence and trends, England and Wales: year ending March 2019. Available from: https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/domesticabuseprevalenceandtrendsenglandandwales/yearendingmarch2019. [Accessed 2021-02-10].

— (2019b) UK homelessness: 2005 to 2018. Available from: https://www.ons.gov.uk/peoplepopulationandcommunity/housing/articles/ukhomelessness/2005to2018\#reasons-for-homelessness. [Accessed 2021-02-09].

Ogata, Y. (1988) Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association*, **83**, 9–27.

O'Neal, E. N. (2019) "Victim is not credible": The influence of rape culture on police perceptions of sexual assault complainants. *Justice Quarterly*, **36**, 127–160.

Osborne, S., Lau, I. and Britton, A. (2012) *Homicides, Firearm Offences and Intimate Violence 2010/2011: Supplementary Volume 2 to Crime in England and Wales 2010/11*. London: Home Office.

Park, J., Schoenberg, F. P., Bertozzi, A. L. and Brantingham, P. J. (2021) Investigating clustering and violence interruption in gang-related violent crime data using spatial–temporal point processes with covariates. *Journal of the American Statistical Association*, **116**, 1674–1687.

Pate, A., Hamilton, E. E. and Annan, S. O. (1992) *Metro-Dade Spouse Abuse Replication Project: Technical Report*. Washington: Police Foundation.

Ratcliffe, J. H. and Rengert, G. F. (2008) Near-repeat patterns in philadelphia shootings. *Security Journal*, **21**, 58–76.

Reinhart, A. (2018) A review of self-exciting spatio-temporal point processes and their applications. *Statistical Science*, **33**, 299–318.

Reinhart, A. and Greenhouse, J. (2018) Self-exciting point processes with spatial covariates: modeling the dynamics of crime. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **67**, 1305–1329.

Rincke, J. and Traxler, C. (2011) Enforcement spillovers. *Review of Economics and Statistics*, **93**, 1224–1234.

Robinson, A. L., Pinchevsky, G. M. and Guthrie, J. A. (2018) A small constellation: risk factors informing police perceptions of domestic abuse. *Policing and Society*, **28**, 189–204.

Rogers, M. et al. (2016) Barriers to help-seeking: older women's experiences of domestic violence and abuse - Briefing note. Available from: https://usir.salford.ac.uk/id/eprint/41328/1/Barriers\%20to\%20older\%20women\%27s\%20help-seeking\%20Briefing\%20note\%2014\%207\%202016.pdf. [Accessed 2021-02-10].

Rotton, J. and Cohn, E. (2001) Temperature, routine activities, and domestic violence: A reanalysis. *Violence and Victims*, **16**, 203–15.

SafeLives (2015) *Insights Idva National Dataset 2013-14*. Bristol: SafeLives. Available from: https://safelives.org.uk/sites/default/files/resources/Insights\%20Idva\%20national\%20dataset\%202013-2014.pdf. [Accessed 2021-02-10].

Sampson, R. J. and Groves, W. B. (1989) Community structure and crime: Testing social-disorganization theory. *American Journal of Sociology*, **94**, 774–802.

Sanders, A. N., Kuhns, J. B. and Blevins, K. R. (2017) Exploring and understanding differences between deliberate and impulsive male and female burglars. *Crime & Delinquency*, **63**, 1547–1571.

Sanders-McDonagh, E., Neville, L. and Nolas, S.-M. (2016) From pillar to post: understanding the victimisation of women and children who experience domestic violence in an age of austerity. *Feminist review*, **112**, 60–76.

Schoenberg, F. P. (2002) On rescaled Poisson processes and the Brownian bridge. *Annals of the Institute of Statistical Mathematics*, **54**, 445–457.

Schuster, E. F. (1985) Incorporating support constraints into nonparametric estimators of densities. *Communications in Statistics—Theory and Methods*, **14**, 1123–1136.

Shaw, C. R. and McKay, H. D. (1942) *Juvenile Delinquency and Urban Areas : A Study of Rates of Delinquents in Relation to Differential Characteristics of Local Communities in American Cities*. Chicago, Ill. : The University of Chicago Press.

Short, M. B., D'orsogna, M. R., Brantingham, P. J. and Tita, G. E. (2009) Measuring and modeling repeat and near-repeat burglary effects. *Journal of Quantitative Criminology*, **25**, 325–339.

Sobel, M. E. (2006) What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *Journal of the American Statistical Association*, **101**, 1398–1407.

Stark, L., Warner, A., Lehmann, H., Boothby, N. and Ager, A. (2013) Measuring the incidence and reporting of violence against women and girls in liberia using the 'neighborhood method'. *Conflict and Health*, **7**, 20.

Strube, M. J. (1988) The decision to leave an abusive relationship: empirical evidence and theoretical issues. *Psychological Bulletin*, **104**, 236.

Taylor, S. C. and Gassner, L. (2010) Stemming the flow: challenges for policing adult sexual assault with regard to attrition rates and under-reporting of sexual offences. *Police Practice and Research: An International Journal*, **11**, 240–255.

Tjaden, P. G. and Thoennes, N. (2000) *Extent, nature, and consequences of intimate partner violence: Findings from the National Violence Against Women Survey*. Washington: National Institute of Justice.

Trute, B., Adkins, E. and MacDonald, G. (1992) Professional attitudes regarding the sexual abuse of children: Comparing police, child welfare and community mental health. *Child Abuse & Neglect*, **16**, 359–368.

van Lieshout, M.-C. (2010) Spatial point process theory. In *Handbook of Spatial Statistics* (eds. S. Frühwirth-Schnatter, G. Celeux and C. P. Robert). Boca Raton, Florida: CRC Press.

Walby, S. and Allen, J. (2004) Domestic violence, sexual assault and stalking: Findings from the British Crime Survey. *Home Office Research Study 276*, Home Office Research, Development and Statistics Directorate.

Walby, S. and Towers, J. (2012) Measuring the impact of cuts in public expenditure on the provision of services to prevent violence against women and girls. Report for Northern Rock Foundation and Trust for London.

Weisburd, D. (2015) The law of crime concentration and the criminology of place. *Criminology*, **53**, 133–157.

Women's Aid (2009) *The Survivor's Handbook*. Bristol: Women's Aid. Available from: `https://www.womensaid.org.uk/wp-content/uploads/2016/05/Full-Survivors-Handbook-English-2009.pdf`. [Accessed 2021-02-10].

Zhuang, J. (2011) Next-day earthquake forecasts for the japan region generated by the etas model. *Earth, Planets and Space*, **63**, 207–216.

Zhuang, J. and Mateu, J. (2019) A semiparametric spatiotemporal Hawkes-type point process model with periodic background for crime data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **182**, 919–942.

Zhuang, J., Ogata, Y. and Vere-Jones, D. (2002) Stochastic declustering of space-time earthquake occurrences. *Journal of the American Statistical Association*, **97**, 369–380.