Assault Crime Dynamic Chain Event Graphs

Jim Q. Smith[1] and Aditi Shenvi[2]

[1]Department of Statistics, The University of Warwick and The Alan Turing Institute
[2]Centre for Complexity Science, The University of Warwick

June 12, 2018

Abstract

In this paper a new class of graphical Bayesian discrete semi-Markov models is developed to describe the various pathways that might lead someone into perpetrating various kinds of crime involving assault and violence. Our discrete probability models are crafted to embody various theory and empirical studies by psychologists and sociologists explaining and describing this development. This probability model is then used to formally structure a new decision support system to help authorities evaluate public risks guided by not only archived but real time data. We argue here that such systems will be able to provide provisional quantified evaluations of the impacts of various policy and policing decisions into the short and medium term. The construction of such probability models is illustrated throughout by examples. We end the paper with a more detailed description of a model built to support authorities to frustrate populations of criminals who have been radicalised into violent extremism.

Acknowledgement We are indebted to Rob Procter for pointing us towards critical related references in Machine Learning.

1 Introduction

In this paper we develop a new probabilistic framework for describing how a subpopulation Ω_t of the general public might be drawn into an assault crime. This framework enables us to leverage recent advances in forensic science, especially in the evaluation of the strength of activity level evidence see e.g. [4, 19, 22, 45] and reapply these to this domain. Our models describe how different sequences of generic *events* - either unplanned or as the result of life decisions - encourage an individual $\omega \in \Omega_t$ into or discourage them away from these criminal activities.

Because of the unbalanced nature of the evidence in this domain it is necessary to embed structured information in the form of sociopsychological theory.

These theories view the adoption of a life of crime as a process. Each individual responds in a rational way to their environment and is driven by their needs and aspirations. An extensive library of information that has been collected about the life history of different classes of violent criminals informs these studies. These need to be embedded into a Bayesian statistical modelling framework to underlie any formal tool developed to support the evaluation of policy options to prevent as many people as possible from getting involved in violent crime. This statistical framework needs on the one hand to be sufficiently detailed to be able to capture the wide variety of possible developments in the lives of criminals that make them dangerous and be able to distinguish these from the various innocent developments in the lives of the vast majority of people. On the other hand it must be sufficiently coarse and generic to classify those diverse subpopulations of individuals who might in the future pose a threat to public safety. Then standard Bayesian decision analytics [56] can be used to formally evaluate evidence in this domain and to support operational and resource deployment decisions in this area. This paper describes and illustrates how this challenge can be met.

We argue here that expert judgments about criminal population dynamics are well described through event trees and dynamic time inhomogeneous semi-Markov processes customised to this domain. To do this it is necessary to further develop the graphical technologies of Chain Event Graphs (CEGs), and especially their dynamic variants [8, 19, 21].

Standard dynamic extensions of the now established CEG methodology have been Markov rather than semi-Markov [21]. Except for a short discussion of technical apparatus in [8] the current development of these models is in its infancy. This is the first paper to explore the use of this technology to the type of domain we face here. Semi-Markov models rather than Markov models are needed because social and psychological models of criminal behaviour are usually articulated in terms of ordered sequences of events within an unfolding process. Within such descriptions - where evidence from different case studies is drawn together - the time it takes to move between adjacent states in the unfolding process tends to vary from subject to subject within a study and can be dependent on their environment - often invisible to an observer. On the other hand, an observer's probability that an individual makes his *next transition* from one situation to another can often be assumed to be more stable across the population. It therefore makes sense to structure expert judgments around a direct description of the list of transitions a person might make when in a given state. At least within the class of problems we have examined this structural information appears less contentious and links directly to the natural language explanations various criminologists use to express their theories.

Once a framework (or competing frameworks) of this kind have been specified, the next delicate task is to somehow harness relevant information and expert judgments to specify a distribution on the probability vector of a particular type of individual transitioning into an adjacent state. These probabilities are clearly highly dependent not only on the population being examined but the type of assault being considered. In this paper we illustrate various different ways this can be achieved through a sequence of examples, before developing certain templates for general use. Within the developed class of models, the final component concerning the length of time it might take before an individual makes such a transition is typically added only at the end of analysis when and if this most fragile component of the analysis is needed.

One critical challenge is that some of the individuals within the general public will be well known to the authorities whilst others may be totally invisible. For known individuals we are able to perform a person centred Bayesian analysis. However for the others we have to rely on a population study where population statistics are uncertainly inferred from the life histories of similar individuals and general criminological hypotheses, and these situations often need to be handled differently [38].

The next section motivates this work through two illustrative person centred Bayesian analyses of different assault crimes. We illustrate how models of individual crimes, previously developed estimation and model selection methodologies can be straightforwardly modified and translated into this domain and demonstrate the use of a new class of models called the Reduced Dynamic Chain Event Graphs (RDCEGs). Drawing on a selection of sociological and psychological studies in Section 3 we proceed to define a generic framework around which such crimes can be discussed and formalised into a probability model based on a new family of Dynamic CEGs (DCEGs). This involves us building a number of random variables that indicate on an ordinal scale several components loosely related to the means, motivation and opportunity of an individual criminal. In Section 4 some of the inferential challenges associated with this model class are explored.

Leaning on some current sociological insights, in Section 5 we then focus down the study to a specific domain that lies within these broader classes. More detailed RDCEG frameworks are built for a specific genre of assault crime: Radicalisation leading to Violent Extremism (RVE). We show how the graphical framework enriches the expression of early naive escalation models used for this domain whilst being sufficiently structured to admit a later embellishment into a full probabilistic class. We end the paper by illustrating these constructive frameworks with reflections on a real study. We outline how to perform this probabilistic embellishment to a full class of Bayesian models needed to inform a particular decision support system: one designed to evaluate different policies designed to discourage a progression into RVE. In particular we illustrate how to seamlessly transition between population models and generic person centred models within this class. The paper ends by describing how this work is currently being applied to build new support systems needed to pursue violent criminals.

The standard depictions of this family of DCEGs, whilst still applicable to this domain, tend to become rather cluttered. So in the next section we begin this study by defining a simple new set of bespoke representational rules.

2 The activity graph of a criminal

2.1 The RDCEG to model a single criminal

We begin by reviewing the CEG and the established extensions of this statistical family of models. We then develop these methodologies so that it can be applied to describe the evolution of criminal behaviour. Chain Event Graphs [6, 19, 57] and their dynamic extensions [8, 21, 32] have now been established to provide a simple but powerful representation that can be used to describe a semi-Markov process. The class of models is based on an event tree. In its dynamic form this is an infinite tree whose vertices and edges are coloured to represent various types of context-specific conditional independence assertions. These can not only form the framework for principled personal models of a single individual but also of a model of a population of potential criminals. For a detailed description of this particular aspect of this technology for the standard CEG see Chapter 5 of [19].

Within our context the potential development of each person within the population Ω_t at time t of people who might be drawn into assault crime at some point in the future is described by a path along the edges of this tree beginning at its root. The vertices of this tree express various discretely classified *situations* where each individual might lie at any given time. Elicited criminological models are often concerned with how someone at a particular situation in their life might transition into either a new situation of escalated threat or into a more benign state. These positions - defining a set of comparable situations - then form the states of the associated semi-Markov process of the crime. Within the DCEG technology and semantics two situations or vertices of the event tree and their emanating edges in the tree share a colouring if the set of transition probabilities associated with colour identified edges are hypothesised to be the same.

The event tree so coloured is called a *staged tree*. In the cases we consider here this tree can be of infinite depth. One finite graph that describes this infinite tree is called a DCEG. This means that the entirely formal features of the hypothesised discrete stochastic model specifically its underlying semi-Markov structure and additional conditional independence assertions can be expressed as a finite coloured directed graph which transparently and evocatively represents the underlying criminal development. The construction of a DCEG is formally and extensively described in [8, 20, 21, 32]. The vertices of a DCEG are called its *positions*. These are sets of comparable situations in the original infinite depth staged tree. When two situations lie in the same position we hypothesise that the probability law describing the future unfolding of events for a person arriving at either of these situations can be identified.

This graph has already proved to be a very useful framework for faithfully eliciting hypotheses concerning how things might happen. This framework can then, after close discussions with domain experts be examined and adapted until the expert is satisfied that the framework expresses a plausible hypothesis. This discussion can be conducted using only natural language and so does not disempower the domain expert - here the expert in the developmental stages describes the development of a person into being prepared to enact a certain type of assault crime. Only once this framework is in place is a probability model built around it.

Thus here we demonstrate how the graph of a particular variant of a DCEG, the RDCEG can provide not only a transparent description of the possible evolutions of someone in the given population - a collection of unfolding histories which take them from one broad category to another as we move along the paths of the graph - but also one which admits a probabilistic embellishment that can then support a quantification of the degree of risk presented by a particular individual. The vertices of the graph correspond to a critical point in the evolution of someone within the population whilst hypotheses about such transitions are expressed *directly* by the edges of the graph between a finite number of broadly defined vertices. For dynamic models with homogeneous transitions this diagram turns out to be a simple but powerful embellishment of the familiar state transition diagram of a finite Markov process. Here are two very simple examples of how this process can be used to describe two very different sorts of assault crime.

2.1.1 A non dynamic CEG directly applied to a criminal process

The first example is of the simplest possible case where the repeated acts of a single identified potential perpetrator of a crime can - at least in the first instance - be legitimately expressed in terms of a finite tree. This enables us to use as our framework a vanilla CEG [19] with no adaptation. The simplicity of this example allows us to transfer directly Bayesian methods for forensic inference concerning activity level evidence - see [4, 22, 26, 27, 56].

Example 1 *A man is suspected of child entrapment. He is accused of initiating a sequence of long electronic chat sessions with a young girl. Over the 50 day period he was observed, 18 of these days he engaged in an electronic chat lasting over an hour with a particular minor. Police were suspicious that these calls were designed to groom this child into an eventual meet. The suspect claims to the contrary that the conversations which actually initiated by the girl, were entirely innocent and that he was simply trying to humour her. To investigate the issue further, police have obtained from the suspect's and also his wife's employer a record of which of those 50 days each of the couple had been away on business. The unfolding events on each day can be represented using the staged tree given in Fig. 1.*

In such cases the path followed on any one day is uncertain. Here, for example, contacts by either party may not lead to a chat. One party may not be able to chat for some unrecorded reason and just because someone was recorded as being away or not does not necessarily imply that a chat was or was not possible. However, given their suspicions, police would expect that proportionately more chats would occur when the suspected man and his wife were recorded as being together - i.e. on days leading to situations given in green (s_4, s_5, s_6) rather than for the situation given in pink (s_3) in the staged tree in Fig. 1. The CEG in Fig. 2 is based on the police suspicions that the node w_1 which constitutes the

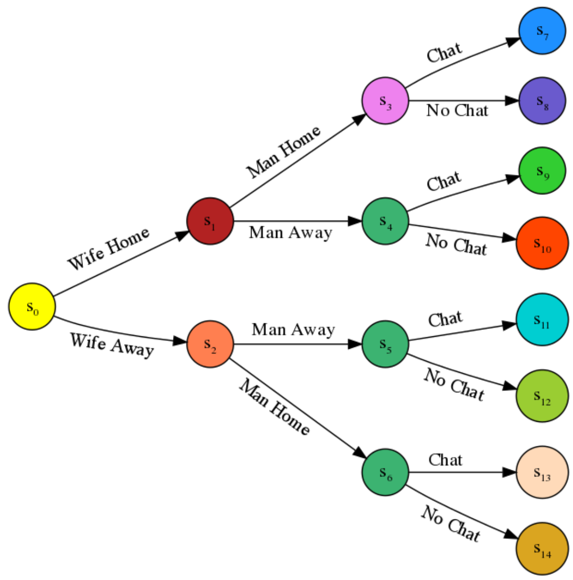

Figure 1: Staged tree for example 1.

situations s_4, s_5, s_6 lies in a different stage than the node w_2 which constitutes the situation s_3 indicating that the probability of a chat happening would be much lower, the suspect being frustrated from making contact by the presence of his wife. However - unless the girl has some way of divining when the suspect would be remote from his wife - under the suspect's explanation of events situations s_3, s_4, s_5 and s_6 would lie in the same stage as shown in Fig. 3. So we have two competing probability models shown by the two CEGs in Fig. 2 and Fig. 3 that provide competing hypothesised explanations of events. Note that, following good practice [19] the CEGs below we have only included events whose distribution might be different under the competing hypotheses. So for example we have drawn together into a single vertex all the ways the wife and suspect could be separated when the suspect might be free to contact the girl. This is possible because under both competing hypotheses these different scenarios appears to have no discriminatory information about the innocence or guilt of the suspect.

Note that both hypotheses - expressible in natural language - are represented accurately and succinctly by the two graphs. However both graphs also formally express the structures of the probability models underpinning the two hypotheses. Collected data can be used both to estimate the edge probabilities on the critical events under either hypothesis. These help to guide the calculation of the extent to which one hypothesis is more plausible than the other: e.g. in this simple case a hypothesis test on a 2×2 table or a standard Bayes Factor score.

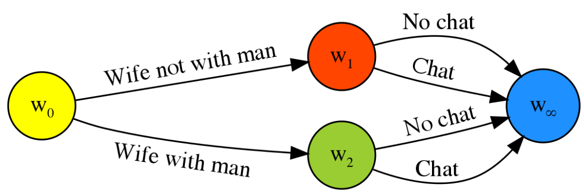

Figure 2: CEG for example 1 based on police suspicions.

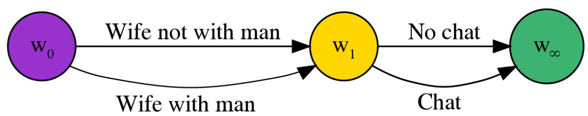

Figure 3: CEG for example 1 based on the suspect's claims.

Assume for example that the only relevant or admissible data available is given in the units of days the man and wife are together or separated on that day and whether or not a long chat happened expressed via the contingency table below. The hypothesis of independence would then be consistent with the hypothesis that child was initiating the contacts whilst the (one sided) alternative of the positive association between a chat taking place and the separation of the man and his wife would be consistent with the police hypothesis.

	No Chat	Chat	
Couple Together	25	2	27
Couple Separated	7	16	23
	32	18	50

Applying Bayes Factor model selection methods described in detail in Chapters 5 & 6 of [19] with a default uniform prior on the leaves of the staged tree above and assuming random sampling over the 50 days the marginal likelihood ratio of this activity data is:

$$l \triangleq \frac{P(\text{Data}|\text{Suspect's Innocence})}{P(\text{Data}|\text{Suspect's Guilt})} = \frac{\Gamma(29)\Gamma(25)\Gamma(34)\Gamma(4)\Gamma(20)}{\Gamma(26)\Gamma(3)\Gamma(8)\Gamma(17)\Gamma(54)} = 0.000085729$$

For the data above this would give what Aitken [4] would advise that there was strong evidence in the support of the police suspicions. An elementary application of Bayes rule gives us that if p_0 is a juror's prior probability of the suspect's guilt then

$$\frac{P(\text{Suspect's Innocence}|\text{ Data})}{P(\text{ Suspect's Guilt}|\text{ Data})} = l\frac{p_0}{1-p_0}$$

It can be checked that even by giving the suspect the benefit of the doubt, any formal analysis of the data above provides strong evidence in the favour of the police version of the process for most rational jurors. This illustrates how data can be formally incorporated into an analysis to discriminate between innocent interactions and guilty ones and the sorts of arguments applied by forensic scientists in preparing their submission to court about strength of evidence. For the decision support we envisage here - where the primary aim is not to convict but to determine whether or not resources should be deployed to frustrate the activities of a particular suspect - such analyses are even more clear cut.

This model could of course be embellished into a more nuanced CEG. For example we could choose the parameters of the beta prior to reflect information that might be available concerning the number of days the players would normally be away from home. More subtly the police may conjecture that the couple would try to be together as much as possible in normal circumstances leading them to be together more than randomly which is what would be expected if the suspect were innocent. On the other hand if he were grooming the girl he may be trying to optimise his time away from his wife. So separation would then be more probable than by chance. The scope for choice would need to be informed by experts and so any analysis would be very context-specific. But all these embellishments could be explored and populated with hard evidence if this was thought necessary. Notice that all these embellishments can be specified as answers to questions asked in natural language - as can its outputs [56]. It is only when it is necessary to determine the *strength* of evidence for and against contending hypotheses that any numerical elicitation or analyses need to be performed. This is what makes this technology so useful to domains like the ones we discuss here.

The point of this type of analysis is two fold. First it guides the systematic collection of data. If the man is eventually arrested then natural language translations of this analysis can form the basis of a prosecution case in court. This application is now becoming widely used in the presentation of forensic evidence in courts of law: albeit usually using the framework of the BN. Perhaps even more useful is to apply this sort of analysis to help determine how to best allocate resources to prevent such unfolding events culminating in an actual assault. In the context above where there are a number of suspicious engagements being observed the framework can be used to support police prioritise the pursuit of collections of different suspects to maximise the number of potential perpetrators they are able to frustrate. Generalisations of this latter type of use will be the focus of this paper.

2.1.2 Extending the crime to model with the RDCEG

A typical criminal $\omega \in \Omega_t$ may at some point drops out of a life of crime - for example because of changing circumstances, repentance, arrest or death. In fast moving single sequences of engagements like the one above this absorbing state has little impact on inferences. However in other examples it will have. Furthermore the man in the example above may have the opportunity to repeat his suspicious behaviour patterns with different victims indefinitely often. When this is a significant possibility an alternative dynamic framework needs to be used see e.g. [8, 20, 21, 32].

In this paper it is convenient to modify the established graphical semantics because we can assume an absorbing state - called here the *immune* state i - always exists and can be transitioned to from many states in these criminal processes. The new graphical representation does not depict i explicitly because by definition edges can only point into this state. A *reduced DCEG (RDCEG)* can therefore be constructed from the subgraph of a DCEG obtained by omitting the immune vertex and all edges into it in the DCEG. The RDCEG therefore has one less vertex and far fewer edges than the standard graph of a DCEG. This simple device turns out to be very useful in this criminal setting because the depiction of the processes is much less cluttered. It also enjoys some pleasing formal properties [52]. Note that when there is always a small probability that a person becomes immune then the full graph of a DCEG can always be reconstructed from the RDCEG. In distinction to the DCEG, when we add edge probabilities to quantify the description these probabilities emanating from a particular vertex these often sum to a value strictly less than one. This loss of probability weight then implicitly quantifies the probability that at their next transition a suspect will cease this particular criminal activity for ever, i.e. will transition into i, the immune state.

Example 2 *Typically in cases like the ones above the entrapment will end at some point. The child may end the communication, the man gives up or the child and the man meet. In any of these circumstances, the man may repent or be arrested - transitioning into an immune state - or he may begin to explore a new contact. This process will continue until such time as the man reaches i, where by definition he will stay. This typical scenario can model such extended hypotheses using an RDCEG. The simplest such model is given in Fig. 4.*

When the other party to a chat with the suspect is known to be a child then under the model hypothesis depicted above leads us to construct a number of separate records of the type collected in the first example. Under the simplest hypothesis that the innocence or otherwise of the chat does not depend on the child, these can be aggregated into a common contingency table just like the one given above. Such evidence can then be collected until the current time or the absorbing state is reached and used to further strengthen the case against the suspect. In this way expected probabilities can be added to each of the directed edges corresponding to the transition probabilities of an associated semi-Markov process. So now well developed inferential methodologies [19] then transfer seamlessly into this new framework: see [52] for a more explicit description of this

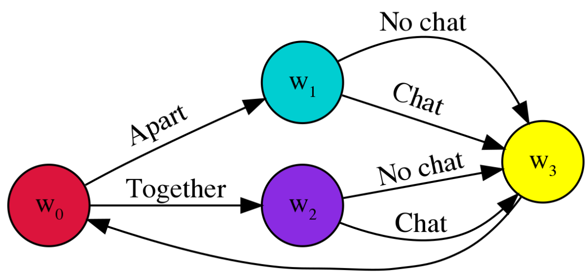

Figure 4: RDCEG for example 2. After reaching position w_3, the unit traverses back to the position w_0 under the assumption that it does not enter the immune state i in the next transition.

transfer.

The next example is of a very different assault crime and illustrates a more nuanced dynamic. To keep things simple we again begin the modelling process just before a potential crime threatens to be committed and where there is a single potential perpetrator of the crime.

Example 3 *Police have strong evidence that suggests that woman A has decided to take revenge on man B by shooting him. At any given time she will be trained to shoot (s) or not ($\overline{s}$) and will own a gun (g) or not $\overline{g}$. She can take various courses of action. She could choose to try to learn to shoot now - if she currently does not know how to - in a condition where she currently owns a gun $t_s(\overline{s}, g)$ or not $t_s(\overline{s}, \overline{g})$. Alternatively she could immediately try to acquire a gun at a time when she has either trained to shoot ($t_g(s)$) or not ($t_g(\overline{s})$). At any point in this process she may enter the immune state i: for example B may be killed by someone else or A may have a change of heart or be arrested and deported. Note that A can only perpetrate the crime when in state (s, g). A possible model of this process is depicted by the RDCEG in Fig. 5.*

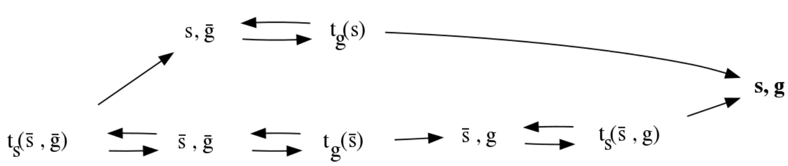

Figure 5: RDCEG for example 3.

Its vertices represent all the possible states in the process other than i. Note that this RDCEG is simply the state transition graph of a semi-Markov process *defining the dynamic with the absorbing state i and all edges into it deleted. When ordering the states - called its* positions - *of this Markov process as*

$$(t_s(\overline{s},\overline{g}),(s,\overline{g}),(\overline{s},\overline{g}),t_g(s),t_g(\overline{s}),(\overline{s},g),t_s(\overline{s},g),(\boldsymbol{s},\boldsymbol{g}),i)$$

the configuration of zeroes in the transition matrix of the corresponding semi-Markov process is given below

$$\begin{bmatrix}
0 & * & * & 0 & 0 & 0 & 0 & 0 & \star \\
0 & 0 & 0 & * & 0 & 0 & 0 & 0 & \star \\
* & 0 & 0 & 0 & * & 0 & 0 & 0 & \star \\
0 & * & 0 & 0 & 0 & 0 & 0 & * & \star \\
0 & 0 & * & 0 & 0 & * & 0 & 0 & \star \\
0 & 0 & 0 & 0 & * & 0 & 0 & 0 & \star \\
0 & 0 & 0 & 0 & * & 0 & 0 & * & \star \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \star \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \star
\end{bmatrix}$$

. So the uncoloured graph of an RDCEG given above directly depicts the structure of the transition matrix implicit in our natural language description of the process. Note that the sum Σ_k of the probabilities in the $*$ entries on row k of this matrix will be no larger than 1 with $1-\Sigma_k \geq 0$ being the probability A enters the absorbing state i from the k^{th} position shown by $\star$ in row k. A full probability model of the process - always consistent with the natural language description encoded by the RDCEG - is then provided by simply adding probabilities to the starred entries along the rows of the structured transition matrix above. This is a device we use throughout this paper to translate criminological theory and instances into a probability model.

There are various important points to make about this class. First although the graph of the RDCEG is evocative and expresses a hypothesised model expressed in natural language fully and precisely its implicit Markov assertions are nevertheless very substantive. In particular the states = positions define the parts of the history that need to be remembered to forecast the future. A critical issue here is to get these states right. Technically a random variable which takes different values at each state at any given time has been assumed to be predictively sufficient. So in the example above, how A arrived at any of the positions has no relevance to the subsequent development once she gets there.

If for example on the contrary it is believed that if A tried to shoot twice and failed then she would give up for ever then the state space given above would not be valid. The set of positions would then need to be increased and redefined so that they remembered whether or not the suspect had tried twice or not giving a different RDCEG representing a competing hypothesis. Note the benefits of one hypothesis over the other can be discussed in *natural language*. So - just as for the BN - these graphs can be used *directly* to analyse the logical and qualitative predictive consequences of such hypotheses. The type of dialogue between expert and analyst is intrinsic to the process of structural decision analysis and the examination of competing hypotheses central to the elicitation processes discussed in this paper. A technological development to do this for a CEG, is now well established. The RDCEG can be used in exactly the same way and interrogated without having to first elicit any probabilities: see e.g. [8, 19, 58] until a structurally requisite model is discovered [56].

Once an RDCEG has been elicited and embellished with transition probabilities, with complete random sampling on the transition events alone data can be used within a conjugate analysis for fast estimation and model selection just as in the entrapment case above. The formulae for these are given in an appendix and discussed in more detail in [52].

Note that the RDCEG is much more than a state space diagram of a semi-Markov process because the *colouring* of its vertices and edges is able to convey critical additional conditional independence information about the process. This type of additional information is often intrinsic to the domain specific models we describe here. Fortunately it is precisely because of this property that the RDCEG inherits most of the technologies originally developed for DBNs. On the other hand one advantage the RDCEG has over BN technologies is that it explicitly depicts zero transitions into all states other than the immune state. Within crime models such zero transitions typically form a central part of the unfolding story.

For the purposes if this paper it is only necessary to illustrate this colouring - a formal detailed exposition of this process can be found in [8, 19, 20, 52]. Depending on our construction protocol these colours are either inherited from an elicited staged tree or added to positions after the graph has been elicited. In the example above suppose the criminologist hypothesises that the two probabilities $P(t_s(\overline{s},\overline{g}) \to s,\overline{g})$ of A learning to shoot when not owning a gun and $P(t_s(\overline{s},g) \to s,g)$ when owning a gun are the same as are $P(t_s(\overline{s},\overline{g}) \to \overline{s},\overline{g})$ of not learning to shoot with no gun and $P(t_s(\overline{s},g) \to \overline{s},g)$ when owning a gun. This qualitative information can be depicted on the RDCEG simply by assigning the same colour to positions $t_s(\overline{s},\overline{g})$, $t_s(\overline{s},g)$, to the edges $t_s(\overline{s},\overline{g}) \to s,\overline{g}$ and $t_s(\overline{s},g) \to s,g$ and to the edges $t_s(\overline{s},\overline{g}) \to \overline{s},\overline{g}$ and $t_s(\overline{s},g) \to \overline{s},g$ [8, 19]. Within DBN technologies such local structure is sometimes referred to as a "context-specific conditional independence" assumption or as defining an "object" [41].

Ideally we would like to construct default RDCEG templates of collections of different criminal process indexed by type of criminal, and type of crime, that build on existing criminological models. To build such a library takes some effort but is quite possible. This type of technology has for example already well developed for BNs within the context of forensic science, see [4, 27] and established frameworks of processes linking activities with evidence. In this paper we illustrate how some of these RDCEG templates can be built to help model crimes associated with assaults or violence against the general public.

A template RDCEG associated with a single criminal, particular crime and demographic may have many more positions than the ones illustrated above and defined more generically so that information across different past cases can be drawn together. In the next section we demonstrate how this can be done.

3 Generic features of a criminal RDCEG

3.1 Means, Motive, Opportunity

An old legal check list for examining the culpability of a suspect is given by their "means, motive and opportunity". Traditionally all of these should be present if a prosecutor has a chance of convincing a jury of a suspect's guilt in a court of law. Here we will use this simple aide memoir to help template some of the structural features of general RDCEGs associated with violent or assault crimes. We then illustrate how these templates can be refined to model radicalisation leading to violent extremism.

First in such crimes someone needs to be *motivated* to enact an assault. In our first example the issue of the man's motivation under the two different hypotheses was implicit for both versions of the recent history. Under the police hypothesis he was motivated in the second he was not. In the gun example we had assumed that the potential killer was already motivated. In practice the motivation of a person however may be much more nuanced. If the aim is to prevent an individual being drawn into crime motivational escalation will be of primary interest.

Second such a person needs to have the *means* or capability to carry out an attack. In the first example he simply needed a laptop or electronic devise and know how to use this in an appropriate way. But in the second we saw that she needed the training (to shoot) and tools (a gun) to perpetrate the act. In more complex crimes an assailant often needs to embed themselves in a gang or at least be in intimate contact with like-minded criminals. For example if someone chooses to act as a mule in a bombing incident then this will usually need the potential perpetrator to embed themselves in a gang with the skills to make a bomb. Such a gang will often be made up of different, often diversely motivated and skilled criminals.

The final component of the process is the suspect has the *opportunity.* This feature is much more immediate and critical to police operations but less so for example for social services. It is again convenient to split this attribute into two components. The first relates to the victims. In the entrapment example, the suspect need to find someone to trap. In the second he must have access to his target. The bomber - possibly through his gang - must have identified a vulnerable place to perpetrate the incident. The second component associates to the assailant. In the first the man's wife or some other potential observer cannot be present or he will be frustrated. In the second the attacker needs to currently own a gun and be somewhere near the planned victim. For the bomber the materials to make the bomb acquired, the bomb needs to be built and then transported to the target. It is often convenient to think of opportunity only to be relevant if motive and means are already in place. The frustration of opportunity is either directed at general threats and systematically enacted, for example the introduction of new laws, codes of practice and protocols that better protect the general public from such an assailant. This is often termed a *Protect* policy. Alternatively it could be frustrated through the active *Pursuit* of potential perpetrating individuals or gangs of criminals. The examples above illustrate these points well.

Only with all three elements in place at least to a sufficiently high degree can a crime be perpetrated. Furthermore the impact of the crime also depends on the levels of all three of these components. So again we have quite a complicated set of mutually exciting states which determine the severity of the threat posed by a particular individual. Here we first construct escalating scales of threat guided by elaborations of the three principles above. The complexity of these scales needs to trade off the specificity of the template to enable critical evidence informing judgments in a given instance to be accommodated against the extreme variety of incidents and actors that make it sufficiently generic to be useful. We illustrate the simplest generic RDCEG templates of two types of assault crime and describe how these can be embellished and customised to a particular real scenario. We then discuss how this structure can be populated with a formal probability model to accommodate expert judgments and any available sample information. We end by discussing and illustrating how this analysis can be used for decision support - especially about resource allocation by combining information in these distributions with the utility of a user to score various portfolios of policies. This helps assess the effectiveness of various resource deployment strategies by providing tentative scores to various portfolios of future remedial actions for domain scrutiny, adaptation and adoption.

Of course no one single template will fit all types of assault crime. However by focusing on states that capture motivation, means and opportunity we find that surprisingly few generic templates are needed to describe a wide class of different processes and to help inform the types of decision making discussed above. In the remainder of this paper for simplicity we will focus mainly on the first two attributes.

3.2 Escalation into assaults on the general public

3.2.1 Motivation

There is now a vast corpus within both the psychological and sociological literature describing a vast variety of theories about how someone can be drawn into a life of violent crime: see for example, [11, 15] and references therein. A very brief review of the literature associated with one sort of violent crime is given below. At the most basic level there are two critical attributes of such people: their preparedness to engage in crime and their preparedness to assault another person. Obviously both need to be in place before the criminal is motivated to perpetrate an assault.

Two escalating scales classifying someone's self perception in relationship to general society are given below. These components are labelled *alienation* and *violence* and measure of someone's preparedness to sympathise with or engage in assault of a fellow human being [59]. Here, for convenience, these scales are described in terms of answers to an imaginary questionnaire a suspect might be asked to use to best describe their attitude. This scale has been informed by similar escalation lists within Social Movement Theory (SMT) [66] and its application to a particular criminal group [14, 28] now translated in a way that enables the scale that can be readily transformed into an event space of a probability model. In particular the alienation scale below expands the escalating SMT scale given in [65] who label these states of mind "openness to new world views" - a_3, a_4, "religious/political seeking" - a_5, "frame alignment" - a_6 and "indoctrinated" - a_7, a_8.

3.2.2 Alienation

1. I am content with my life and its trajectory and the normal evolution of mainstream society, a_1.
2. I am dissatisfied with my current life and am currently looking for better opportunities within it, a_2.
3. I am disaffected with my life chances; there seems to be no future for me, a_3.
4. Society must change; it cannot currently support right thinking people like me, I am an alien and need to align myself to a group of other outsiders a_4.
5. Mainstream society could never change appropriately within usual democratic processes. Revolutionary change is needed to transform it so that people with whom I am aligned can participate, a_5.
6. Current mainstream society - its agents and its participants - is my enemy; its processes and people need to be actively undermined, a_6.
7. Any acts - criminal or not - that attack and undermine the current order are fully justified; I wholeheartedly support those with whom I am aligned who are prepared to take such acts, a_7.
8. Those who do not align with my perspectives are to be despised. My personal vocation is to align everyone to my world view using all means possible within my own moral compass, a_8.

3.2.3 Violence

1. I would never knowingly physically harm someone or support anyone else doing this, v_1.
2. I could envisage circumstances when I could support someone assaulting someone else but I could never perpetrate this act myself, v_2.
3. In extreme circumstances I might be prepared to assault someone myself, v_3.

4. Whenever necessary I will reluctantly personally assault another person, v_4.
5. I am quite happy to physically assault someone if necessary, v_5.
6. I am eager to assault someone if given the opportunity, v_6.

In practice a suspect could not be assumed to complete these two questionnaires truthfully if really asked! However it is still useful to refect on what someone might report if they were honest. This then provides an event that passes the Howard's Clarity Test [37, 56] that can be then legitimately assigned a probability by someone else. (The suspect if asked and motivated to be honest could provide their preferred answer at any given time.)

An outsider will often not usually be able to supply the suspect's answer to these questions, unequivocally. However government should be able to specify its joint probability distribution on these two attributes and so calculate its probability distribution on the motivation scale for criminal assault discussed below. In fact it often turns out that someone's observed behaviour is strongly informative about where they are positioned on both of these scales: a suspect's own social network communications might even volunteer this information!

Even for an unknown person within the population some data will inform a probability distribution defined over similar individuals. For example people within the population who have been out of work or are underemployed are more likely to feel alienated from society. Similarly anyone recently having visited a war zone is much more likely to be unusually acculturated to violence and assault. Note that for both lists most combinations of pairs of states in these escalating gradation a person is no threat whatever to the general public. Indeed many religious people would identify themselves at the point a_4, whilst a soldier would typically place themselves on rung v_4 or v_5.

3.2.4 Combining Alienation and Violence

Someone's motivation to violence is not usually a linear function of these component attributes. Except for the last moral setting we need *both* a sufficient level of alienation *and* an acquiescence to personally commit a violent act. This means that any simple combined additive score of these two attributes will not work well. However it is straightforward to propose a plausible non-linear function such as the one below - introduced for illustrative purposes. This function expresses coarse categories of violent motivation - listed in escalating order - (*immune,* m_0, *benign*, m_1, *open to adopting*, m_2, *aligned,* m_3, *enactor,* m_4) as a function of their alienation and violence - that might be used to inform a potential perpetrator's threat.

Alienation\Violence	v_1	v_2	v_3	v_4	v_5	v_6
a_1	m_0	m_0	m_0	m_0	m_0	m_0
a_2	m_0	m_0	m_1	m_1	m_1	m_4
a_3	m_0	m_0	m_1	m_1	m_1	m_4
a_4	m_0	m_0	m_1	m_1	m_2	m_4
a_5	m_0	m_0	m_2	m_2	m_2	m_4
a_6	m_0	m_0	m_2	m_2	m_2	m_4
a_7	m_0	m_0	m_2	m_3	m_4	m_4
a_8	m_0	m_0	m_4	m_4	m_4	m_4

3.2.5 Means

A potential perpetrator needs means as well as motive to enact a violent crime. In this context means can be helpfully split into two elements - one we call *training/skills* and the other an *embeddedness* within a facilitating group and a remoteness from other inhibitory human contacts.

Training to perpetrate a particular crime personally: Note that there may be various skills to add here such as ability to fight, use a knife, shoot, make a bomb or develop a biological weapon. So depending on context we may need a refinement of type of training. We illustrate this type of customised embellishments in later examples. Here we present the simplest and most generic scale. Training is a resilient attribute - once trained someone cannot be untrained, at least in the time frames we envisage in our applications.

1. Untrained in playing a role r in a criminal assault of type g.
2. Partially trained to enact r in a criminal assault of type g.
3. Fully trained to enact r in a criminal assault of type g.

Embedded in criminal fraternity: Assault criminals often at least benefit and sometimes need facilitation. One simple measurable scale of this social dynamic - which in any given context can in principle and with enough information be formally and unambiguously determined is given below.

1. Not meeting every two weeks with other similarly or more motived like minded criminals and embedded in contacts with immune people.
2. Meeting every two weeks remotely, for example electronically with similarly or more motived like minded criminals and embedded in contacts with immune people.
3. Meeting every two weeks electronically and physically with similarly or more motived like minded criminals while in full contact with immune people.
4. Meeting every two weeks electronically and physically with similarly or more motived like minded criminals and contact with immune people reduced by at least 50% from two years ago.

5. Meeting regularly electronically and physically with similarly or more motived like minded criminals and contact with immune people reduced to less than 10% from two years ago.

In many contexts "like minded" individuals will equate with being affiliated to some criminal organisation constituting a gang or brotherhood of believers. Embedding of course goes hand in hand with motivation each excites the other in the perpetration of violent crimes. However for modelling purposes it is useful to separate these two. Remedial acts or external events can change or disrupt a suspect's embedding immediately whilst decreasing motivation tends to move only slowly. We note that such embedding can already be in place because of the suspect's kinships or naturally evolving friendship groups.

For most assault crime perpetrated by a group, for example for a bombing, skills and embedding will both need to be present to an appropriate level before it is possible to commit the crime.

3.2.6 Opportunity

Opportunity in the sense we use here presents immediately before a crime is committed. So especially in a gang crime this links closest to the Pursuit, Protect and Recovery phase of an attack. Almost by definition an opportunity can develop or be frustrated very suddenly. In court cases the use of an alibi - often based on time and location - is often central to establishing no opportunity could have existed for someone to have committed a particular crime. On the other hand if a suspect can be shown that it is likely he was at the scene of the crime then this can be central to a prosecution case.

Here how we might choose to define opportunity depends very closely on the nature of the crime in question as are the countermeasures that might be taken to minimise the possibility of presenting the suspect with an opportunity. A generic template can therefore only be quite coarse. However for most assault crime we nearly need two indicators to be positive:

1. A vulnerable target of the assault has been identified by the potential perpetrator (perhaps through his fraternity).
2. A team (possibly of one) of suitably skilled assailants that can plan, prepare and carry out the assault is currently in place.

Opportunity can of course be affected by prophylactic measures to safeguard potential targets whether these are children - as in the first example - or protecting the potential victim in the second. However actions to frustrate opportunity in the short term often require intervention by the police. This is summarised in Fig. 6.

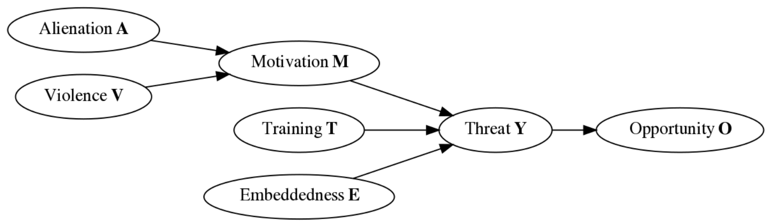

Figure 6: The computation graph summarising the process described above.

3.3 RCEG Templates for Assault Crime: Pursuit

3.3.1 Introduction

The Markov nature of these paths of an RDCEG building on the motivation, current modus operandi of perpetrators, histories of past criminals and knowledge of developing environments ensures two different individuals $\omega, \omega' \in \Omega_t$ classified as having reached the same "position" in their lives, i.e. for whom $w_t(\omega) = w_t(\omega')$ at t, will by definition respect the same probability law over their future participations in incident cluster if given the same subsequent external stimuli.

For simplicity the next two examples coarsen the above scales and our population Ω_t is restricted to those currently in the most threatening radicalised states: when suspects need to be actively pursued. In the diagrams below we have contracted our depiction of a directed graph using a common association: two edges pointing in different directions between two vertices will be replaced by a single undirected edge. For both examples, a person in position R is radicalised (m_3), M is personally motivated to attack (m_4), T is personally trained to attack, E is embedded in a gang, O^e has opportunity to perpetrate an incident when embedded and O^l is acting as a loner. Finally S, P, F denote the assault incident of the given type is a success, partial success or a failure respectively.

3.3.2 Template for a lone attack crime

Assume only one skill is needed to perpetrate the crime in question (e.g. ability to use knife, gun or a vehicle as a weapon). An RDCEG an expert might hypothesise is given in Fig. 7.

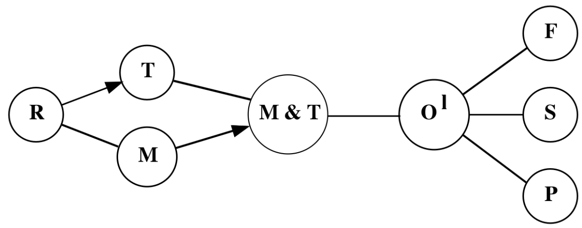

Figure 7: RDCEG for a lone attack.

Note that the representation expresses the hypothesis that some transitions are modelled as impossible. For example once someone is trained they remain so. Note that the RDCEG has one absorbing state and its first two states are non-recurrent. The structure of the semi-Markov transition matrix is given below.

	R	M	T	M&T	O	S	P	F	I
R	$*$	$*$	$*$	0	0	0	0	0	$\star$
M	$*$	$*$	0	$*$	0	0	0	0	$\star$
T	0	0	$*$	$*$	0	0	0	0	$\star$
M&T	0	0	$*$	$*$	$*$	0	0	0	$\star$
O	0	0	0	$*$	$*$	$*$	$*$	$*$	$\star$
S	0	0	0	0	$*$	0	0	0	$\star$
P	0	0	0	0	$*$	0	0	0	$\star$
F	0	0	0	0	$*$	0	0	0	$\star$
I	0	0	0	0	0	0	0	0	$\star$

For some agencies interest might focus on the *population* of such criminals at a given location. At this population level danger at time t of $\omega \in \Omega_t$ can be characterised by the *number* of ω in each of the positions above. These random variables can obviously be associated with the positions of this graph. Alternatively identified people can either be assigned a single current position or more commonly to a distribution over a collection of possible positions.

3.3.3 A template for a straight forward assisted crime

Next consider a slightly more complicated scenario where the trained person needs to be embedded within a cell of conspirators to perpetrate the assault, a cell a suspect can in principle transition in and out of. One simple hypothesised RDCEG in this case is given in Fig. 8.

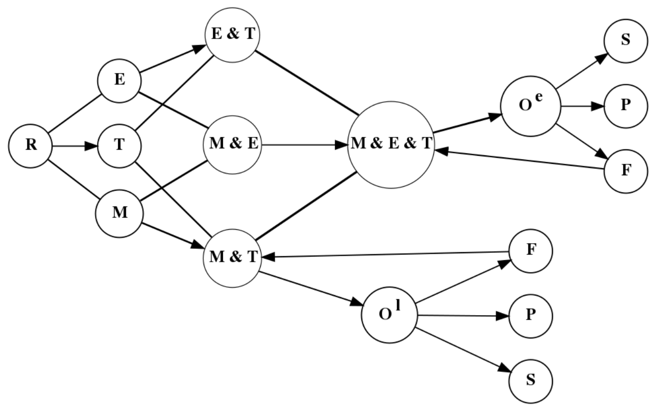

Figure 8:

Again, note that the graph above embeds a hypothesis that training is not forgotten. Also in distinction to the first example and to illustrate the variety of hypotheses such templates can exhibit it is assumed that the perpetrators of a successful or partially successful attack are killed or neutralised but failed attackers may still be free to attack again. We keep positions O^l and O^e separate here because the vectors of edge probabilities determining the success or failure of a loner or an enabled attacker are often quite different. A relatively small number of states describing the critical features of the process ensures the computation associated with learning and model selection between this model and similar competing ones is manageable.

Perhaps most important to note is that even in this very simple model the states of the population do not lie on a single escalating scale contrary to many early provenly naive models of the recent past: our model captured from expert judgments and survey information describes a more nuanced progression to code into a probability state space. This description can be further elaborated until it captures all hypotheses a criminologist might want to hypothesise no matter how complex these hypotheses might be.

4 Sample Distributions and Signatures

4.1 Introduction

There are at least four reasons for collecting data during a criminal investigation:

1. to help estimate - at a *population level* - the numbers of those within it who are at various rungs of escalation into potentially eventually perpetrating a criminal assault. This is to better target various resources both positively dissuading those at particular stages of threat to turn away and also of frustrating people at those rungs from being drawn into crime. An

example of the former would be general public information of the dangers of engaging in these processes and the latter of making access to criminalising instruments, such as a radicalising electronic web personal sites.

2. to help *identify those individuals who are likely to currently be on a higher rung* of the criminal ladder and might be a threat in the medium term. This would be to help and monitor them and their families to prevent or inhibit this progression over the medium term through social services.
3. to identify and *frustrate those individuals who are now a serious threat* of perpetrating a violent assault in the immediate future, through police action.
4. to *provide evidence* about the most threatening individuals *that will stand up in court* and lead to their arrest and conviction.

Although concerning the same population, each one of these types of inference needs to be treated rather differently. RDCEGs of the last two cases have already been illustrated. So for the remainder of this paper we will now focus on the first two of these.

4.2 Bayes Rule and Feature Selection

Bayesian uncertainty handling is in principle straightforward and the generic methodology can be immediately applied to the state space defined by an RDCEG like the ones discussed above. First the modeller helps the domain expert build a prior distribution on the states. Probability judgments concerning a criminal's possible latent positions defined as functions of the scales illustrated above within the threatening population are first elicited. In this context we find that these populations first need to be partitioned into pertinent classes defined as a function of elicited covariates defining the threat posed by a particular person within the population. The second step is to determine the sample distribution that we might have gathered either associated with the activities of known individuals or concerning appropriate population studies. The third step is purely an algorithmic one. Bayes Rule is applied to discover on the basis of the evidence available the client's posterior probabilities about either the distribution of the population of interest or their probability distribution concerning a particular person of interest.

The first step of this process needs some care. However the methods to perform such an elicitation once we have an overarching structure like the RDCEG are now widely available and have been successfully applied in a multitude of contexts see e.g. [19, 41, 56]. We briefly illustrate this procedure for a particular criminal setting below.

Certain aspects within the second and third steps are hazardous. New technologies [6, 19, 24] mainly applied to public health have already demonstrated how relevant data from population surveys or complete observational experiments can be used to search a space of possible plausible models within a CEG class of a given domain. It is possible simply to adapt standard established

Bayesian model selection techniques to this class embellishing the method we have illustrated in our first example. With the necessary time homogeneity assumptions it is straightforward to adjust these methods so that they apply to a DCEG [8, 20, 21, 52].

However these fast closed form selection methods have mostly been applied only to settings based on data which are complete samples of the whole population - for one exception see [7]. In the settings considered in this paper the data sets are usually extremely biased and furthermore data is missing not at random. In principle this does not cause a problem - the Bayesian method described above still applies, although these issues mean that much more care is needed to model sample distributions accurately and to perform inference. Two issues are especially important. First model selection must perforce lean heavily on expert judgments about the parts of the population about whom we have little empirical information. This is a logical necessity. Thus routine methods used in machine learning need significant adaptation before they are fit for purpose.

Of course successful policing is all about making inferences based on unbalanced information and intelligent but subjective judgments of this kind. However it does mean that it is essential for insights from criminology to be embedded into any methodology if that methodology is to be successful. Otherwise unsafe inferences will be made. Policing inferences cannot be substituted by automatic data driven inference some would like to see. Rather collected data can only legitimately be used to bring into better focus inferences already structure by theory and established practice using a framework that is rich enough to embed standard inferences at its core. The RDCEG is one framework that can aid this process.

If sufficient studies of the general population were to exist to inform the corresponding probability table assignments then methods based on likelihood ratio techniques [56] exist. In this idealised setting with care data centric inferences can provide statistics that can give more or less unequivocal numerical measures of the strength of evidence for or against certain hypotheses. Indeed many such experiments and studies have been performed that enable formal - and almost objective - uncertainty handling in a related field to combine forensic activity evidence to determine evidence to support the guilt or innocence of a suspect. The CEG and BN framework is one of a number of a useful frameworks around which to construct these estimates [4, 19]. However even there the extensive studies that have been performed only comprehensively inform certain types of criminal cases.

Note that at a population level whilst quite detailed information about criminals that are actually caught and convicted of crimes - through data collected leading to someone's arrest and conviction - evidence is parse concerning people who commit undetected crimes or the innocent. Fortunately the technologies for performing both the elicitation and the Bayesian learning given such biased information are now available. A precedent for an analogous translation is provided by the achievements in forensic science [4, 27, 45]. These methodologies are now ready to be translated into crime prevention. The potential bias of the information available may be coloured by police preconceptions about the profile of a potential perpetrator and have recently well aired in relation to the controversies behind predictive policing e.g. [34]. Despite the output of many studies from a formal inferential point of view it is not enough that samples of observed suspects in a given position share a particular property with high probability. We also need that someone in each of the other alternative benign positions performs those acts with low probability. Good machine learners working in these area acknowledge these hazards - see e.g. [30, 63] who express extreme caution about using even state of the art machine learning techniques to this domain.

Example 4 *Consider the following type of data that is collected from surveys. Here the hypothetical life histories of 100 known violent criminals of a particular type have been studied and it is noted that 34 of these have a relation who has a criminal history. Because of this it might be suggested that being embedded in a criminal family is a predictor of whether or not another person might be such a violent criminal.*

	Criminal	Not criminal	
Related to criminal	**38**	962	1000
No relation	**62**	x?	?
	100	?	

The probability $P(C|R)$ - a person is a criminal (event C) given he has a relative who is also a criminal (event R) - is the probability of interest. However the sample proportion estimates the conditional probability $P(R|C) = 0.38$ i.e. conditions the wrong way round: the Prosecutor Fallacy [56].

One common attempt to solve this issue within machine learning is to then extend the study by taking the criminal relations we have found in the study in this way: the so called snowball sample [63]. We then find the number of these that are not a violent criminal. This gives us a new entry which in our table we let be 962. However this naive sampling method can still be seriously flawed since by definition the relatives who seed this analysis are related to a violent criminal! Unless placed within the context of a development which includes embeddedness - for example through using an RDCEG - this biases the sample severely. Note the use of the obvious sample proportion 0.038 as an estimate of $P(C|R)$ in this context is likely to be far too large. These problems persist and indeed accumulate once we search for more than one explanatory cause. Although snowball sampling commonly used in machine learning can sometimes be debiased, this is only possible by making severe assumptions which in our contexts are usually untenable [33].

A further complication here is that people of interest are *rational*. So - at least as they enter more dangerous stages - we must allow for the fact that a suspect disguise or limit the more direct indicators of their threat. So it is essential that within the feature function selection we bear this in mind [5]. People can sometimes self identify themselves in the classes above. However if these are extreme stages we may need to fold in the possibility that statements are bravado. On the other hand if someone appears to lie in a benign stage appears then this might suggest that person is in fact actually dangerous but in hiding. Some recent systematic ways of introducing the gaming element for classes of CEGs is given in [61].

A final challenge is presented by the sheer volume of potentially relevant data currently collected - most of which is irrelevant - streaming in real time through for example online activities and CCTV images. Unless appropriately filtered this data can be totally overwhelming. We need to first develop new statistical and machine learning technologies to determine the *functions* of the data that might be relevant to whether or not someone is in each of the given positions in a hypothesised RDCEG. We illustrate below that once elicited the RDCEG can be used to guide the careful choice of these functions so that state of the art Natural Language and Image Processing technologies to be trained to this application, for example, to estimate the hyperparameters associated with these sample distributions [40].

So statistics with *discriminatory power* to differentiate the signatures of different people in each different position are essential. Here - because a priori it is much more likely that anyone - regardless of their indexing covariates is much less likely to be honest than be a criminal sample are unbiased - this is especially important. Various authors have begun to do this [39, 63] although in an unsupported way.

4.3 Population models

Especially when appraising the effectiveness of measures designed to reduce the probability that any individual $\omega \in \Omega_t$ will become a danger to the public, there are some useful summary statistics we use later. These focus on the individuals within the population and marginalise out information about interactions between $\omega \in \Omega_t$. A useful partition of Ω_t separates those whose identity is known to the authorities $\omega_i \in K_t$ and those $\omega_i \in \overline{K}_t$ who are not. A selection of the covariates of ω_i might typically be available, for example from those recently released from prison. With the caveats above this information can be used to inform our estimates. This is developed in more detail in [55].

One obvious projection on to each time slice $t > 0$ - conditions on certain relevant subsets of covariates $\boldsymbol{x}_t$ defining a subpopulation $\Omega_t(\boldsymbol{x}_t)$ of interest - focuses on the random vector of numbers $\boldsymbol{N}_t(\boldsymbol{x}_t)$ of people in the population of people in the different threat positions in the subpopulation. This vector can be written as

$$\boldsymbol{N}_t(\boldsymbol{x}_t) \triangleq \{N_t(w_t, \boldsymbol{x}_t) : w_t \in W_t\}\,.$$

We illustrate later how when the dynamics are slow enough that the aggregated semi-Markov process on individuals can be well approximated by a Markov one then the topology of the RDCEG can be straightforwardly transformed into a better studied Markov process [**Collazo 18**, 21]. We note that from the comments above we are typically informed much better about people in the more threatening states than those who are not so such samples are very far from being multinomial. However various features of its joint distribution can be simply calculated using the topology of the RDCEG and be used within any decision support tool to be used on this population.

5 Radicalisation to Violent Extremism

5.1 Introduction

In the next two sections we illustrate how to structure a particular application of the technology above into a particular domain: religious and political radicalisation leading to attacks on the general public. Henceforth call this radicalisation to violent extremism (RVE) currently a serious threat of such assaults on the general public in the UK - most recently from Al-Qaeda and then ISIS - and so there is a point to such a study. Secondly there is now a vast sociological and psychological literature about this domain giving deep understanding of some of the processes behind this development. For a summary of some of these see [25]. Thirdly there have recently been efforts especially within machine learning to develop tools to harvest real time evidence to support the identification of particular individuals who might perpetrate such crimes. This paper shows, by using the framework of the RDCEG how these technologies can be embedded in sociological models. For simplicity here we will focus on only those people who will actually be part of a perpetrated attack and not those who are simply inciters.

The study of radicalisation processes is now quite advanced. A good review of some of the more established work in this area is given in [11]. There he points out that most theorists think of radicalisation to violent extremism as a process - like the one modelled above - quoting [44] "'[Radicalisation is] the end point in a dialectic process that gradually pushes an individual towards a commitment to violence over time". Our first step is to translate the broad categorisations we developed above into categories customised to this particular subpopulation. For example, in a detailed study of one class of violent extremists [35] points to four elements in their recruitment and capability which closely parallel the four components of our motivation and means scales. He argues as we have that these do not in constitute components of an escalating scale but rather four separate parallel escalating scales. He used the term "indoctrination" for what we have called "alienation", "violent acculturation" for "violence", "relations building" for "embedding" and "training" for "training and skills". The NYPD scales of "pre- radicalisation" lead to "self identification" and then "indoctrination" [53].

In [12] the study of motivation for radicalisation in Denmark by [49] is discussed. These have also informed the scales we produce below. They unpick what might be meant by alienation in the context of Islamic extremism. They discover from their sample certain indicators of the higher alienation rungs in our scale: unemployment/ underemployment, lack of a role in society. Experiences of discrimination often pushes someone up this scale. Ideological justifications pulls them into a group further alienating them from mainstream society. With other authors they discover acculturation of violence as another defining feature.

The support of agencies aiming to discourage or prevent people from those rungs which make someone a danger to society is informed by a series of detailed interviews with another group. Horgan [36] found that those in their sample dropping down our scales tended to do so not so much from a decrease of feeling of alienation but through intense disillusionment that violent struggle was an appropriate way to address injustices. This would often be paired with a realisation that the brotherhood with whom they associated were far from being upright and sometimes people whose attitudes and behaviours disgusted the subject. This suggests that encouraging a subject to modify their disruption rung - for example trying to convince them that violence against the general public is never justified whatever its perceived inadequacies/sinfulness. Convincing them of the ideological impurity of their brotherhood might provide compelling and immediate dissuasions to some, thus disturbing their embeddedness. Such tools might then go in tandem with attempts to rehabilitate a person - for example empowering them to find a fulfilling job or enjoy a normal family life: the latter associated with much longer processes and never attractive or attainable to some. Again the scales we give below identify when and where such tools might be most effective.

Various documents provide advice to professionals both to identify radicalisation leading to violent extremism with suggestions about how to address these issues: see for example [18, 29, 46, 47] and are used below where we relate activities to our scales. So for example in [29] indicators of being in a "concerning" point on our alienation scale on rungs scales a_5 and above is being more removed from normal social networks than they would normally be, stated commitments to a radical ideology and irritation or anger about contrary views. "Concern" associated with violence (greater than v_3 on our scale) would be the use of language that advocated violence or aggression. Indicators of people deserving "attention" would be a complete commitment to and engagement in the ideology and almost complete disengagement with previous friends, family, strong expressions of hostility to their "enemy" including law enforcement and government, and seeing violence as a necessary and legitimate instrument (rungs a_7, a_8) and being prepared to plan and prepare a violent act v_5, v_6. We note that these scale indicators if existing together would lead to a the highest rung m_4 in our motivation scale.

5.2 Adjusting terminology to apply to radicalisation processes

The Alienation scale below expands the escalating scales based on SMT and given in [65] who label states of mind openness to new world views - a_3, a_4, religious/political seeking - a_5, frame alignment - a_6, indoctrinated - a_7, a_8. A path orientates a person towards some *specific alignment* with a particular ideal - for example an Islamic one - although we do not restrict ourselves to this instance here.

1. *Benign*: I am content with my life and its trajectory and the normal evolution of mainstream society, a_1.
2. *Dissatisfied*: I am dissatisfied with my current life and am currently looking for better opportunities within it, a_2.
3. *Disaffected*: I am disaffected with my life chances; there seems to be no future for me, a_3.
4. *Converted*: I am not part of this society and do not share its ideals and aspirations. I instead have a different set of ideals aligned to my brothers and sisters. I identify with the suffering of my kin, a_4.
5. *Radicalised*: Some sort of revolutionary change is needed in the UK to transform it to be consistent with my ideology and transform its politics to be sympathetic with our beliefs and life styles. I love my kin and only them, a_5.
6. *Hating*: Those people representing or acquiescing with the current order are my enemies; we must fight these people, they are filth a_6.
7. *Affiliated*: Any acts - criminal or not - that provoke a complete and total transformation of a society so that its politics wholly supports my ideology and codes of actions and tolerates no other are fully justified; I wholeheartedly support those who are prepared to bring about such a transformation; they represent me, a_7.
8. *Embodying*: I, personally, yearn to take an active role (within my own moral compass) to instigate this transformation, I am prepared to be imprisoned or to die for this cause, a_8.

The generic violence scale needs no modification, neither does the matrix function defining our motivation scale. The terminology describing the motivational states can however be usefully adjusted: *immune*, m_0, "I would find such a narrative of RVE abhorrent and unacceptable"; *benign*, m_1, "I am not dwelling on these issues"; *open*, m_2, "I might respond positively to an RVE narrative"; *aligned*, m_3, "I wholeheartedly support RVE"; *indoctrinated*, m_4, "I embrace my particular RVE narrative. This is my world view and I am determined to be a soldier for its cause."

5.2.1 Means

The RVE embedded, encouraged and enabled: There are two types of RVE. Lone wolf attacks where someone acts individually and is only enabled electronically. The second are ones where RVE need the active engagement of other like minded people within a collection of people drawn from others affiliated to the same ideological group. We call this local collection of people that person's *brotherhood*. This second type is especially interesting.

As has been demonstrated through benign religious groupings, for an ideological movement to have traction it usually needs be constituted by groups who meet, care for and encourage one another. This social networking increases the commitment and courage of those following the ideology. For those described as RVE it is these ends the police need to frustrate. Even those planning to act alone therefore are most dangerous when embedded in a brotherhood. Engagement within a brotherhood, whilst not necessary to perpetrate and act of violence increases the extent of the potential harm a potential perpetrator might induce. An attempt to capture this social dynamic for RVEs below in measurable categories.

1. Not meeting regularly with others affiliated to the same ideology - embedded in contacts with other friends and family.
2. Meeting regularly remotely e.g. electronically with affiliates but most contacts are with other friends and relatives.
3. Meeting regularly electronically and physically with others in the brotherhood of believers but also in ordinary contact with other friends and relatives.
4. Meeting regularly electronically and physically with others in the brotherhood of believers and have reduced contacts with family and friends.
5. Meeting regularly electronically and physically with others in the brotherhood of believers and have minimal contacts with family and friends.

As noted before, embedding goes hand in hand with motivation each exciting the other - reinforcing the ideological framework and acceptance of violence that provides motivation. However for modelling purposes it is useful to separate these two. Remedial acts or external events can change or disrupt a suspect's embedding immediately whilst decreasing motivation tends to move only slowly. We note that such embedding can already be in place because of the suspect's kinships or naturally evolving friendship groups.

Skills to be part of an act of violence: The final element we need to address in this model is the capabilities of an individual motivated to be part of an attack on the general public. There are a number roles r in forms of attack g used by radicals, each needing different levels of skill from the potential perpetrator and making him or her more or less useful. We list the more common roles below.

1. A mule: Need someone to demonstrate (physically or electronically) what you do but usually little training required. These roles might be acid attackers, or bomb mules.
2. Vehicle as a weapon: Need to be able to drive the vehicle.
3. A knife attacker: Some training and practice to do this well but no real technical skills.
4. Attacker using fire arms: Need some training on a firing range.
5. Bomb maker: Some specific skill needed here - although such training is available on the web if you know where to look.

6. A strategic planner: Able to identify targets, formulate a plan of attack and draw together participants to enact the plan - needs logistic and organisational skills.
7. Making a biological or nuclear weapon: More specific training needed here.

These abilities reflect the extent of the danger presented by an individual. A single person can of course have many skills: for example a bomb mule may well have also made the bomb. Note that someone who has received military training over at least 6 months by the affiliating group is likely to be skilled in the first 6 bullets and so is especially dangerous. Also note however that the first three bullets are skills that need very little training. A significant proportion of recent attacks in the UK have been of this type. These are especially difficult for the police to frustrate because the lead time between being motivated to make an attack and enacting that attack can be very short.

In most circumstances it should in principle be give a probability distribution of the skills of a particular person of interest based on their likely history and the three point scale of training. How this may be done is illustrated in the next section.

Opportunity: Finally the individual needs the opportunity to commit the crime. Sometimes targets will be heavily protected so this may be an issue or increased police presence may frustrate an attack - at least forcing its postponement. Certainly in all cases someone currently held in custody cannot be party of an attack force and is one of the instruments a government has to frustrate an attack. These elements typically inform police operations and link closely to the pursuit of criminal cells. The modelling of this feature is somewhat more complex and therefore developed in a later paper see [3].

5.3 Evidence based analysis of radicalising processes

5.3.1 The RDCEG

We need to set up a description of the RDCEG which is sufficiently detailed to capture the varieties of different ways a person can become a potential perpetrator but sufficiently coarse to be able to draw evidence from other individuals who have followed life paths that have at least some features in common. It is helpful therefore to consider first the variety of types some illustrated below and see that the categorisations below with the right RDCEG can still be applied to all of these. These examples are based loosely on a categorisation of young Al Qaeda fighters in [62].

Example 5 ***For status:*** *Soldier A is an immigrant and has recently returned from 2 years of action fighting for the cause to which he is known to be affiliated. He continues to be in electronic contact with people known to proselytise RVE. His brother is still in action abroad and he has just made local contacts with fellow sympathisers. On arriving in the UK he cannot find work and suffers discrimination. He realises he is an alien from UK society which he now despises.*

This man scores high on all scales and is a clear and potentially enduring threat. He may well believe in the justice of the movement, feel it is something he should engage in himself and of course he will be familiar with the possibility of risking death. So he scores highly on the alienation scale. From his experiences in the conflict he is likely to be desensitised to the immorality of violent acts, so score high on the disruption scale. He is very likely to be skilled enough to be an active participant in an attack because of his experiences abroad. He also appears to be embedding himself within a brotherhood in the UK through which he could coordinate an attack on the UK public. Note that for this person being brought up in a war zone has acculturated him to violence early in his life. By fighting abroad he has acquired skills that enable him to perpetrate an attack. The final element - his embedding has come last - he has found people in the UK he can relate to. $(V \to T \to A \to E)$

Example 6 ***For revenge:*** *A female UK recruit B's lover has been savagely killed by allies of the UK whilst he fought abroad. She now hates these allies and is determined to take revenge. She embraces death because life no longer holds a meaning. She is currently saving and borrowing money and plans to travel to fight for her lover's cause to then return to take retribution against the nation she hates. She has had long contacts with her lover's friends within the UK with whom she is embedded. This woman is clearly also a threat. However the potential threat she presents will be more delayed. It may be possible to prevent her from travelling to be fully equipped. Her commitment to a violent act might be entrenched or her ideological fervour for the cause may be flaky. In time her hate provoked plan might be reappraised. Note the different order here. She was first embedded, became violent, alienated and is about to train.* $(E \to V \to A \to T)$

Example 7 ***To belong:*** *An angry alienated young man C feels it is something he should be doing and is familiar with the possibility of risking death. He is intelligent but has been discriminated against from a young age and has many grievances. He had been engaging in criminal activities since his youth ending up in prison a number of times. Since leaving prison the last time he has met up with others who have become his close friends. He has converted to Islam and found order and structure in his life and a mission. He now embraces his new found faith which he understands as encouraging the violence which made up so much of his past life and redirects this to a "good" cause. The idea that by dying from suicide bombing he will go straight to heaven and all his sins be forgiven appeals. Here high motivation and a brotherhood has been added to an earlier acculturation to violence. This person is a real threat. However having a better understanding of Islam and a separation from his current brotherhood, replaced by another more benign version of his new religion might defuse his threat to himself and others. Note again the different order has occurred.* $(V \to A \to E \to T)$

Example 8 ***For thrills:*** *A young man D lusts after violence. He laps up material like beheadings on the web and has recently discovered a brotherhood* *living locally who also want to perpetrate violent acts. He is happy to join this gang who will train him to perpetrate other violence and engage with its ceremonials and diktats so that he can get hands-on experience of a local war. This man is dangerous. But the radicalisation is just a vehicle for expressing his violent urges and comes as a late add on. He will remain a danger to society whatever is done to mitigate the causes of RVE.* ($V \to E \to T \to A$)

The second and third cases might well respond to different stimuli both from each other and from the first and last, so they may well be combined. The first and last cases should be kept apart.

5.3.2 Explanatory variables

Several detailed but mainly qualitative studies of different groups of radicalised people and their pathways to eventual crime are now available and many of their categorisations - derived from an honest interviewee would be impossible to observe externally on an identified but remote person. However a short list of generic classifications could be observed and some of these are summarised below.

Relevant explanatory variables - such as those found on a passport have been the following. All these link in some way or other to vulnerability to RVE.

1. *Home location*: This is most critical because policies to counter RVE are typically enacted locally. This also is relevant to the ability of an individual to embed themselves in a brotherhood.
2. *Gender*: The majority of people who have perpetrated an attack have been male. It depends on the type of attack but as a ballpark figure perhaps only 5% of these have been perpetrated by women.
3. *Age*: About a third of people convicted of incidents associated with RVE are between 24 and 28 with their arrest happening about 2-3 years after first entering the highest of our motivation scale. In the West children less than 16 and people over 50 rarely commit an act.
4. *Religion*: Obviously by definition all Islamic terrorists would self identify as Muslim, usually Sunni. Similarly members of the Provision IRA self identified as Catholic Christians.
5. *Nationality*: Certain national cultures seem to be more consistent with the adoption of a violent ideology than others.

Prior probabilities for any identified person will normally reflect these categories. However it cannot be emphasised too strongly that although such categories inform our processes they provide a very biased picture of the potential of someone being identified as a threat - see the discussion of the Prosecutor Fallacy in the previous section. In particular prior probabilities must be informed by the available complementary evidence concerning the general public.

5.3.3 Activity level data, medium term

The next type of externally collectable data concern past events that can help inform our scales. They apply to known or partially identified individuals although local demographic information is collected on some of these and can be used as background information.

1. Is the person an immigrant to the UK? If so are they an immigrant from a very different culture? Have they emigrated from a risk region? These all signify a higher probability of being higher on the **alienation** scale. To the last we must add that such a person is more likely to either abhor violence or be much more acculturated than the general public. What is the probability that they are a skilled fighter (**training** scale) and acquainted with local people with radical views (so have a higher **embedding** score)?
2. Are they fluent in a language other than English? Such people can be more effectively covertly targeted by radicalisers and covertly communicate with others radicalised within an affiliation. So they can be higher on the **embedded** and **alienation** scores.
3. Do they have a prison conviction? These suggest a higher probability of being further up the **alienation** score than the general public and also to have been in contact with those with radical views (so possibly more **embedded**). If the conviction was for a violent crime or for a crime with a high probability of associated violence (such as being part of a drug gang) then this gives a higher probability on the **violence** scale.
4. Are they unemployed, underemployed or a student not engaging in their programme? If so this suggest a higher **alienation** score than that of the general public.
5. Have they travelled to risk regions and stayed for more than 6 months? Then they are more probable to have a high **training** score (need to offset this if they have relatives in the area which is an alternative explanation) and return more acculturated with violence (a high **violence** score).
6. Do they enjoy a kinship relationship with someone known to be radicalised (higher score on the **embedding** scale)?
7. Have they at some point in the past accessed and read radicalising web sites on more than one occasion, or been in contact with radicalisers (a higher **alienation** score)?
8. Has this person been in active communication with known radicalisers (high **alienation** and **embedding** score)?

5.3.4 Activity level, current behaviour

These are about an identified or partially identified person's current activity.

1. Currently accessing radicalising websites and meetings at least once a fortnight (high **alienation**) [29].

2. Openly evangelising RVE (high **alienation**), e.g. reposting with a positive sentiment on public electronic forums or physical meetings [43]. This can now be analysed directly from post/repost messages by finding features to discriminate between extremist and anti-extremist posts [64] and combinations of behaviour changes that signal RVEs from posts [51].
3. Openly aggressive to family and past friends (high **violence**) [29].
4. Currently in energetic contact with known radical friends and family - at least once a week - (high **alienation** and **embedding**) unless this is a frequent past state in which case a sudden reduction of this might indicate a rational act - see above - associated imminent involvement in an incident [43]. Methods that weigh electronic links to seed individuals already known RVE's have been studied [13, 16].
5. Shunning usual communications with non-radicalised friends and family (high **alienation** and **embedding**) known as "hardening" [43].
6. Saving money, trying to secure a false passport, researching travel options, websites for weapon making and use (high **training** and **violence**).

Changes in behaviour are especially useful; they indicate previous as well as new position. Note that some of these classifications need intelligence concerning the RVE in this person's network.

5.4 Sampling, questionnaires and position probabilities

Typically demographic information will be available about a particular population as a whole so that the probability distribution of the numbers in these categories in a given catchment can be made fairly tight. Analogous setting suggest one good way of eliciting this prior information is to first estimate the size of the whole population $\Omega_t(\boldsymbol{x})$ for each of the explanatory categories $\boldsymbol{x} \in \mathbb{X}'$ perhaps 6-20 of these categories. These numbers will typically be estimated with a variance which can be treated as very small and known. Typical categorisation is by geographical region, age, gender, nationality. We have found it useful to keep separate categories for recent immigrant from war torn countries and from prisons because these two categories can have rather different transition probabilities than the rest of the population.

Then for each $\Omega_t(\boldsymbol{x}) : \boldsymbol{x} \in \mathbb{X}'$ we estimate the pertinent functions of the explanatory indicators above and then for each $\boldsymbol{x} \in \mathbb{X}'$ the expected numbers currently in a position and of those transitioning from one position to an adjacent one in the RCEG over the next transition. Usually there will be less than 20 positions and as we have seen these transitions are sparse so this is feasible. Various refinements of this using the Markov prior equilibrium probabilities when we believe the transition probabilities are slow moving over time are outlined in [8].

For population inferences propensities need to be translated into sample distributions over the relevant populations for each position in our RDCEG, by covariate. Here we can use the first two lists of prompts to toggle the estimates towards reality. This takes a little time to set up but there are now many studies that can lead us into reasonable ballpark figures which are usually sufficient to discriminate policies. Recall we will have already elicited a partition of the population into the positions relevant to the crime. One of the points of using the RDCEG rather than just the positions is of course that fact that there is strong information about transition structure which the RDCEG hard wires into the statistical model. The implied covariance structure across the positions is for example very heterogeneous but also easily calculated using simple applications of tower rules found in basic probability modelling. If we need to appropriately accommodate sampling information then we will need at least the second moment properties of the probabilities in each positions. These can again be well approximated from eliciting directly quartile bounds on the numbers, and then using the same approximating formulae or alternatively by directly estimating the hyperparameters of Dirichlet priors on the relevant transition probabilities: see [19] for details.

Data we have concerning known individuals $\omega_i \in K_t$ within the population then can be embedded either formally by proposing sample distributions to various measures derived from the types of activity data mentioned above or where we simply use this information to populate the probabilities associated with these known individuals. Using this information we might choose to review our beliefs about $\omega_i \in \overline{K}_t$ especially those in regions with high concentrations of the most threatening of the $\omega_i \in K_t$.

Of course the assessments need to be regularly revisited. But revising assessments so that they better fit the developing environment is a much easier task: typically only a few of these will need adjusting driven by changes in environments in selected geographical or demographic regions. These will typically result in various transition probabilities changing or be the result of changes in numbers in various positions due to immigration or emigration.

Current activity levels inform at a population level via open source social media. Activities and especially changes in activities can be informative about changing transitions and provided that based round an expert validated RDCEG can often be naively accommodated using standard updates of the relevant transition probabilities. Often the inferences we undertake concern an individual's potential to act violently or to repent. The RDCEG can be used to start building a model of the threat posed by that particular individual into the future. More is known about the position of that given individual through their current activities. Those that are not known can be informed by the population probabilities elicited as above.

6 Agency RVE for Decision Support

6.1 Introduction

We have outlined a fine grained categorisation above. However this is often usefully coarsened by different government agencies who will have different responsibilities and focus. For example policy implementations directed at simply dissuading as many of the general public in involving themselves in RVE need to be targeted at those on the lower rungs of the ladders while remedial policies aimed at identified individuals are quite rare. For this purpose we therefore tend to focus on the *early signs* of radicalisation whilst the most threatening manifestations are handed over to other agencies. We illustrate the value of a statistical model for this type of decision support in the example below.

In 2010 the UK Home Office needed to evaluate the effectiveness of past programmes and the promise of future ones. Various types of portfolios of Prevent programmes were designed to enact early interventions designed to discourage vulnerable people being recruited into RVE. Some programmes were implemented through statutory requirements by various organisations. Examples of these included new policies and protocols enacted by social workers and probation officers and demands for attendance registers at universities. Others resourced the development of tools to identify people in danger of being (further) radicalised through their online behaviours, development of new methods to disrupt RVE propaganda and others to promote citizenship from those potential clients, educational programmes for schools and threatened neighbourhoods, training of personnel at mosques and other activities such as the English language training of Muslim monks. Government sought guidance in the wisest deployment of its resources to best mitigate threats of RVE into the medium and long term using such instruments.

Any decision support system had on the one hand to be formal enough to be transparent, defensible and so able to accommodate any available data [56]. On the other it needed to be credible to sociologists. This section builds on early work outlined in [50]. We show how the RDCEG - whose development was prompted by this application - can form the basis of a fit for purpose model of each *individual* in the process. We then demonstrate how a population model can be constructed from this RDCEG. Since the time of our original analyses a much richer data set has become available through publicly available posts on social media [1, 2, 10, 17, 39, 48, 60]. We will briefly illustrate how such information, most commonly used to police the internet, can also be used to better inform our models so that the overall efficacy of the evaluations above can be enhanced.

Information concerning training and the opportunities for criminal attacks could be truncated at boundaries of threat. Anyone who was in the course of the process reaching positions leading to these would be referred to other agencies and the police. Decision support would therefore focus on states informing countermeasures to address alienation and inhibit too many in the population being embedded in a threat group. Any population data used in the analysis needs to be in the public domain. In 2010 this was provided by some very scant survey information. Currently there is more public social media data although this data is patchy and very biased. So the most vital information comes from the prior information elicited from experts about the nature and extent of drifts within the positions of the population and the success or otherwise of the programmes themselves, as briefly discussed above.

6.1.1 An Initial RDCEG and its positions

An RDCEG of a typical $\omega \in \Omega_t$ who might benefit from a Prevent scheme is given below where we have omitted the demographic and geographic categories $\boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}$. We use our RVE scales discussed above.

Numbers	Positions	Description
N_{t1}	$w_1 = (m_1, e_0)$	benign & not embedded
N_{t2}	$w_2 = (m_2, e_0)$	open to radicalisation & not embedded
N_{t3}	$w_3 = (m_2, e_1)$	open to radicalisation & partially embedded
N_{t4}	$w_4 = (m_2, e_2)$	open to radicalisation & embedded
N_{t5}	$w_5 = (m_3, e_0)$	aligned & not embedded
N_{t6}	$w_6 = (m_3, e_1)$	aligned & partially embedded
N_{t7}	$w_7 = (\mathbf{m_3}, \mathbf{e_2})$	aligned & embedded
N_{t8}	$w_8 = (\mathbf{m_4}, \mathbf{e_0})$	indoctrinated & not embedded
N_{t9}	$w_9 = (\mathbf{m_4}, \mathbf{e_1})$	indoctrinated & partially embedded
N_{t10}	$w_{10} = (\mathbf{m_4}, \mathbf{e_2})$	indoctrinated & embedded

Assume that when there is a high probability someone lay in the last 4 states at any future time then they would be referred to other agencies as a potential danger. Suppose that we frame Prevent strategies in terms of minimising the number of people who at some point in the future end up in one of these four states, and a severe danger when arriving in the last. The simplest possible transitions where double directions are represented by an undirected edge is given below. The full RDCEG simply copies the same structure for each possible value of covariates. Most of the positions in this system are associated with the same (m, e) label but different covariates will be in the same stage and so enjoy the same colours. The most critical positions are given in bold and the severe one is boxed.

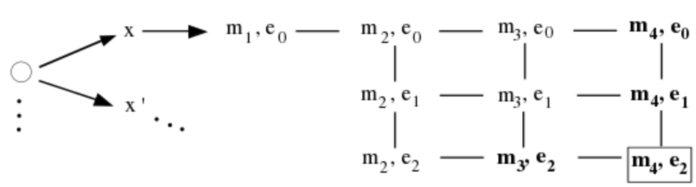

Figure 9: An initial RDCEG for modelling the portfolios which are part of the Prevent strategies.

Let $G_t = \{\boldsymbol{g}_{ti}(d, \boldsymbol{x})\}$, $i = 1, 2, \ldots, 11$, $d \in D$, $\boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}$ be a matrix which depends on the suite $d \in D$ of policies adopted and the covariates $\boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}$

but whose structure of zeroes is assumed here to be policy independent and covariate independent. Thus for all values $d \in D$, $\boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}$ the shape of the semi-Markov transition matrix G_t takes the structural form given below.

rungs	m_1, e_0	m_2, e_0	m_2, e_1	m_2, e_2	m_3, e_0	m_3, e_1	m_3, e_2	m_4, e_0	m_4, e_1	m_4, e_2	i
m_1, e_0	0	$*$	0	0	0	0	0	0	0	0	$\star$
m_2, e_0	$*$	0	$*$	0	$*$	0	0	0	0	0	$\star$
m_2, e_1	0	$*$	0	$*$	0	$*$	0	0	0	0	$\star$
m_2, e_2	0	0	$*$	0	0	$*$	0	0	0	0	$\star$
m_3, e_0	0	$*$	0	0	0	$*$	0	$*$	0	0	$\star$
m_3, e_1	0	0	$*$	0	$*$	0	0	0	$*$	$*$	$\star$
m_3, e_2	0	0	0	$*$	0	$*$	0	0	0	0	$\star$
m_4, e_0	0	0	0	0	$*$	0	0	0	$*$	0	$\star$
m_4, e_1	0	0	0	0	0	$*$	0	$*$	0	$*$	$\star$
m_4, e_2	0	0	0	0	0	0	$*$	0	$*$	0	$\star$
i	0	0	0	0	0	0	0	0	0	0	1

Let $\rho_{ti}(\boldsymbol{x},d)$, $i = 1, 2, \ldots, 10$ denote the probability any $\omega \in \Omega_t$ lies in the i^{th} state at time t given they are not in the immune state for each considered subset of covariates $\boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}$, each possible suite of decisions $d \in D$ and in quarter t. Now write

$$\begin{aligned}
\boldsymbol{\rho}_t(\boldsymbol{x},d) &\triangleq (\rho_{t1}(\boldsymbol{x},d), \rho_{t2}(\boldsymbol{x},d), \ldots, \rho_{t10}(\boldsymbol{x},d)) \\
\boldsymbol{\rho}_t^-(\boldsymbol{x},d) &\triangleq (\rho_{t1}(\boldsymbol{x},d), \rho_{t2}(\boldsymbol{x},d), \ldots, \rho_{t6}(\boldsymbol{x},d)) \\
\boldsymbol{\rho}_t^+(\boldsymbol{x},d) &\triangleq (\rho_{t7}(\boldsymbol{x},d), \rho_{t8}(\boldsymbol{x},d), \rho_{t9}(\boldsymbol{x},d), \rho_{t10}(\boldsymbol{x},d))
\end{aligned}$$

All the states of this Markov chain are transient other that the immune state which is absorbing.

6.2 Population Modelling and Policy Appraisal

The policies of this agency will be often directed at those who lie in the first 6 states. The aim of the policies are to prevent as many individuals as possible entering the other states at future time step when a suite of policies are in place. The predicted effect of the policies can then be measured.

On the other hand the last 4 states will define the extent of failure of any portfolio of policies. A failure could be said to occur when anyone reaches the last 4 states at some time and an extreme failure when they reach the last. So $\boldsymbol{\rho}_t^+(\boldsymbol{x},d)$ and particularly $\rho_{t10}(\boldsymbol{x},d)$, $t \geq t_0$, for an as yet unknown $\omega \in \Omega_t$ will be of central interest. The weight given to each of these will depend on the client's utility function - an internal issue which needs to balance both the extent of the problem with its immediacy. Here it is sufficient to point out that these quantities are simple functions of the terms above.

Example 9 *Suppose we consider only policies associated with the subpopulation* $\boldsymbol{x} = \boldsymbol{x}^* \in \mathbb{X}'$ *- perhaps several potential policies aimed at reducing alienation of men who are currently in the age range 16-26 years who live in a particular part of a given city. Government say they want to finance the new policy* $\delta = \delta^* \in \Delta$ *to add to its current suite that minimises the probability* $\theta^*(d \cup \delta)$ *someone in the* *current population at some point will reach position w_{10}. This is a simple function of the current probabilities $\boldsymbol{\rho}_t(\boldsymbol{x},d)$ and the elicited entries of the sequence of future semi-Markov transitions. Let $G^*(d \cup \delta) \triangleq \{\boldsymbol{g}^*_{ti}(d \cup \delta, \boldsymbol{x}^*)\}_{1 \leq i \leq 11}$ where*

$$\boldsymbol{g}^*_{ti}(d \cup \delta, \boldsymbol{x}^*) \triangleq \boldsymbol{g}_{ti}(d \cup \delta, \boldsymbol{x}^*) \ 1 \leq i \neq 10 \leq 11$$

but where we make w_{10} an absorbing state by setting $\boldsymbol{g}^*_{t10}(d \cup \delta, \boldsymbol{x}^*)$ to be a row with all components zero but the 10^{th} which is a one. This new semi-Markov transition matrix now has two absorbing states w_{10} and i. It is then easily checked that $\theta^*(d \cup \delta)$ is exactly the probability of being eventually absorbed into w_{10} rather than i under the action of $G^*_t(d \cup \delta)$ from the starting vector $\boldsymbol{\rho}_t(\boldsymbol{x},d)$.

Example 10 *Now suppose the aim is to prevent as many as possible of the population entering a failure rather than extreme failure state. This calculation is just as simple. The new policy $\delta = \delta^* \in \Delta$ minimising the probability $\theta^+(d \cup \delta)$ is now the sum of the probabilities of landing in the 4 new absorbing states associated with states w_7, w_8, w_9, w_{10} associated with the current probabilities $\boldsymbol{\rho}_t(\boldsymbol{x}, d)$ and the different transition matrix $G^+_t(d \cup \delta) \triangleq \left\{\boldsymbol{g}^+_{ti}(d \cup \delta, \boldsymbol{x}^*)\right\}_{1 \leq i \leq 11}$ elicited entries of the sequence of future semi-Markov transitions. Let $G^*(d \cup \delta) \triangleq \{\boldsymbol{g}^*_{ti}(d \cup \delta, \boldsymbol{x}^*)\}_{1 \leq i \leq 11}$ where*

$$\boldsymbol{g}^+_{ti}(d \cup \delta, \boldsymbol{x}^*) \triangleq \boldsymbol{g}_{ti}(d \cup \delta, \boldsymbol{x}^*) \ 1 \leq i \neq 7, 8, 9, 10 \leq 11$$

but where we make w_{10} an absorbing state by setting $\boldsymbol{g}^+_{ti}(d \cup \delta, \boldsymbol{x}^*)$ to be a row with all components zero but the i^{th} which is a one, $i = 7, 8, 9, 10$.

Note that if the decision maker wanted to trade off these two attributes and she had preference independent attributes then she would simply choose the policy δ to minimise a linear function of these two scores that function a criterion weight that would reflect their relative priority [56]. Usually each Prevent project had its own small demographic so its portfolios could be disaggregated and impacts assessed individually and the scores then sum, weighted by the numbers believed to be in each explanatory class $\boldsymbol{x} \in \mathbb{X}'$. All such calculations are simple to make once the the elicited scores are there and can be used to initialise and explain why potential acts are good or bad.

But is it feasible to elicit the * entries in the different transitions? For each $\boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}$ and suite of policy options $d \in D$ there are 25 such probabilities if we use the row sum to one condition to determine probabilities in the immune column i given by $\star$. However only at most a small number of these transition will depend on $\boldsymbol{x}$ and d. So we expect less than 100 such elicitations to be necessary. This order of magnitude makes the elicitation non-trivial but certainly feasible [9, 56]. Adjustment of probabilities in the light of unpredicted dramatic events occurring in the wider world not embedded in the model structure such as pertinent outbreaks of war, police outrages especially when such events are given wide news coverage are sometimes needed. However if curated, such adaptations are usually straightforward to perform once the early quantifications are in place. We emphasise that the given RDCEG is illustrative of the simplest plausible model and others might be better supported. However it is of about the order of magnitude most agencies might need and so adequate for illustrative processes.

6.3 Population Modelling

Deliberative decision support outlined below can be fashioned to respond to changing circumstances. It has been noted that the transition times between positions in the early stages of radicalisation are also usually quite slow. In this case this gradual development can therefore be well approximated as a quarterly time series. Any $\omega \in \Omega_t$ would remain or move up or down one rung of the measures in that three month window or during that time renounce the process completely and drop into the immune category. So for this model of early developments into RVE, because the probabilities of transitioning into a non-absorbing state were not large, the RDCEG above can in fact be well approximated by a simpler Markov rather than semi-Markov process. This simplifies further evaluation.

The early process to full affiliation often takes time. So little is lost especially when used for a population study where interest lies in the overall effect of a Prevent scheme rather than its effect on a given person by assuming this is actually a Markov rather than semi-Markov process: all the adjacent position except the immune state which is absorbing will have small associated probabilities and so are likely to occur at most once in the quarter. Most of the non-zero transition probabilities in this table need to be elicited using expert judgments. This simplifies the analysis even further because then for many purposes the shape of transition matrix $\widetilde{H}_t$ of this Markov approximation of G_t which we assume here applies to transitions made in a 3 month period, for each given value of the covariate vector $\boldsymbol{x}$ would then take the structural form given below.

rungs	m_1, e_0	m_2, e_0	m_2, e_1	m_2, e_2	m_3, e_0	m_3, e_1	m_3, e_2	m_4, e_0	m_4, e_1	m_4, e_2	i
m_1, e_0	$*$	$*$	0	0	0	0	0	0	0	0	$\star$
m_2, e_0	$*$	$*$	$*$	0	$*$	0	0	0	0	0	$\star$
m_2, e_1	0	$*$	$*$	$*$	0	$*$	0	0	0	0	$\star$
m_2, e_2	0	0	$*$	$*$	0	$*$	0	0	0	0	$\star$
m_3, e_0	0	$*$	0	0	$*$	$*$	0	$*$	0	0	$\star$
m_3, e_1	0	0	$*$	0	$*$	$*$	0	0	$*$	$*$	$\star$
m_3, e_2	0	0	0	$*$	0	$*$	$*$	0	0	0	$\star$
m_4, e_0	0	0	0	0	$*$	0	0	$*$	$*$	0	$\star$
m_4, e_1	0	0	0	0	0	$*$	0	$*$	$*$	$*$	$\star$
m_4, e_2	0	0	0	0	0	0	$*$	0	$*$	$*$	$\star$
i	0	0	0	0	0	0	0	0	0	0	1

Note that we simply now have an additional set of 10 diagonal elements. It may be useful to write

$$\widetilde{H}_t = \begin{bmatrix} H_t & \boldsymbol{i}_t \\ \mathbf{0} & 1 \end{bmatrix}$$

where $\mathbf{0}$ denotes a 10-row vector of zeros and $\boldsymbol{i}_t$ a 10-column vector of the probability someone will transition into an immune state from the other states. The individual analysis given above translates seamlessly into a population model where the focus moves onto the relationships between the numbers in the population that might lie in each position at any one time. For the decision support for this agency for a subset of covariates $\boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}$, each possible suite of decisions $d \in D$ and in quarter t the authorities utility function will usually depend upon $\boldsymbol{N}_t(\boldsymbol{x},d) \triangleq (N_{t1}(\boldsymbol{x},d), N_{t2}(\boldsymbol{x},d), \ldots, N_{t10}(\boldsymbol{x},d))$, for $t \geq t_0$ where t_0 is the current time. Let $N_t(\boldsymbol{x},d)$ be the total population of people with the value of covariates $\boldsymbol{x}$ under the decision $d \in D$. Usually $N_t(\boldsymbol{x},d)$ with mean $\mu_t(\boldsymbol{x},d)$ will depend on d only as a response to immigration/emigration policies, replacements and detention orders expressed within the suite d. Let $\boldsymbol{\mu}_t(\boldsymbol{x}, d) \triangleq (\mu_{t1}(\boldsymbol{x}, d), \mu_{t2}(\boldsymbol{x}, d), \ldots, \mu_{t10}(\boldsymbol{x}, d))$. In this context $N_t(\boldsymbol{x}, d)$ and $(\mu_t(\boldsymbol{x}, d))$ will be relatively well known, will have a small variance and can be based on publicly available geographical and demographic information. This means there are simple formulae that approximate well the dependence structure. For example assuming $\widetilde{H}_t$ is a good approximation we note that these proportions s quarters ahead will be given by

$$\boldsymbol{\rho}_{t_0+s}(\boldsymbol{x}, d) = \boldsymbol{\rho}_{t_0}(\boldsymbol{x}, d) H_{t_0}^{(s)}$$

where $H_{t_0}^{(s)} \triangleq \prod_{j=1}^{s} H_{t_0+j}$. In particular under the hypothesis of time homogeneity $H_t = H$, $t \geq t_0$ then $H_{t_0}^{(s)} = H^s$.

When considering the current states data concerning $\boldsymbol{\rho}_t(\boldsymbol{x}, d) \triangleq (\rho_{t1}(\boldsymbol{x}, d), \rho_{t2}(\boldsymbol{x}, d), \ldots, \rho_{t10}(\boldsymbol{x}, d))$ where $\rho_{ti}(\boldsymbol{x}, d) \triangleq \frac{\mu_{ti}(\boldsymbol{x},d)}{\mu_t(\boldsymbol{x},d)}$ is often poor. This is where expert judgments are vital. However these parameters can be inherited from our person centric analysis by rolling forward the transition matrix in the way we described above. These conditional probabilities will of course be uncertain. However many Bayesian models allow this to be taken into account - see the Dirichlet analysis in Appendix 8.1 and see [20, 21] so that this uncertainty appropriately informs the expected utility scores in any analysis.

At a population level, the population process $\{\boldsymbol{N}_t(\boldsymbol{x},d) : t \geq t_0, d \in D, \boldsymbol{x} \in \mathbb{X}' \subseteq \mathbb{X}\}$ can be embedded into a flow graph [54] and then if necessary a 2 time slice DBN [41]. Indeed this is what we did for a much simpler class in [50]. However this final translation is now unnecessary. Having established the formal semantics of the RDCEG the processes can be represented much more transparently directly in most cases.

6.4 Embedding recent Twitter analyses: examples of data accommodation

Although the data concerning the progression of the small proportion of the general public in the early stages of a path towards RVE used to be sparse, publicly available data from social media has provided new indications of both the depth of support for various such organisations and also those specific individuals that might be at risk of being drawn into violence against the general public. A proper review of some of this exciting new work is beyond the scope of this article. So we will focus on just one stream of this work which concerns the use of Twitter. After a brief review of some pertinent observations made by the authors of [63, 64], we reinterpret their findings so that they can be used to inform our particular inferential schemes.

First these authors distil the enormous data concerning Twitter communications between ISIS supporters into three binary measures concerning a particular individual. They label these as Sentiment Tendency (ST), Ego Network Extremism Support (ENES) and Mention Network (MN). ST measures rely on the idea that if a person is an ISIS supporter then they are likely to show positive sentiments towards ISIS whilst others will tend to display neutral or negative sentiment. They therefore count the proportion of positive tweets concerning ISIS. If this proportion is > 0.5 that person is classified as having positive sentiment for ISIS. The second ENES measure is based on an observation that a user high on our alienation scale is likely to be followed by at least one with similar sentiment whilst others like journalists and researchers would not. ENES is an indicator of whether a given person has at least one *follower* with a positive sentiment. Finally the MN measure is based on the idea that extremist tweets are likely to be mentioned by someone high on the alienation scale. If that user also has a high sentiment then the indicator on MN is one. Validation experiments against the results of machine learning techniques simply based on these scores using simple naive Bayes methods - actually one of the more successful and entirely compatible with our approach - achieved around an 80% success against the test set of expertly classified Tweets.

This work is very interesting and demonstrates the promise of this type of data synthesis. However the choice of these statistics have been chosen through common sense rather than the consideration of the possible pathways Tweeters might be on. The measures all reflect where the Tweeter might be currently positioned and so in the notation above inform $\boldsymbol{\rho}_{t_0}$. Here we note that within the semantics developed above the ST feature broadly corresponds to a measure of someone's motivation position $\geq m_2$. Both ENET and MN indicators mentioned above are indicators of embedding levels $\geq e_1$.

How could we use the development above to provide more promising indicators. Well the first suggestion would be to use a second sentiment indicator which not only promoted ISIS but also violence to the general public associated with anything. This would help tighten the motivation score. If it is possible to discover whether the person of interest is geographically close to the follower or the mentioner then this is further evidence that this person is more likely to be embedded.

A second measure would more closely link to transitions. It is well known that longitudinal statistical analyses are far more informative about populations than cross sectional studies. So here - as emphasised by [51] changes in behaviour is more indicative. Notice that someone who begins to retweet ISIS propaganda might have made a transition. Similarly someone who stop Tweeting materials seems to have made a transition. However the transition could be one where they have become immune or alternatively have gone into electronic hiding because they are close to perpetrating a violent incident. Each of these data concerning change helps us to re-estimate transition probabilities perhaps using the estimate of that person's current state. Such information can be used formally provided that we have elicited probabilities of these measured electronic activities. So given current ISIS operandi what are their current instruction concerning electronic communications to those about to attack civilian targets? This will inform the choice of probability. Such information is ephemeral and so needs regular revision. However by assuming the adversary is rational even without internal intelligence it is possible to make predictions of what such instructions might be [61].

Note that these sort of arguments enable the construction of more nuanced statistics which nevertheless admit volume processing and are still open to all the validation methodologies applied to those chosen by [63] because they are constructed through prior out of sample methods.

Finally by processing this type of date in a Bayesian way we can import many of the useful factorisations of a problem that have been found useful in forensic science. Here we will restrict ourselves to the mention of just one such example. In forensic science trace evidence is critical. For example a number of glass fragments are found by a forensic scientist when examining the jacket of a suspect. The question is the strength of this evidence in indicating that this person was at the scene of a crime where a window was broken. It has been strongly argued that to assess the strength of this evidence it is important to consider three questions: - What were the probabilities such glass fragments come from the window broken in the incident or from somewhere else (transfer)? Given other information about the incident concerning the activities of the suspect between the incident, the arrest and retrieval what is the probability that this number of fragments were discovered (persistence) and how many of the fragments can we expect the scientist to retrieve given the number of fragments on the jacket (retrieval)? Asking these questions about the evidence actually found conditional on the person being innocent and then conditional on being guilty, in conjunction with the prior information, can then by Bayes Rule be used to construct the suspect's guilt or innocence. The innocent probabilities are informed by physical models of shattering glass plus sample surveys measuring the number of glass fragments on different types of people, the persistence of glass fragments on jackets of various kinds and on experiments concerning percentage retrieval of fragments known to be on a jacket [4].

Analogous constructions can be crafted in this domain. Consider for example evidence provided by a number of retweets successfully identified as coming form an ISIS source made by a suspect. These incriminating items have been the result of the same three processes. Working backwards these need to have been retrieved by the researcher at the time of the search, these had to have persisted between the time of the original ISIS post and these must have been sourced from this particular place. Each of these processes has its own probabilities for someone who is on each of the positions in our RDCEG above. Just as for forensic evidence, by carefully unpicking each of these aspects in turn we can hope to properly populate the sample distribution of the evidence and hence calculate our new position probabilities for this individual. Detailed numerical illustrations of this process will be reported later.

7 Conclusions

This paper has demonstrated how a simple discrete Bayesian analysis can support the systematic forecasting of assault crimes and the evaluation of suites of policy designed to mitigate these threats. It has been shown how detailed and thoughtful sociological modelling of this area can be. This can then be transformed into a well tuned quantitative tool, which to able encode this expert judgment. It can also be calibrated and adjusted in the light of the types of patchy data that might be available in ways illustrated above. The formal embedding then provides the appropriate inputs to a full Bayesian decision support system. This in turn is able to quickly produce initial scores concerning the efficacy of different candidate policies for discussion, adaptation and adoption by a variety of different agencies. One critical feature of these models is that they are fashioned around constructs that are commonly used within this domain. This means that the rationale behind each evaluation can be fed back to the user in a transparent way for criticism and modification. In this way the analysis can support rather than override these often sensitive decision processes, and decisions can be scrutinised and if adequate, justified to other government agencies as well as the general public. We also illustrated how the granularity of a probability model so constructed could be customised to the needs of a particular agency so that the supporting model was not overly complicated.

There are two critical issues that emerged from this study. First it has been shown that it is essential for informed prior judgments to be drawn into any analysis. Of course such judgments should be informed by as much empirical evidence that can be made available and the Bayesian paradigm guides this accommodation. However for absolutely justified confidentiality and ethical constraints, and indeed the depth of current survey information there are currently many gaps in this empirical evidence base we would ideally use in this domain. Furthermore the interpretation of what data we can collect tends to be ephemeral because of the rationality of the adversary. These conditions are unlikely to change in the near future. It follows that judgments need to be made based on informed reflections of criminologists. Bayesian models are especially useful in this type of domain because they provide an established methodology for systematically addressing such environments and accommodate what evidence we do have into the support system.

Secondly to actually implement the technology we describe here proper suites of template RDCEGs need to be developed in close collaboration with the appropriate agencies. This will obviously take a little time: the templates provided here are meant as simple exemplars of properly elicited structures. And once these structures are elicited these still need to be populated with probabilities. It is hoped that the efficacy of embarking on this activity has the potential of being extremely fruitful and worthwhile.

Here attention has focused on a single type of crime. However it is important to note that wider resource allocation across domains can also be addressed by these models. Tools to perform Bayesian decision analysis to best mitigate crime of many different types can be built in a similar way. To forecast the consequences of allocating resources against one criminal activity to another we can then simply compare their calculated consequences. Provided that there is an agreement as to the utility function to use, a direct application of an RDCEG Bayesian analysis can guide this resource allocation. For example if the aim is to reduce the number of civilian deaths in the life of the current parliament then once the necessary RDCEGs have been elicited and populated with the necessary edge probabilities, this becomes a simple calculation. Of course such probabilities can be contentious as can the appropriate choice of utility function. But at least we have a framework within which rational discussion of the various options can take place.

In 2010 we explored the use of the DBN for evaluating Prevent models [50]. Within this short study it became apparent that an alternative Bayesian model to the DBN that was based on a semi-Markov process would be more consistent and compatible with the elicited expert judgments we needed to make and would provide an even better platform for decision support. This prompted the development of the RDCEG and the work described here.

Of course the RDCEG is only one tool of many to apply to the study of criminal behaviour. It is especially useful when the focus of the study is on the threat posed by individual criminals and the nature and extent of a particular criminal population. However especially when supporting the pursuit of criminal gangs there are at least two other vital components of the processes: the real time networking activities of the gangs and the geographical nature of threats as these develop over time. The population models described in this paper only indirectly use such information which needs to be imported from these two other processes.

So our next task is to report the design specific real time integrating system that draws the three different decision support systems into a single coherent structure, symbiotically allowing information flow between the three systems and harmonising their outputs. An Integrating Decision Support System [9, 42] that provides a single consistent predictive tool to support police operations in real time in their pursuit of violent criminals is now conceived and its implementation is currently being investigated. This work will be reported in the forthcoming paper [3].

References

[1] Agarwal, S. and Sureka, A. “A focused crawler for mining hate and extremism promoting videos on YouTube.” In: *Proceedings of the 25th ACM conference on Hypertext and social media.* ACM. 2014, pp. 294–296.

[2] Agarwal, S. and Sureka, A. “Using knn and svm based one-class classifier for detecting online radicalization on twitter”. In: *International Conference on Distributed Computing and Internet Technology*. Springer. 2015, pp. 431–442.

[3] Aglietti, V., Damoulas, T., and Smith, J. Q. “An integrating Decision Support System for the Pursuit of Violent Crime”. In: ((in preparation)).

[4] Aitken, C. G. and Leese, M. *Statistics and the evaluation of evidence for forensic scientists*. 04; QA276. A2, A5. J. Wiley, 1995.

[5] Banks, D. L., Aliaga, J. M. R., and Insua, D. R. *Adversarial risk analysis*. CRC Press, 2015.

[6] Barclay, L. M., Hutton, J. L., and Smith, J. Q. “Refining a Bayesian network using a chain event graph”. In: *International Journal of Approximate Reasoning* 54.9 (2013), pp. 1300–1309.

[7] Barclay, L. M., Hutton, J. L., Smith, J. Q., et al. “Chain event graphs for informed missingness”. In: *Bayesian Analysis* 9.1 (2014), pp. 53–76.

[8] Barclay, L. M. et al. “The dynamic chain event graph”. In: *Electronic Journal of Statistics* 9.2 (2015), pp. 2130–2169.

[9] Barons, M. J., Wright, S. K., and Smith, J. Q. “Eliciting probabilistic judgements for integrating decision support systems”. In: *Elicitation: The Science and Art of Structuring Judgement*. Ed. by Dias, L. C., Morton, A., and Quigley, J. Springer, 2018. Chap. 17, pp. 445–478.

[10] Bermingham, A. et al. “Combining social network analysis and sentiment analysis to explore the potential for online radicalisation”. In: *Social Network Analysis and Mining, 2009. ASONAM’09. International Conference on Advances in*. IEEE. 2009, pp. 231–236.

[11] Borum, R. “Radicalization into violent extremism I: A review of social science theories”. In: *Journal of Strategic Security* 4.4 (2011), pp. 7–36.

[12] Borum, R. “Radicalization into violent extremism II: A review of conceptual models and empirical research”. In: *Journal of Strategic Security* 4.4 (2011), pp. 37–62.

[13] Bouchard, M., Joffres, K., and Frank, R. “Preliminary analytical considerations in designing a terrorism and extremism online network extractor”. In: *Computational models of complex systems*. Springer, 2014, pp. 171–184.

[14] Brady, H. E., Schlozman, K. L., and Verba, S. “Prospecting for participants: Rational expectations and the recruitment of political activists”. In: *American Political Science Review* 93.1 (1999), pp. 153–168.

[15] Carrabine, E. et al. *Criminology: A sociological introduction*. Routledge, 2009.

[16] Chatfield, A. T., Reddick, C. G., and Brajawidagda, U. "Tweeting propaganda, radicalization and recruitment: Islamic state supporters multisided twitter networks". In: *Proceedings of the 16th Annual International Conference on Digital Government Research*. ACM. 2015, pp. 239–249.

[17] Chen, H. et al. "Uncovering the dark Web: A case study of Jihad on the Web". In: *Journal of the Association for Information Science and Technology* 59.8 (2008), pp. 1347–1359.

[18] Christmann, K. "Preventing religious radicalisation and violent extremism: A systematic review of the research evidence". In: (2012).

[19] Collazo, R. A., Görgen, C., and Smith, J. Q. *Criminology: A sociological introduction*. Chapman & Hall, 2018.

[20] Collazo, R. A. and Smith, J. Q. "An N Time-Slice Dynamic Chain Event Graph". In: *CRiSM Res. Rep.* (2018).

[21] Collazo, R. A. and Smith, J. Q. "The Dynamic Chain Event Graph". In: *Proceedings of ISI Marrakech* ((to appear)).

[22] Cook, R. et al. "A hierarchy of propositions: deciding which level to address in casework". In: *Science & Justice* 38.4 (1998), pp. 231–239.

[23] Cowell, R. et al. *Probabilistic networks and expert systems*. Springer, 1999.

[24] Cowell, R. G., Smith, J. Q., et al. "Causal discovery through MAP selection of stratified chain event graphs". In: *Electronic Journal of Statistics* 8.1 (2014), pp. 965–997.

[25] Crossett, C. and Spitaletta, J. A. *Radicalization: Relevant psychological and sociological concepts*. Ft. Meade, MD, U.S. Army Asymmetric Warfare Group, 2010.

[26] Dawid, A. P. and Evett, I. W. "Using a graphical method to assist the evaluation of complicated patterns of evidence". In: *Journal of Forensic Science* 42.2 (1997), pp. 226–231.

[27] Dawid, A. P. and Mortera, J. "Graphical models for forensic analysis". In: *Handbook of Graphical Models*. Ed. by Drton, M. et al. CRC Press, (to appear).

[28] Della Porta, D. *Social movements, political violence, and the state: A comparative analysis of Italy and Germany*. Cambridge University Press, 2006.

[29] Department, A. .-.-. G. *Preventing Violent Extremism and Radicalisation in Australia*. Report GTReC, Monash University & Australian Government, 2015.

[30] Ferrara, E. et al. "Predicting online extremism, content adopters, and interaction reciprocity". In: *International conference on social informatics*. Springer. 2016, pp. 22–39.

[31] Freeman, G. and Smith, J. Q. "Bayesian MAP model selection of chain event graphs". In: *Journal of Multivariate Analysis* 102.7 (2011), pp. 1152–1165.

[32] Freeman, G., Smith, J. Q., et al. "Dynamic staged trees for discrete multivariate time series: forecasting, model selection and causal analysis". In: *Bayesian Analysis* 6.2 (2011), pp. 279–305.

[33] Glynn, A. N. et al. "Alleviating linear ecological bias and optimal design with subsample data". In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171.1 (2008), pp. 179–202.

[34] Harcourt, B. E. *Against prediction: Profiling, policing, and punishing in an actuarial age*. University of Chicago Press, 2008.

[35] Hegghammer, T. "Terrorist recruitment and radicalization in Saudi Arabia". In: *Middle East Policy* 13.4 (2006), p. 39.

[36] Horgan, J. "From profiles to pathways and roots to routes: Perspectives from psychology on radicalization into terrorism". In: *The ANNALS of the American Academy of Political and Social Science* 618.1 (2008), pp. 80–94.

[37] Howard, R. A. "Decision analysis: practice and promise". In: *Management science* 34.6 (1988), pp. 679–695.

[38] Imrey, P. B. and Dawid, A. P. "A commentary on statistical assessment of violence recidivism risk". In: *Statistics and Public Policy* 2.1 (2015), pp. 1–18.

[39] Johnson, N. F. et al. "New online ecology of adversarial aggregates: ISIS and beyond". In: *Science* 352.6292 (2016), pp. 1459–1463.

[40] Karampelas, P. and Bourlai, T. *Surveillance in Action: Technologies for Civilian, Military and Cyber Surveillance*. Springer, 2017.

[41] Korb, K. B. and Nicholson, A. E. *Bayesian artificial intelligence*. CRC press, 2010.

[42] Leonelli, M. and Smith, J. Q. "Bayesian decision support for complex systems with many distributed experts". In: *Annals of Operations Research* 235.1 (2015), pp. 517–542.

[43] McCauley, C. and Moskalenko, S. "Mechanisms of political radicalization: Pathways toward terrorism". In: *Terrorism and political violence* 20.3 (2008), pp. 415–433.

[44] McCormick, G. H. "Terrorist decision making". In: *Annual Review of Political Science* 6.1 (2003), pp. 473–507.

[45] Mortera, J. and Dawid, A. P. "Forensic identification then and now". In: 27.2 (2016), pp. 145–172.

[46] Office, H. *The Prevent Strategy: A Guide for Local Partners in England - Stopping People becoming or Supporting Terrorists and Violent Extremism*. 2008.

[47] Pressman, D. E. "Risk assessment decisions for violent political extremism". In: (2009).

[48] Qi, X. et al. “A hierarchical algorithm for clustering extremist web pages”. In: *Advances in Social Networks Analysis and Mining (ASONAM), 2010 International Conference on*. IEEE. 2010, pp. 458–463.

[49] Rezaei, S. and Goli, M. *House of War: Islamic Radicalisation in Denmark*. Centre for Studies in Islamism and Radicalization (CIR), Aarhus University, 2010.

[50] Rigat, F. and Smith, J. Q. *Study of Key Interventions into Terrorism using Bayesian Networks*. (SRG/09/44) Home Office Final Report (44 pages), 2011.

[51] Rowe, M. and Saif, H. “Mining Pro-ISIS Radicalisation Signals from Social Media Users.” In: *ICWSM*. 2016, pp. 329–338.

[52] Shenvi, A. and Smith, J. Q. “Elicitation, estimation and model selection for the reduced dynamic chain event graph”. In: ((in preparation)).

[53] Silber, M. D., Bhatt, A., and Analysts, S. I. *Radicalization in the West: The homegrown threat*. Police Department New York, 2007.

[54] Smith, J. Q. and Figueroa, L. J. “A causal algebra for dynamic flow networks”. In: *Advances in Probabilistic Graphical Models*. Springer, 2007, pp. 39–54.

[55] Smith, J. Q. et al. “A Structured Bayesian Decision Support System to Pursue Radicalised Violent Extremists”. In: ((in preparation)).

[56] Smith, J. Q. *Bayesian decision analysis: principles and practice*. Cambridge University Press, 2010.

[57] Smith, J. Q. and Anderson, P. E. “Conditional independence and chain event graphs”. In: *Artificial Intelligence* 172.1 (2008), pp. 42–68.

[58] Smith, J. Q. and Queen, C. M. “Bayesian models for sparse probability tables”. In: *The Annals of statistics* (1996), pp. 2178–2198.

[59] Smith, J. “The prediction of prison riots”. In: *British Journal of Mathematical and Statistical Psychology* 33.2 (1980), pp. 151–160.

[60] Sureka, A. and Agarwal, S. “Learning to classify hate and extremism promoting tweets”. In: *Intelligence and Security Informatics Conference (JISIC), 2014 IEEE Joint*. IEEE. 2014, pp. 320–320.

[61] Thwaites, P. A. and Smith, J. Q. “A Graphical method for simplifying Bayesian Games”. In: *Reliability Engineering & System Safety* (2017).

[62] Venhaus, J. M. *Why youth join al-Qaeda*. US Institute of Peace, 2010.

[63] Wei, Y. and Singh, L. “Detecting users who share extremist content on twitter”. In: *Surveillance in Action*. Springer, 2018, pp. 351–368.

[64] Wei, Y. and Singh, L. “Using network flows to identify users sharing extremist content on social media”. In: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer. 2017, pp. 330–342.

[65] Wiktorowicz, Q. *Radical Islam rising: Muslim extremism in the West*. Rowman & Littlefield Publishers, 2005.

[66] Zald, M. N. and McCarthy, J. D. "Social Movements in an Organizational Society. Transaction". In: *Inc.: New Jersey* (1987).

8 Appendix

8.1 Furnishing the topology of an RDCEG with probabilities

To embellish an infinite event tree into a probability tree which fully specifies the evolution of a unit in the described population we need to elicit only two components from the domain expert:

1. For each situation $v \in S(\mathcal{T})$ a distribution for the *holding time* $T(v,t)$, i.e. the time the unit stays in that situation before moving to one of its children after arriving at v at time t.
2. the *conditional probability vector* (cpv) $\boldsymbol{\pi}_j(t)$ associated with one *representative* floret $\mathcal{F}(v_{ij})$, $v_{ij} \in S(\mathcal{T})$, $v_{ij} \in u_j$ from each stage $u_j \in U$ of the elicited staged tree. This vector specifies the probabilities of the next step of a unit from v_{ij} to one of its children as a function of the time t it arrived at v_{ij}. Note that by definition of a stage, $\boldsymbol{\pi}_j(t)$ can depend on the index j of the stage but not the index i of the representative situation.

We then start to colour the tree. Often situations which share characteristics of escalation but with different background covariates will be identified into the same stage.

The time taken for transitions from one state to the next will typically be slow but will accelerate up the tree. So we have a genuine semi-Markov process here and what is technically known as an RDCEG. This holds for both known and unknown individuals.

The stages $\{u_i : 0 \leq i \leq k\}$ are simply the non-sink superpositions. For each component of $\boldsymbol{\pi}_i$ the transition pobabilities $\boldsymbol{\pi}$ are defined by π_i, $i = 0, 1, \ldots, 5$. As time advances on each of the units and their progress evolves, the posterior distributions of the probabilities $\boldsymbol{\pi}_i$ associated with stage u_i can be summarized by k_i parameters $1 \leq i \leq k$, which can be expressed as the posterior means $\alpha^+_{ij}/\sum_{j'} \alpha^+_{ij'}$ of the probabilities for each of the emanating edges and an effective sample size measure $\sum_{j'} \alpha^+_{ij'}$ of the accuracy of these assessments (see [19, 20]). Thus setting $\boldsymbol{\pi}_1 \amalg \boldsymbol{\pi}_2, \ldots \amalg \boldsymbol{\pi}_k$ and assigning a Dirichlet $D(\boldsymbol{\alpha}^0_j(\boldsymbol{\pi}_i))$ prior density to each where $\boldsymbol{\alpha}^0_j = (\alpha^0_1(\boldsymbol{\pi}_i), \alpha^0_2(\boldsymbol{\pi}_i), \ldots, \alpha^0_{k_i}(\boldsymbol{\pi}_i))$ a priori, when $\sum_{j=1}^{k_i} \pi_{ij} = 1$, $\pi_{ij} > 0$, $1 \leq j \leq k_i$, we have,

$$p(\boldsymbol{\pi}_i) = \frac{\Gamma(\alpha^0_1(\boldsymbol{\pi}_i) + \alpha^0_2(\boldsymbol{\pi}_i) + \ldots + \alpha^0_{k_i}(\boldsymbol{\pi}_i))}{\Gamma(\alpha^0_1(\boldsymbol{\pi}_i)) \ldots \Gamma(\alpha^0_{k_i}(\boldsymbol{\pi}_i))} \prod_{j=1}^{k_i} \pi_{ij}^{\alpha^0_j(\boldsymbol{\pi}_i)}$$

on each component $\boldsymbol{\pi}_i$, $i = 1, 2, \ldots, k$ gives a conjugate analysis. Here the separation of the likelihood above gives $\amalg_i \boldsymbol{\pi}_i | \boldsymbol{x}$ a posteriori. Each of these stage parameters $\boldsymbol{\pi}_i$ also have a Dirichlet distribution where the hyperparameters $\boldsymbol{\alpha}_i^+(\boldsymbol{\pi}_i)$ of the posterior density are linked to the corresponding prior hyperparameters $\boldsymbol{\alpha}_i^0(\boldsymbol{\pi}_i)$ by the linear equation

$$\boldsymbol{\alpha}_i^+(\boldsymbol{\pi}_i) = \boldsymbol{\alpha}_i^0(\boldsymbol{\pi}_i) + \boldsymbol{n}_i$$

where $\boldsymbol{n}_i = (n_{i,1}, n_{i,2}, \ldots, n_{i,k_i})$ which is the vector of the number of incidents $n_{i,r}$ when a unit arriving at i then passes along the edge r, $r = 1, 2, \ldots, k_i$. Note that in particular if we are comparing model classes of $\mathcal{C}$ a fast model selection algorithm is available which uses the sum of the log Bayes Factor as a score function to find the MAP model in this class.

Conjugate learning in non-dynamic CEG's can accommodate not only sampling schemes but also causal experimental data and is now well documented [19, 31, 56]. Learning schemes can be devised - even if these probabilities are believed to drift in time often within a closed form analysis [31, 32]. These methods and the formulae closely resemble analogous learning in discrete Bayesian Networks under full sampling of the net [23, 56]. For estimation of transition times which can be undertaken independently within this model class see [8].