Appendix

Appendix A

Implementation in WinBUGS

A.1 Additive model, fixed location of m latent kernels

WinBUGS code for Poisson–Gamma model assuming an additive influence of benzene and a fixed number of latent risk sources at fixed locations.

```
model
{
###########################################################################
# Data
#=========================================================================
# Variable      Description
#----------    -----------------
#  benzene     benzene emission in the area each ward covers
#  E           population numbers (vector of length I)
#  count       observed number of cases (vector of length I)
#  J           number of grid cells (benzene emissions)
#  I           number of wards
#  Source      number of latent covariates involved
#  Sx.grid     coordinates of the grid cells used for evaluation of latent
#              risk factors (vector of length nx.grid)
#  Sy.grid     similar for y direction
#  nx.grid     number of grid cells in x direction
#  ny.grid     similar for y direction
#  Sx.source   location of the latent covariates (vector of length nx.source)
#  Sy.source   similar for y direction
```

```
#  nx.source   number of sources in x direction
#  ny.sorces   number of sources in y direction
#  X.coords    Index of grid cell the ward center belongs to
#  Y.coords    similar for y direction
#  S.index     nx.source times ny.source matrix giving index number (1:Source)
#              of latent grid cell
#  area        area of region in km^2.
#############################################################################
# Parameters
#============================================================================
# required constants
#===================
 expect <- mean(E[])

# Priors for additive risk factors
#=================================
 beta.0      ~ dgamma(a.0     , tau.0)        # prior for intercept
 beta.benz   ~ dgamma(a.benz  , tau.benz)     # prior for benzene
 beta.latent ~ dgamma(a.latent, tau.latent)   # prior for latent coefficient
 a.0 <- 0.575                                 # shape for intercept
 tau.0 <-  a.0 * 3 * expect                   # scale for intercept
 a.benz <- 0.575                              # shape for benzene
 tau.benz <- a.benz * 3 * expect              # scale for benzene
 a.latent <- 0.575                            # shape for latent coefficient
 tau.latent <- a.latent * 3 * expect          # scale for latent coefficient

# no multiplikative risk factors
#===============================
# Priors for gamma[m]'s
#======================
# magnitudes of the latent risk factors

 for (s in 1:Source)
    {
    gamma[s] ~ dgamma(a.gamma, tau.gamma)
    }
 a.gamma <- area * tau.gamma
 tau.gamma <- 1/Source

# Kernel: calculate kernel in WinBUGS to allow rho to be uncertain
#=================================================================
 logrho ~ dnorm(0,2)
 rho <- exp(logrho)
```

```
 for(sx in 1:nx.source)              # loop over nx.sources
    {
    dFx[sx,1:nx.grid]<- eval.grid(Sx.grid[], Sx.source[sx], rho)
    # eval.grid() is a Black Box function. Code is given below.
    }
 for(sy in 1:ny.source)
    {
    dFy[sy, 1:ny.grid] <- eval.grid(Sy.grid[], Sy.source[sy], rho)
    }
 # calculates kerX[ward.number, source.number]
 # X.coords contains the number of the grid cell
 # (grid for evaluation of latent risk factors)
 # in which the ward center is situated
 # (similar for Y.coords)
 for(sx in 1:nx.source)                      # loop over sources
    {
    kerX[sx,1:I] <- belong(dFx[sx, ], X.coords[1:I])
    # belong() is a Black Box function. Code is given below.
    }
 for(sy in 1:ny.source)                      # loop over sources
    {
    kerY[sy,1:I] <- belong(dFy[sy, ], Y.coords[1:I])
    }
 # merge together for 2 dimensional risk surface
 for(i in 1:I)                                          # loop over areas
     {
     for(sx in 1:nx.source)
        {
        for(sy in 1:ny.source)
            {
            kernel[S.index[sx,sy] ,i] <- kerY[sy,i] * kerX[sx,i] * 1000
            }
        }
     }

# Intensities
#=============
 for (i in 1:I)
    {
    count[i] ~ dpois(lambda[i])
    lambda[i] <- p[i]*E[i]
    benz.term[i] <- beta.benz*benzene[i]
    for (s in 1:Source)
```

```
        {
        latent[s,i] <- delta[s]*kernel[s,i]
        }
    latent.term[i] <- beta.latent*sum(latent[,i])
    p[i]<- beta.0 + benz.term[i] + latent.term[i]
    }

#expected number of cases attributed to each source
#=================================================
 E.base <- beta.0
 E.benzene <- mean(benz.term[])
 for (i in 1:I)
    {
    latentsum[i] <- sum(latent[,i])
    }
 E.latent <- beta.latent*mean(latentsum[])
 Total <- E.base + E.benzene + E.latent
 percent.base <- E.base / Total *100
 percent.benzene <- E.benzene / Total * 100
 percent.latent <- E.latent / Total * 100
}
```

A.2 Black Box function `eval.grid()`

Function to calculate the intensity of a latent risk source `source` in grid cell `grid`.

```
MODULE WBDevEvalGrid;
IMPORT
   WBDevVector, WBDevSpecfunc, Math;
TYPE
   Function = POINTER TO RECORD (WBDevVector.Node) END;
   Factory = POINTER TO RECORD (WBDevVector.Factory) END;
VAR
   fact-: WBDevVector.Factory;
   PROCEDURE (func: Function) DeclareArgTypes (OUT args: ARRAY OF CHAR);
   BEGIN
       args := "vss";
   END DeclareArgTypes;
   PROCEDURE (func: Function) Evaluate (OUT dFx: ARRAY OF REAL);
   CONST
```

```
15          grid = 0; source = 1; rho = 2;
16      VAR
17          numGrid, i, j: INTEGER;
18      val: REAL;
19      Fx: ARRAY 1000 OF REAL;
20      BEGIN
21      numGrid :=LEN(func.arguments[grid]);
22      i:=0;
23      WHILE i < numGrid DO;
24        val :=0.001*(func.arguments[grid][i].Value()-
25                    func.arguments[source][0].Value());
26        val:= val/func.arguments[rho][0].Value();
27        Fx[i] := WBDevSpecfunc.Phi(val);
28        INC(i);
29      END;
30      j:=0;
31      WHILE j < (numGrid-1) DO;
32      dFx[j] :=Fx[j+1] - Fx[j] ;
33      INC(j);
34      END;
35      END Evaluate;
36
37  PROCEDURE (f: Factory) New (option: INTEGER): Function;
38  VAR
39     func: Function;
40  BEGIN
41     NEW(func); func.Initialize; RETURN func;
42  END New;
43
44  PROCEDURE Install*;
45  BEGIN
46     WBDevVector.Install(fact);
47  END Install;
48
49  PROCEDURE Init;
50  VAR
51     f: Factory;
52  BEGIN
53     NEW(f); fact := f;
54  END Init;
55  BEGIN
56    Init;
57    END WBDevEvalGrid.
```

A.3 Black Box function `belong()`

To ease and to speed up calculations in WinBUGS, the function `belong()` is written. This function can be replaced by the lines

```
for(i in 1:I)
 {
 for(sx in 1:nx.source)
   {
   kerX[sx,i] <- dFx[sx, X.coords[i] ]
   }
 }
```

in the WinBUGS code.

```
MODULE WBDevBelong;
IMPORT
   WBDevVector,
       Math;
TYPE
   Function = POINTER TO RECORD (WBDevVector.Node) END;
   Factory = POINTER TO RECORD (WBDevVector.Factory) END;
VAR
   fact-: WBDevVector.Factory;
   PROCEDURE (func: Function) DeclareArgTypes (OUT args: ARRAY OF CHAR);
   BEGIN
       args := "vv";
   END DeclareArgTypes;
   PROCEDURE (func: Function) Evaluate (OUT val: ARRAY OF REAL);
   CONST
       dFx = 0; Xcoords = 1;
   VAR
       numTimes, i, a: LONGINT;
   BEGIN
   numTimes :=LEN(func.arguments[dFx]);
   i := 0;
       WHILE i < numTimes DO;
   a:=ENTIER(func.arguments[1][i].Value())-1;
   val[i]:= func.arguments[0][a].Value();
   INC(i);
   END;
```

```
   END Evaluate;

PROCEDURE (f: Factory) New (option: INTEGER): Function;
VAR
   func: Function;
BEGIN
   NEW(func); func.Initialize; RETURN func;
END New;

PROCEDURE Install*;
BEGIN
   WBDevVector.Install(fact);
END Install;

PROCEDURE Init;
VAR
   f: Factory;
BEGIN
   NEW(f); fact := f;
END Init;

BEGIN
  Init;
  END WBDevBelong.
```

A.4 Multiplicative model, random location of m latent kernels

```
model
{
#############################################################################
# Parameters
#============================================================================
# required constants
#===================
 expect <- mean(E[])

# no additive risk factors
#=========================

# multiplikative risk factors
```

```
#============================
 beta.0      ~ dgamma(a.0     , tau.0)         # prior for intercept
 beta.benz   ~ dgamma(a.benz  , tau.benz)      # prior for benzene
 beta.latent ~ dgamma(a.latent, tau.latent)    # prior for latent coefficient
 a.0 <- 0.575                                  # shape parameter for intercept
 tau.0<-  a.0 *3 * expect                      # scale for intercept
 a.benz<- 0.575                                # shape for benzene
 tau.benz <- a.benz * 3 * expect               # scale for benzene
 a.latent<- 0.575                              # shape for latent coefficient
 tau.latent <- a.benz * 3 * expect             # scale for latent coefficient

#Priors for gamma[m]'s
#======================
 for (s in 1:Source)
    {
    delta[s] ~ dgamma(a.gamma, tau.gamma)
    }
 a.gamma <- area * tau.gamma
 tau.gamma <- 1/Source

#Location of the Sources: allow uncertainty
#=========================================
 for(i in 1:Source)
    {
    moveX[i] ~ dunif(dist1[i], dist2[i])
    moveY[i] ~ dunif(dist3[i], dist4[i])
    }
 Sx.sourceMove[1:Source] <- Add(moveX[], Sx.source[])
 Sy.sourceMove[1:Source] <- Add(moveY[], Sy.source[])
 # Add() is a Black Box function. Code is given below.

#Kernel: calculate kernel in WinBUGS to allow rho to be uncertain
#================================================================
 for (sx in 1:Source)                  #loop over areas
    {
    for (i in 1:I)
        {
        distanceX[sx, i] <- abs(wardXcenter[i] - Sx.sourceMove[sx])
        distanceY[sx,i]  <- abs(wardYcenter[i] - Sy.sourceMove[sx])
        kernel[sx,i]     <- exp(-(pow(distanceX[sx,i]/(2*rhoX[sx]), 2) +
                                  pow(distanceY[sx,i]/(2*rhoY[sx]), 2)) )
        }
    logrhoY[sx] ~ dnorm(0,wert)
```

```
57      logrhoX[sx] ~ dnorm(0,wert)
58      rhoY[sx] <- exp(logrhoY[sx])
59      rhoX[sx] <- exp(logrhoX[sx])
60      }
61  #Intensites
62  #==========
63   for (i in 1:I)
64      {
65      count[i]  ~ dpois(lambda[i])
66      lambda[i] <- p[i] * E[i]
67      benz.term[i] <- beta.benz * benzene[i]
68      latent[i] <- inprod2(delta[], kernel[,i])
69      latent.term[i] <- beta.latent * latent[i]
70      p[i]<- (beta.0 + latent.term[i]) * exp( benz.term[i])
71      }
72  }
```

A.5 Black Box function `Add()`

The WinBUGS code given in Section A.4 uses in lines 43-44 a vectorised version of the following lines:

```
for(i in 1:Source)
{
Sx.sourceMove[i] <- Sx.source[i] + moveX[i]
}
```

This procedure speeds up calculations.

```
1   MODULE WBDevAdd;
2   IMPORT
3      WBDevVector, WBDevRandnum,
4         Math;
5   TYPE
6      Function = POINTER TO RECORD (WBDevVector.Node) END;
7      Factory = POINTER TO RECORD (WBDevVector.Factory) END;
8   VAR
9      fact-: WBDevVector.Factory;
10     PROCEDURE (func: Function) DeclareArgTypes (OUT args: ARRAY OF CHAR);
11     BEGIN
12         args := "vv";
```

```
13      END DeclareArgTypes;
14
15      PROCEDURE (func: Function) Evaluate (OUT SxsourceMove: ARRAY OF REAL);
16      CONST
17          move=0; Sxsource=1;
18      VAR
19          Number, i: INTEGER;
20      BEGIN
21          Number := LEN(func.arguments[Sxsource]);
22          i := 0;
23          WHILE i < Number DO;
24          SxsourceMove[i] := func.arguments[Sxsource][i].Value()
25                                  + func.arguments[move][i].Value();
26              INC(i);
27          END;
28      END Evaluate;
29
30  PROCEDURE (f: Factory) New (option: INTEGER): Function;
31  VAR
32      func: Function;
33  BEGIN
34      NEW(func); func.Initialize; RETURN func;
35  END New;
36
37  PROCEDURE Install*;
38  BEGIN
39      WBDevVector.Install(fact);
40  END Install;
41
42  PROCEDURE Init;
43  VAR
44      f: Factory;
45  BEGIN
46       NEW(f); fact := f;
47  END Init;
48
49  BEGIN
50  Init;
51  END WBDevAdd.
```

Appendix B

Additional simulation results

In this Chapter we give the result of Poisson–Gamma random field models in the implementation discussed in Section 3.4.3. Selected structures are already discussed in Chapter 8, here we discuss the following structures

- structure A in Section B.1;
- structure B in Section B.2;
- structure D in Section B.3;
- structure F in Section B.4;
- structure G in Section B.5;
- structure H in Section B.6;
- structure J in Section B.7;
- structure M in Section B.8;
- structure N in Section B.9;
- structure P in Section B.10;
- structure Q in Section B.11;

♯ latent factors	0	1	2	3	4	5
model f	339.9 (0.024)	341.7 (0.026)	341.6 (0.027)	342.3 (0.028)	—	—
model m	347.5 (0.049)	349.3 (0.051)	349.6 (0.053)	349.0 (0.053)	348.4 (0.050)	387.0 (0.309)
model o	—	363.2 (0.155)	364.9 (0.150)	365.3 (0.150)	365.5 (0.150)	—
model y	373.2 (0.165)					
model v	353.3 (0.050)					
model z	351.2 (0.050)					

Table B.1: DIC (MSE) values for extended models applied to structure A.

- structure R in Section B.12;
- structure T in Section B.13;
- structure U in Section B.14
- structure V in Section B.15.

For a schematic overview on the used letters see page v.

B.1 Structure A

Data generated according to structure A assumes an additive influence of benzene. In total, 330 observations are generated. Table B.1 reports the calculated DICs and corresponding MSEs for the employed models.

When including benzene in modelling by Poisson–Gamma models, we favour an additive influence based on DIC values. This corresponds to the underlying structure. While the best fitting model of class m, which is the one without any latent risk sources, leads to a DIC of 347.5, model Af0 assuming an additive influence of benzene improves this DIC to be 339.9, see Table B.1. An illustration of the modelled risk surface is given in Figure B.1. A

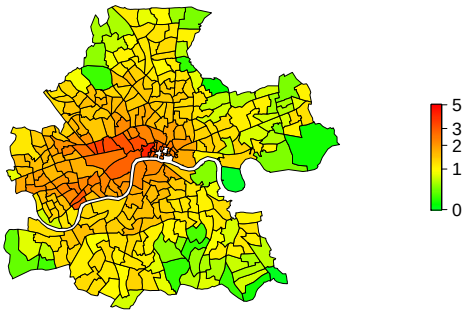

(a: simulated structure)

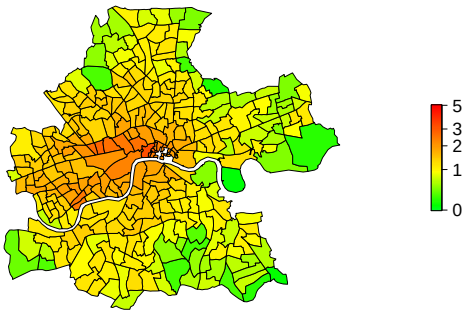

(b: Af0)

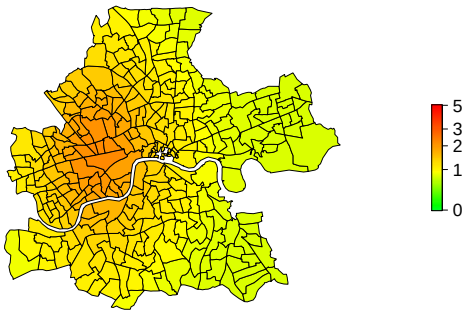

(c: Ao1)

Figure B.1: Simulated Λ_i for structure A (a) and estimated spatial pattern $\widehat{\Lambda}_i$ of the best fitting model Af0 (b), as well as the Poisson–Gamma model with one latent risk source Ao1 (c).

close agreement between the underlying structure and the modelled values is observed. This is reflected by a small MSE of 0.024.

Performance of Poisson–Gamma models not including benzene is inferior to Af0. Here, we calculate DIC values larger than 360, see Table B.1. As presented in Figure B.1 (c) this is caused by larger differences between the parameters of the Poisson distribution Λ_i and the estimated values $\widehat{\Lambda}_i$, which occur along the whole interval. Ao1 is characterised by a smooth risk with a high risk region corresponding to the area with increased benzene observations.

Alternative models show similar performances as those Poisson–Gamma mod-

♯ latent factors	0	1	2	3	4
model f	312.2 (0.016)	313.4 (0.020)	313.8 (0.020)	313.9 (0.020)	313.9 (0.020)
model m	326.8 (0.135)	327.7 (0.136)	328.3 (0.135)	328.2 (0.133)	327.1 (0.136)
model o	—	362.0 (0.494)	366.0 (0.477)	364.9 (0.470)	368.2 (0.440)
model y	372.4 (0.537)				
model v	329.6 (0.151)				
model z	327.9 (0.143)				

Table B.2: DIC (MSE) values for extended models applied to structure B.

els including benzene multiplicatively. Both MRF models gain a DIC of around 350, different neighbourhood structures have no effect on the model fit. The BDCD model shows an inferior performance for this structure reflected by a DIC of 373.2.

For this structure, the best model equals the underlying pattern in data generation.

B.2 Structure B

For structure B, we generate 770 cases assuming an additive influence of benzene only. Table B.2 reports the calculated MSEs and corresponding DICs.

Among Poisson–Gamma models, the one assuming an additive influence of benzene leads to best results among all applied models. Here, the model which does not include any latent risk factors has the lowest MSE of 0.016 and a DIC of 312.2. The corresponding spatial risk surface is given in Figure B.2 (b).

Those DIC results are elevated by about 14 points for model class Bm where a multiplicative influence of benzene is assumed. DIC values are stable for

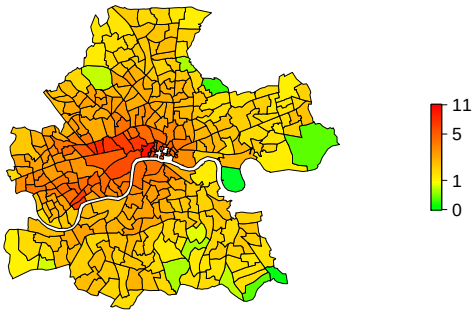

(a: simulated structure)

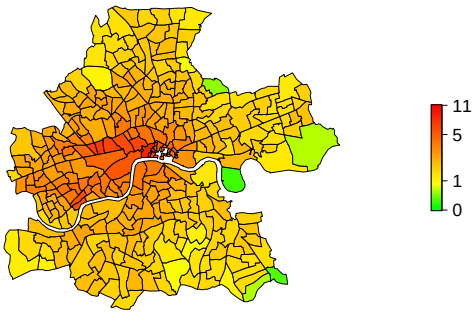

(b: Bf0)

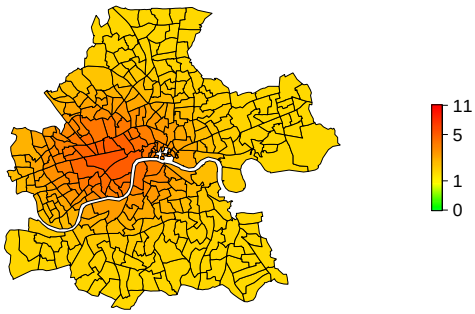

(c: Bo1)

Figure B.2: Simulated Λ_i for structure B (a), estimated spatial pattern $\widehat{\Lambda}_i$ of the best fitting model Bf0 (b) and the best among Poisson–Gamma models not including benzene Bo1 (c).

all numbers of incorporated latent risk sources. Nevertheless, model Bm0 not including additional latent covariates achieves best results. When comparing models with three and four covariates — Bm3 and Bm4 — we see an improvement of Bm4 over Bm3. Nevertheless, the DIC is still higher than the one of Bm0, so we stop adding latent covariates.

If benzene is not included in the Poisson–Gamma model the obtained model fit is inferior. This holds for up to 4 latent risk factors. As we do not observe any convergence of the DIC towards lower values, we do not increase the number of latent risk sources any further. Even if the model fit is inferior, we point out that a MSE of less than 0.500 which is achieved by all of those models is a good match. A plot of the spatial risk $\widehat{\Lambda}_i$ for model Bo1 is given

♯ latent factors	0	1	2	3	4
model f	583.9 (5.767)	326.8 (0.197)	327.3 (0.230)	—	—
model m	618.5 (6.177)	358.0 (0.564)	360.7 (0.511)	361.8 (0.495)	—
model o	—	401.8 (1.175)	399.5 (0.928)	366.6 (0.721)	374.8 (0.822)
model y	395.5 (2.003)				
model v	361.0 (0.779)				
model z	368.4 (0.823)				

Table B.3: DIC (MSE) values for extended models applied to structure D.

in B.2 (c). Model Bo1 gives a very smooth representation of the benzene term involved in generation. The MSE of 0.494 reflects a reasonable model fit as main characteristics of the structure are reproduced.

Other models do not reach the DICs we obtained for the best model Bf0. While both MRF models achieve a similar model fit compared to the Poisson–Gamma models assuming a multiplicative influence of benzene, the BDCD algorithm produces the least appropriate fit for the data set with a DIC of 372.4.

Similar as in structure A, the best model equals the underlying pattern in data generation for structure B.

B.3 Structure D

Data generated by structure D is characterised by an additive influence of benzene accounting for about 770 cases and a latent risk source that accounts for further 330 cases. The spatial pattern is presented in Figure B.3 (a).

DIC and MSE values which are achieved by the chosen models are reported in Table B.3. For Poisson–Gamma models including a benzene term we observe an immense modelling benefit when including a latent risk term in

comparison to modelling benzene alone, although this benefit is not as big as for structure C. The inclusion of more latent covariates neither improves the DIC of the additive nor the multiplicative model. Additive modelling is to favour over multiplicative modelling which reflects the generation scheme. Whereas DIC is 326.8 for model Df1, multiplicative modelling of benzene leads to values of DIC = 358.0 (Dm1) or higher, which is similar to those values we gain when benzene is excluded. Model Do3 which includes three latent covariates performs best among models of class Do. The corresponding risk surface is given in Figure B.3 (c). Compared to the underlying pattern, this model has difficulties to estimate low-risk regions correctly, leading to a DIC value of 366.6 and MSE of 0.721 in comparison to those values we estimate for model Df1.

Similar DICs are estimated by MRF models without a remarkable influence of which neighbourhood structure is applied. We get a DIC of 361.0 for model Dv and a DIC of 368.4 for model Dz, see Table B.3. We conclude comparability of both approaches. Furthermore, for this structure both, the log-link MRF model and the multiplicative Poisson–Gamma model produce comparable pattern. The possibility of Poisson–Gamma models to treat benzene as an excess risk factor by the identity link reveals the advantage of this model class.

The BDCD model performs worse for this structure. The calculated MSE of 2.003 which is more than the double amount of those MSEs calculated for other models. The same holds for the corresponding DIC which is 395.5. We conclude that the BDCD model is able to reproduce the main characteristics of the underlying structure but Poisson–Gamma models and MRF models are to prefer.

All together, we found Df1 to be the best fitting model. A plot of the spatial risk surface is given in Figure B.3 (b). Estimated parameters of the Poisson distribution are close to the generated ones, which is reflected by a small MSE of 0.197. Location and variance of the Gaussian kernel representing latent risk in data generation is estimated well by the model.

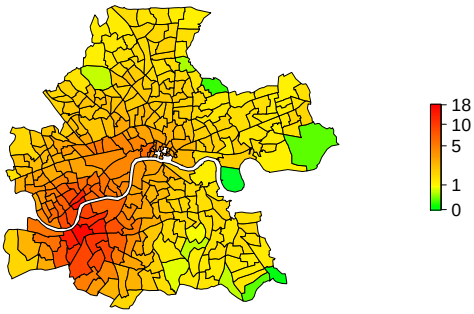

(a: simulated structure)

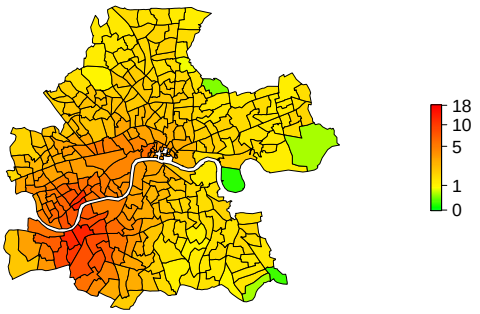

(b: Df1)

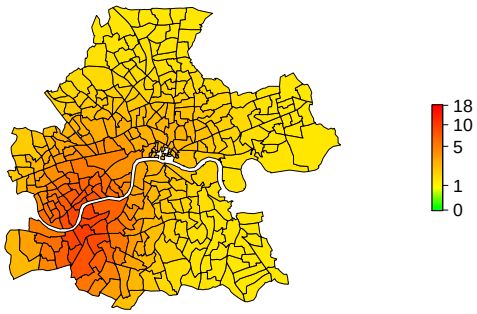

(c: Do3)

Figure B.3: Simulated Λ_i for structure D (a) and results: Estimated spatial pattern of $\widehat{\Lambda}_i$ for the best fitting model Df1 (b), and of the best one without benzene Do3 (c).

B.4 Structure F

Data generation according to structure F assumes an additive influence of benzene (770 cases) and a linearly decreasing trend component (330 cases). A plot of the spatial structure of generated Λ_i is given in Figure B.4 (b).

Modelling these data by Poisson–Gamma models involving benzene additively leads to a DIC of 330.425 (compare Table B.4) when involving five latent covariates which is the best model fit in this class. This is slightly improved over 330.378 when using six latent covariates. Further extension does not lower the DIC. We therefore assume model Ff5 to be the most appropriate one for this structure.

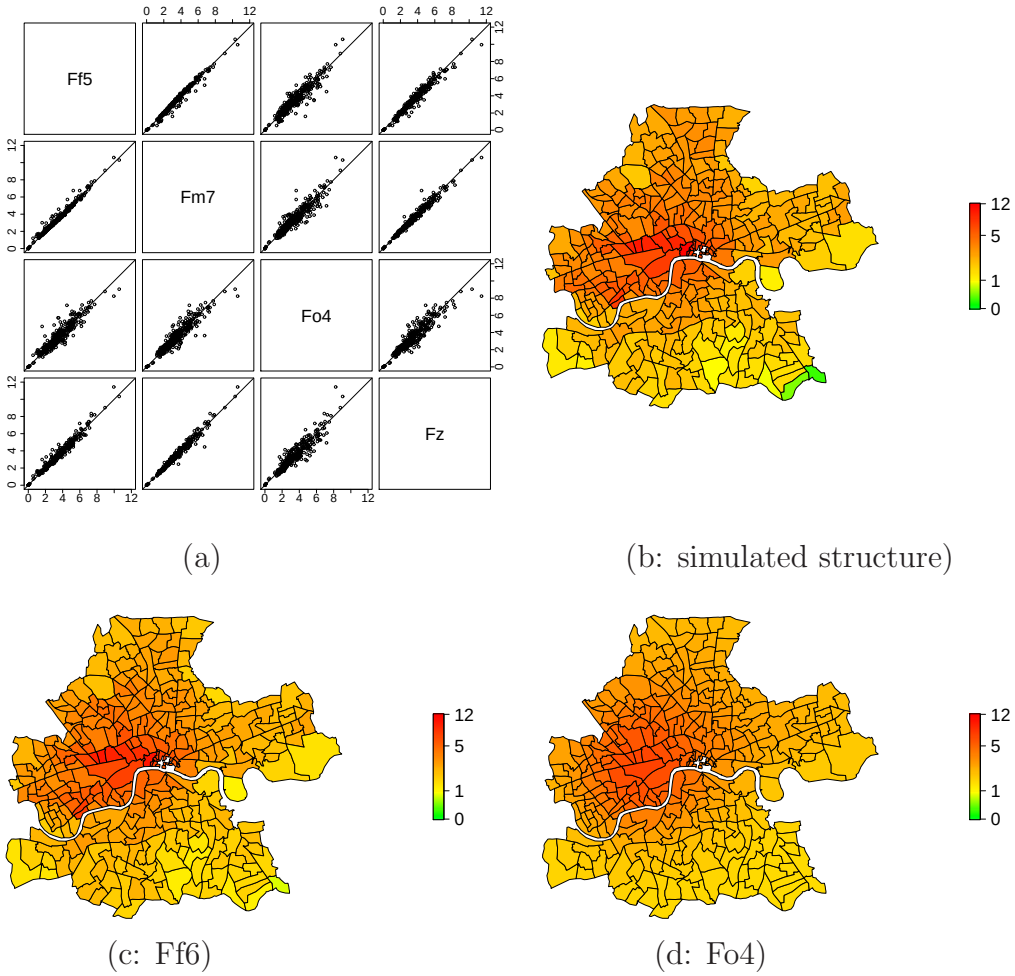

(a) (b: simulated structure)

(c: Ff6) (d: Fo4)

Figure B.4: Results for structure F: Scatterplot matrix of the estimated parameters $\widehat{\Lambda}_i\, E_i$ of Poisson–Gamma models with additive influence of benzene and six latent covariates (Ff6), multiplicative influence of benzene and seven latent covariates (Fm7), four latent covariates without a benzene term (Fo4), and the MRF model assuming wards across the Thames to be neighbours (Fz) (a); simulated structure F (b); estimated rate $\widehat{\Lambda}_i$ of the Poisson–Gamma model with benzene as excess risk factor and six latent risk sources (Ff6) (c) and estimated rate $\widehat{\Lambda}_i$ of the Poisson–Gamma model Fo4 without benzene and four latent risk sources (d).

♯ latent factors	0	1	2	3	4	5	6	7
model f	343.0 (0.388)	336.3 (0.105)	336.3 (0.130)	331.4 (0.086)	330.6 (0.080)	330.4 (0.083)	330.4 (0.085)	330.4 (0.089)
model m	351.9 (0.463)	343.1 (0.174)	339.8 (0.163)	339.5 (0.161)	339.4 (0.166)	338.9 (0.168)	338.6 (0.169)	338.1 (0.171)
model o	—	370.8 (0.654)	360.2 (0.542)	357.4 (0.492)	357.5 (0.477)	—	—	—
model y	360.2(0.627)							
model v	338.6(0.244)							
model z	334.4(0.171)							

Table B.4: DIC (MSE) values for extended models applied to structure F.

	Ff5	Fm7	Fo4	Fz
Ff5	1	0.988	0.936	0.979
Fm7	0.988	1	0.936	0.985
Fo4	0.936	0.936	1	0.934
Fz	0.979	0.985	0.934	1

Table B.5: Structure F: Pearsons correlation coefficient of the estimated parameters $\widehat{\Lambda}_i E_i$ of models plotted in Figure B.4 (a).

When benzene is included multiplicatively DIC is constantly decreasing with an increasing number of Gaussian kernels. As the DIC is decreased by 1.7 points only between model Fm2 (two latent covariates, DIC = 339.8) to model Fm7 (seven latent covariates, DIC = 338.1) we do not proceed further. We compare the estimated values $\widehat{\Lambda}_i E_i$ of model Ff6 with Fm7 in a scatterplot matrix given in Figure B.4 (a). Here, we also include the results of model Fo4, the best model among the Poisson–Gamma models without benzene, and the MRF model Fz assuming wards across Thames to be neighbours.

Model Fz convinces by a DIC of 334.4, which is slightly lower than the one of the MRF model where the Thames disconnects neighbours. BDCD does not convince in modelling data according to structure F. The DIC is 360.2 which is even higher than the one of the class of Poisson–Gamma models not involving benzene.

Figure B.4 (a) reveals a close concordance between the selected models, which is also reflected by the correlation matrix using Pearson's method as presented in Table B.5 which is at least 0.934. The lowest correlation coefficient is calculated for model Fo4 which has also the highest DIC value among the selected models. Additionally, we give plots of the best–fitting model Ff6 and of model Fo4 in Figures B.4 (c) and (d) respectively. We see a high agreement among estimates of both models and the generated values as given in Figure B.4 (b). Both models satisfy by reproducing the characteristics of the generated structure. It follows that models with intermediate correlations lead to good results as well. Even if additive modelling of benzene is to favour, multiplicative modelling as by models Fm7 and Fz is a good substitute.

♯ latent factors	0	1	2	3	4
model f	643.0 (2.462)	393.8 (0.373)	367.3 (0.240)	368.9 (0.224)	370.8 (0.209)
model m	643.2 (2.452)	403.4 (0.400)	379.4 (0.290)	383.9 (0.257)	386.8 (0.249)
model o	—	406.3 (0.461)	391.9 (0.315)	388.1 (0.285)	388.4 (0.365)
model y	345.6 (0.234)				
model v	340.5 (0.126)				
model z	370.6 (0.351)				

Table B.6: DIC (MSE) values for extended models applied to structure G.

B.5 Structure G

Structure G represents data generated assuming an additive influence of benzene and a covariate representing an increased risk in all wards south of the river as plotted in Figure 6.5 (a) on page 68. Both account for about 330 cases. The results of the modelling are reported in Table B.6. Again, we see a huge drop–off in DIC when including one latent covariate in both benzene settings of Poisson–Gamma models. This reflects the presence of risk that is not associated to benzene. Hence, inclusion of latent covariates is required.

Performance of Poisson–Gamma models assuming an additive influence of benzene is better than of those assuming a multiplicative influence which increase the DIC for about 10 points. Best results are achieved when including two latent covariates (Gf2) where the DIC is 367.9 (MSE 0.240). For more than two latent covariates DIC increases while the MSE decreases slightly. As judgement is based on the DIC we choose Gf2 to be the best fitting model. Similar conclusions as drawn for multiplicative models where Gm2 leads to the lowest DIC value of this group which is 379.4.

If benzene is not a covariate in the Poisson–Gamma model we should include three latent covariates leading to a DIC of 388.1 (MSE 0.285). Even if this value is increased by 20 points difference in the MSE is only 0.045 which is

negligible. We compare spatial pattern of models Go3 and Gf2 in Figure B.5.

Risk surface modelled by Gf2 gives the best reproduction when focusing on the appropriate estimation of low-risk regions. The MRF model Gv (Figure B.5 (e)) which treats wards separated by the Thames not to be neighbours performs significantly better (DIC 340.5) than the one where we have neighbours across the Thames (DIC 370.6, Figure B.5 (f)) which agrees with the underlying structure. Model Gv is the one most able to reproduce the abrupt change in Λ_i at the Thames. In comparison to MRF model Gv Poisson–Gamma models have difficulties to estimate the sharp drop-off at the river leading to overestimated risks north of the Thames. For Poisson–Gamma models such as Gf2, a decomposition into latent term and benzene term is possible see Figure B.5 (a) and (b). The latent term reveals the characteristics of Gaussian kernels used for the latent covariates. We observe a small band where the latent influence is decreasing fast. Probably, other kernels like Uniform ones would improve the model fit. Additionally, the estimated pattern is not completely homogeneous in the southern part. Again, the choice of alternative kernels would improve this pattern.

Main differences between model Gf2 and Go3 are observable in the estimation of low risk regions where the Poisson–Gamma model without benzene overestimates. Here a mixture of different kernel functions could lead to improved results.

The BDCD model also convinces by a low DIC value of 345.6. As presented in Figure B.5 (d) the high risk region south of the Thames is identified correctly. Additionally, the model convinces in identifying the sharp border created by the river. For regions north of the river bank the cluster structure estimated by the model does not reflect the generated pattern. Low risk regions in the east are overestimated to a medium risk cluster while the low risk cluster with mean risk of 0.625 detected in the north–west compromises 23 ward whose generated mean risk is 1.258. Using the median instead of the mean Monte Carlo estimates leads to identical conclusions. In such a situation credibility intervals should be considered which are not available by the BDCD software.

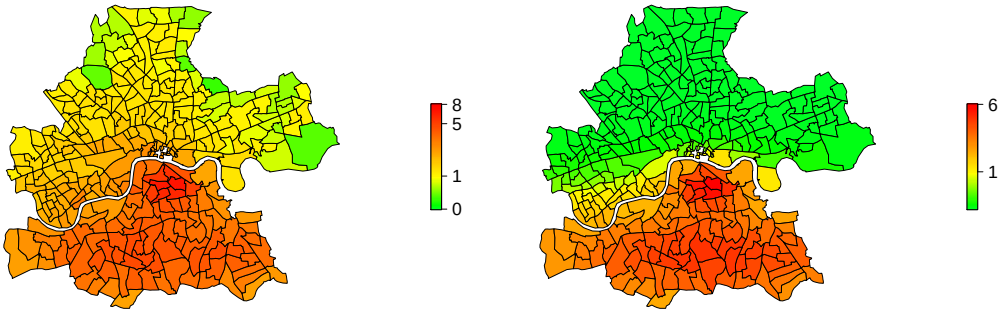

(a: Gf2)

(b: latent term of Gf2)

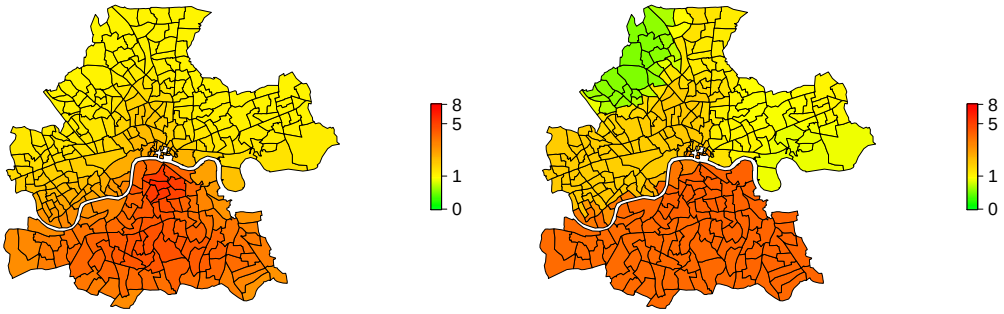

(c: Go3)

(d: Gy)

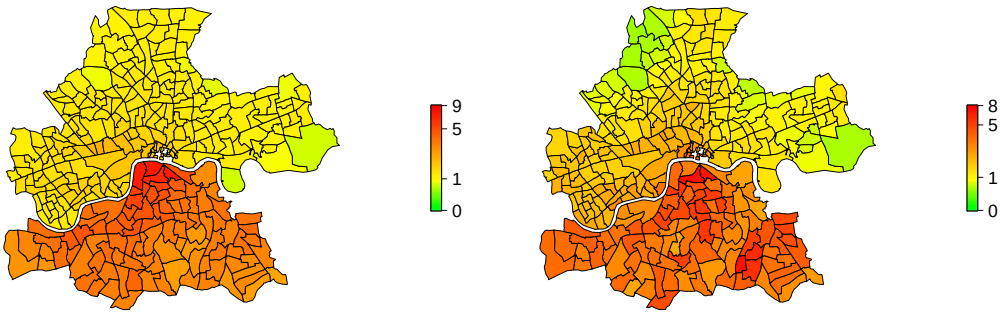

(e: Gv)

(f: Gz)

Figure B.5: Results for structure G: Estimated rate $\widehat{\Lambda}_i$ of the Poisson–Gamma model with additive influence of benzene and two latent risk sources Gf2 (a) and the corresponding latent term (b), without benzene and three latent risk sources Go3 (c), the BDCD model (d), MRF model Gv where the river does part the neighbourhood structure (e) and Gz where the river does not part (f).

♯ latent factors	0	1	2	3	4	5	6
model f	504.4 (2.457)	377.3 (0.436)	373.9 (0.344)	362.0 (0.347)	362.0 (0.331)	362.3 (0.322)	—
model m	508.6 (2.558)	385.9 (0.535)	377.2 (0.497)	370.4 (0.425)	370.2 (0.418)	370.3 (0.416)	—
model o	—	424.9 (0.891)	418.2 (0.704)	416.6 (0.667)	414.3 (0.617)	413.6 (0.604)	412.1 (0.610)
	7	8	9	10	11	12	13
	408.8 (0.610)	407.4 (0.610)	404.9 (0.618)	403.3 (0.620)	402.1 (0.623)	411.7 (0.600)	—
model y	389.9 (0.687)						
model v	353.6 (0.229)						
model z	372.0 (0.430)						

Table B.7: DIC (MSE) values for extended models applied to structure H.

B.6 Structure H

Data generated according to structure H is characterised by an additive influence of benzene accounting for approximately 770 expected cases. For wards situated south of the Thames, an increased risk is generated, accounting for further 330 expected cases.

Poisson–Gamma models including benzene additively such as assumed for data generation lead to satisfying model fits. Best results are achieved when including three to four latent covariates such as in Hf3 and Hf4 leading to a DIC of 362.0. Slightly worse results are estimated by Poisson–Gamma models including benzene multiplicatively. Best fits include four latent sources and have a DIC of 370.2, see Table B.7.

When benzene is not included in the model, the generated data set requires eleven latent covariates to achieve a DIC of 402.1 which corresponds to a MSE of 0.623. Inclusion of more latent covariates decreases the fit. Compared to models including benzene, the DIC is increased by 30 to 40 points indicating an inferior model fit. In real applications this gives an indication that benzene

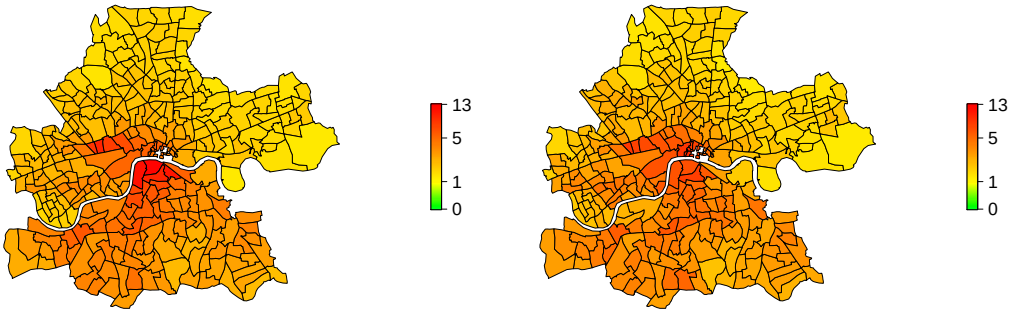

(c: Hv) (d: Hz)

Figure B.6: $\widehat{\Lambda}_i$ by the MRF model Hv where the river does not part the neighbourhood structure (a) and Hz where the river does part (b).

needs to be considered. Additionally, we point out the highly corresponding model fits of the three groups of Poisson–Gamma models, see Figure B.5 (a).

MRF models lead to similar model fits. Here we observe differences due to the chosen neighbourhood structure. The assumption of wards across the Thames to be neighbours leads to a DIC of 372.0 while a separation by the river improves the DIC to 353.6 as this neighbourhood represents the underlying risk pattern much better. Figure B.6 illustrates the differences in $\widehat{\Lambda}_i$ between both models.

Figure B.5 (a) presents a scatterplot matrix of estimated $\widehat{\Lambda}_i E_i$ of Poisson–Gamma models representing the best model fits in each group. Additionally, we include estimated values of the more appropriate MRF model Hv which assumes the Thames to separate neighbours as well as those by the BDCD model. Even if we expected the BDCD model to be a good performing one due to the clustered pattern, the model achives a higher DIC value than Poisson–Gamma models including benzene as well as MRF models.

We give the underlying pattern of this structure in Figure B.5 (b) as well as the estimated spatial pattern $\widehat{\Lambda}_i$ for model Ho11 (Figure B.5 (c)) and for model Hf4 (Figure B.5 (d)). All models lead to comparable results. Most deviations are observable for higher values estimated by model Hv. Never-

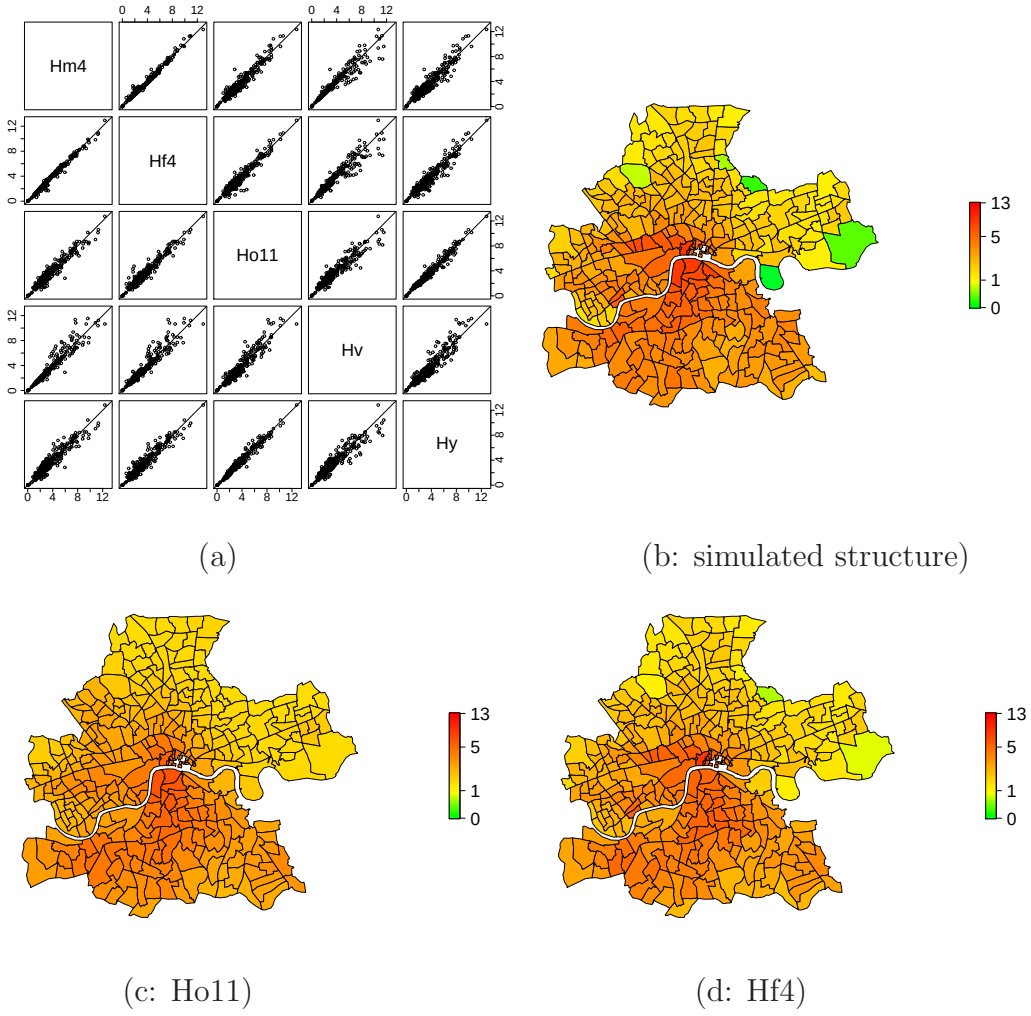

(a) (b: simulated structure)

(c: Ho11) (d: Hf4)

Figure B.7: Results for structure H: Scatterplot matrix of the simulated and estimated parameters $\widehat{\Lambda}_i\, E_i$ of the Poisson–Gamma model with multiplicative influence of benzene and four latent covariates (Hm4), additive influence of benzene and four latent covariates (Hf4) and eleven latent covariates without a benzene term (Ho11) (a); simulated structure H (b); pattern of $\widehat{\Lambda}_i$ for model Ho11 (c) and $\widehat{\Lambda}_i$ for model Hf4 (d).

♯ latent factors	0	1	2	3	4	5
model f	1001.9 (11.405)	670.3 (7.827)	486.3 (3.811)	454.1 (2.811)	410.7 (2.017)	411.4 (2.017)
model m	982.6 (11.286)	700.4 (8.256)	502.1 (4.236)	465.7 (3.063)	478.3 (2.581)	—
model o	—	1006.2 (11.389)	805.2 (6.543)	521.2 (3.703)	487.4 (3.397)	466.6 (2.804)
	6	7	8	9	10	11
	461.3 (2.665)	459.0 (2.583)	457.2 (2.536)	456.0 (2.505)	476.1 (2.495)	—
model y	397.6 (1.820)					
model v	431.7 (2.568)					
model z	431.9 (2.515)					

Table B.8: DIC (MSE) values for extended models applied to structure J.

theless, spatial pattern are close to the generated one. All models lack of the ability to estimate the three generated low risk regions correctly, although model fit is good in general. This is also reflected by the estimated MSEs, see Table B.7. There is no difference between the additive and the multiplicative setting of benzene, even if treating benzene as excess risk factor leads to a better fit with less covariates.

B.7 Structure J

Data generated by the assumption of an additive influence of benzene accounting for about 770 observations in combination with 330 cases in three cluster regions corresponds to structure J. Resulting MSEs and DICs are summarised in Table B.8. We also calculate Pearsons correlation coefficient between the best models in each group, see Table B.9.

Poisson–Gamma models assuming benzene to be an excess risk factor correspond to the generating structure. This model also leads to the lowest DIC which is 410.7 when including four latent covariates. Inclusion of further

	J	Jf4	Jm3	Jo9	Jz	Jy
J	1.000	0.919	0.882	0.905	0.904	0.934
Jf4	0.919	1.000	0.958	0.978	0.935	0.914
Jm3	0.882	0.958	1.000	0.944	0.900	0.857
Jo9	0.905	0.978	0.944	1.000	0.943	0.932
Jz	0.904	0.935	0.900	0.943	1.000	0.953
Jy	0.934	0.914	0.857	0.932	0.953	1.000

Table B.9: Structure J: Pearsons correlation coefficient of the generated (J) and estimated values $\widehat{\Lambda}_i E_i$ of the Poisson–Gamma model with additive influence of benzene and four latent covariates (Jf4), with multiplicative influence and three latent covariates (Jm3), Poisson–Gamma model with nine latent covariates (Jo9), MRF model Jz, and BDCD algorithm Jy.

covariates increases the DIC.

The assumption of benzene to be a relative risk factor does not lead to satisfying results. The correlation coefficient between model Jm3 including three latent covariates and the generated values is 0.882, compare Table B.9. We calculate the DIC to be 465.7 for this model. For higher number of latent sources the DIC increases.

This result can be improved when not including benzene but only latent covariates. Although the number of those is higher now, model Jo9 satisfies by a DIC of 456.0. The correlation between model Jf4 and Jo9 is 0.978, see Table B.9. Both show a high correlation to the generated values of 0.919 and 0.905 respectively. Corresponding spatial plots are given in Figure B.8. By visual inspection, both models are able to identify the high risk regions correctly. As model Jf4 uses only four Gaussian kernels to model the risk in the cluster regions, the ability to identify the plateau-like character is limited. Therefore, the risk is overestimated in some of those regions. Model Jo9 has less problems here as nine kernels provide a more flexible way of estimation. On the other hand, low risk regions present especially in border regions of the generated structure cannot be reproduced by model Jo9. Here model Jf4 has its amenities.

With a value of 397.6, the DIC estimated by the BDCD model Jy is even

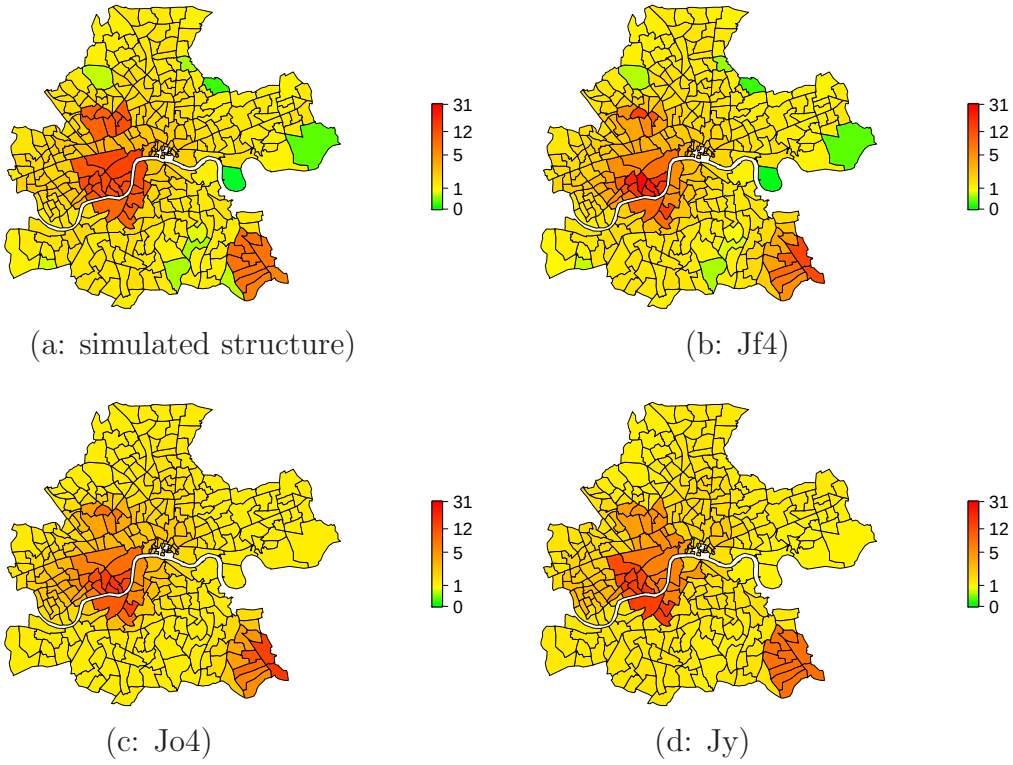

Figure B.8: Simulated Λ_i for structure J (a) and results: Estimated spatial pattern of $\widehat{\Lambda}_i$ of Poisson–Gamma models Jf4 (b) and Jo9 (c) and of the BDCD model Jy (d) .

lower compared to Poisson–Gamma models. The corresponding spatial plot is presented in Figure B.8 (d). Here we obtain similar results as for model Jo9 concerning low risk regions: those are not reproduced by the model. High risk regions are identified correctly, although we see different risk levels than generated in the central cluster.

We also employ MRF models on this structure. Both involved neighbourhood structures lead to similar results. The DIC is about 431 which is slightly worse than Poisson-Gamma models involving benzene additively but leads to better results than other settings for Poisson–Gamma models. They provide an alternative modelling tool in this situation.

♯ latent factors	0	1	2	3	4	5	6
model f	321.9 *2.1* (0.006)	323.6 *4.6* (0.007)	323.3 *4.6* (0.007)	392.3 *20.8* (0.145)	323.6 *6.1* (0.009)	—	—
model m	322.8 *1.9* (0.008)	324.3 *4.0* (0.010)	324.8 *5.5* (0.010)	323.8 *5.9* (0.009)	323.5 *3.0* (0.009)	—	—
model o	—	329.7 *5.9* (0.041)	327.8 *8.0* (0.034)	327.2 *8.0* (0.034)	326.0 *8.9* (0.033)	325.9 *9.1* (0.033)	386.4 *22.0* (0.180)
model y	328.0, *7.5* (0.043)						
model v	321.1, *5.5* (0.021)						
model z	322.4, *5.0* (0.005)						

Table B.10: DIC, p_D and (MSE) values for extended models applied to structure M.

B.8 Structure M

For structure M, we generate 330 observations determined by the amount of benzene only. Contrary to structure A benzene is assumed to be a relative risk factor. The spatial pattern of Λ_i is given in Figure B.9 (b). These are modelled by the class of Poisson–Gamma models with a number of latent risk sources as well as MRF models and the BDCD algorithm. Resulting DICs and MSEs are reported in Table B.10.

In the class of Poisson–Gamma models we estimate similar risk assuming either additive or multiplicative influence of benzene. The introduction of any latent covariates does not improve the model fit. When benzene is not included, the fit is degraded. Five latent risk factors are necessary to replace the benzene covariate and to gain a fit of a similar quantity. This leads to a DIC value of 325.9 which is an increase of 4 points only in comparison to Mf0, compare Table B.10. MSEs are very low for all models, e.g., 0.033 for Mo5 and 0.008 for Mm1.

The lowest DIC values in the class of Poisson–Gamma models are achieved by the

- Poisson–Gamma model assuming an additive influence of benzene and no latent factors (Mf0);
- Poisson–Gamma model assuming a multiplicative influence of benzene and no latent factors (Mm0);
- Poisson–Gamma model with five latent covariates (Mo5).

We compare the estimated parameters of the Poisson distribution by a scatterplot matrix, see Figure B.9 (a). We see almost identical estimates $\widehat{\Lambda}_i E_i$ for model Mm0 and model Mf0. Estimates of model Mo5 are close, although we observe small deviations across the whole range of values, in particular for higher values. In addition to the scatterplot matrix, we calculate the correlation coefficient by Pearson. The matrix for models Mf0, Mm0 and Mo5 is

$$\begin{pmatrix} 1 & 0.998 & 0.964 \\ 0.998 & 1 & 0.958 \\ 0.964 & 0.957 & 1 \end{pmatrix}.$$

The spatial risk surface of model Mm0 is given in Figure B.9 (c), estimated $\widehat{\Lambda}_i$s for model Mo5 are found in Figure B.9 (d).

MRF models lead to similar risk surfaces for both neighbourhood structures, the BDCD algorithm produces slightly higher DIC values.

Altogether, the best model for this structure equals the underlying structure. Furthermore, the assumption of additive influence of benzene or the usage of latent covariates instead of the benzene term reproduces the generated structure well.

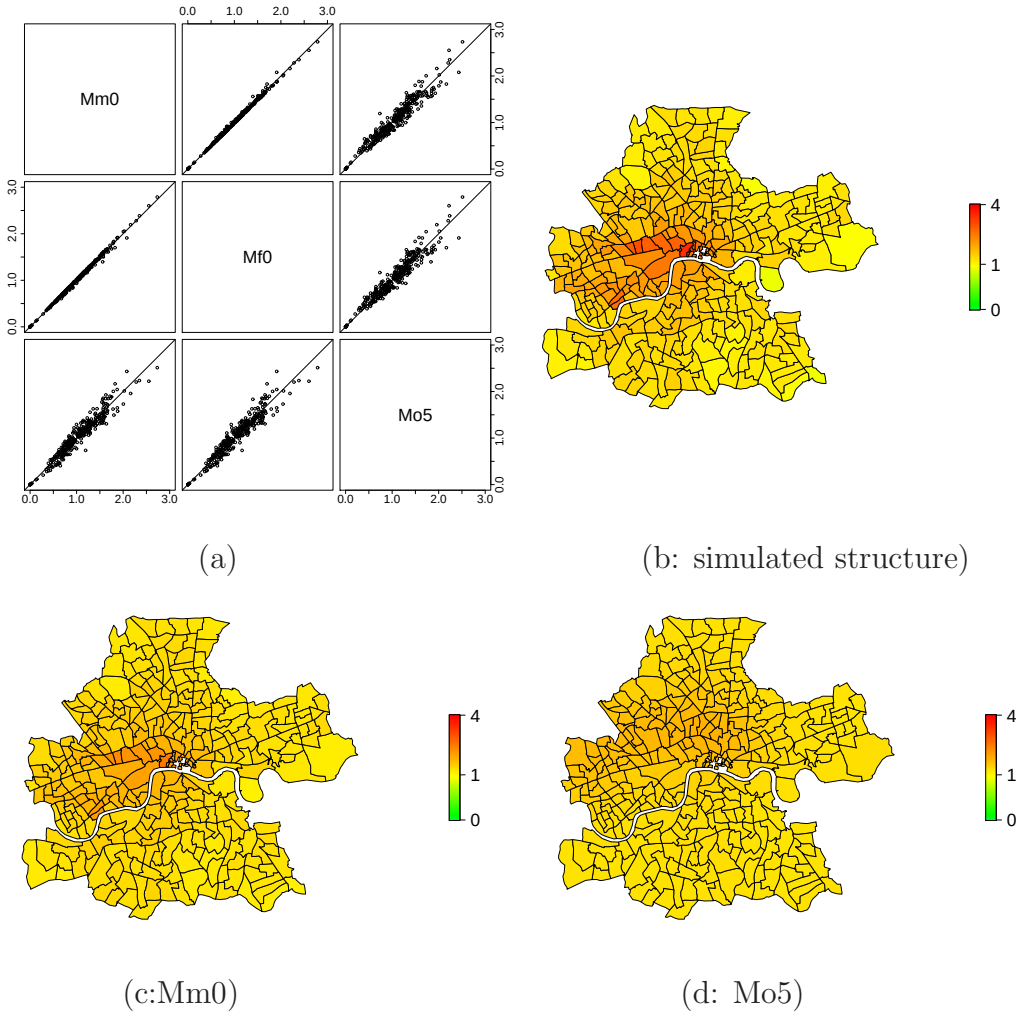

(a) (b: simulated structure)

(c:Mm0) (d: Mo5)

Figure B.9: Results for structure M: Scatterplot matrix of the estimated parameters $\widehat{\Lambda}_i\, E_i$ of the model with multiplicative influence of benzene (Mm0), additive influence of benzene (Mf0) and five latent covariates without a benzene term (Mo5) (a); simulated structure M (b); pattern of $\widehat{\Lambda}_i$ for model Mm0 (c) and $\widehat{\Lambda}_i$ for model Mo5 (d).

♯ latent factors	0	1	2	3	4	5
model f	442.4 (0.893)	387.6 (0.575)	388.8 (0.523)	—	—	—
model m	340.5 (0.055)	341.5 (0.053)	340.7 (0.053)	—	—	—
model o	—	466.0 (1.141)	450.5 (0.870)	433.7 (0.853)	427.6 (0.798)	433.3 (0.750)
model y	420.6 (0.972)					
model v	366.5 (0.277)					
model z	364.6 (0.282)					

Table B.11: DIC (MSE) values for extended models applied to structure N.

B.9 Structure N

We generate observations depending multiplicatively on benzene only. In structure N, this leads to about 770 cases. Generated rates Λ_i are plotted in Figure B.10 (a).

When including benzene into Poisson–Gamma models, we need only a small number of latent risk sources in the model to get the best fit. For the favourable model among the additive ones Nf1 we achieve a DIC of 387.6. In contrast, if we set the influence of benzene to be multiplicatively as assumed in data generation, the model fit is improved reflected by a DIC of 341.5 for Nm1. The DIC of model Nm0 is even lower, namely 340.5, see Table B.11. Therefore, the model corresponding to the underlying structure is to favour.

When not including benzene into the Poisson–Gamma model, we achieve higher values for the DIC, which are at least 427.6. Model No4 that leads to best results among this group includes four latent Gaussian kernels. The corresponding MSE of 0.798 is highly satisfying.

We plot rates $\widehat{\Lambda}_i$ of the multiplicative model not including any further covariates (Nm0) as well as the Poisson–Gamma model consisting of four latent covariates only (No4) in Figure B.10. The resulting pattern are similar, especially in identification of high risk regions. Low risk regions are somewhat

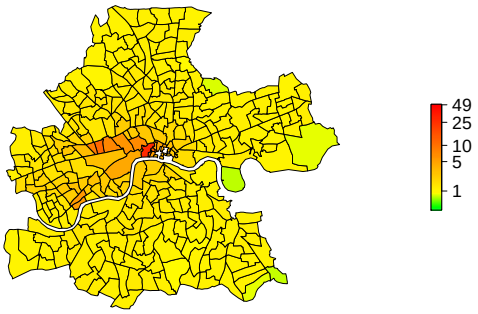

(a: simulated structure)

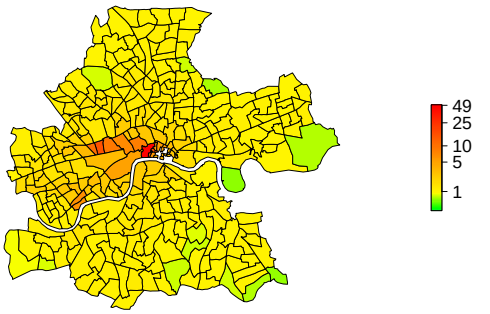

(b: Nm0)

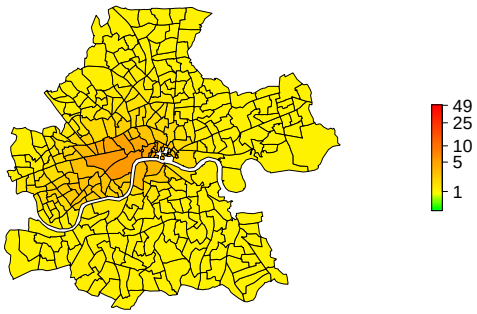

(c: No4)

Figure B.10: Simulated Λ_i for structure N (a) and results: Estimated spatial pattern of $\widehat{\Lambda}_i$ of the best fitting model Nm0 (b) and of the best one without benzene No4 (c).

overestimated by model No4.

In terms of alternative models, we observe a similar performance of the BDCD model as for the Poisson–Gamma models not including benzene. Here we estimate a DIC of 420.6. MRF models including the CAR structure as well as benzene on a multiplicative level perform slightly better. Here we achieve a DIC of about 365.0 depending on the choice of neighbourhood, although its influence is low for this structure. Nevertheless, Poisson–Gamma models convince by even lower DICs as well MSEs.

B.10 Structure P

For structure P, we assume a multiplicative influence of benzene together with the influence of a latent covariate represented by a Gaussian kernel. In this structure, we assume a 1:2 matching for the two covariates in terms of the numbers of cases caused by either. The generated Λ_i are displayed in Figure B.11 (a).

For Poisson–Gamma models assuming benzene to be a relative risk factor we find model Pm1 to be the most suitable, see Table B.12. This model includes one Gaussian kernel and is therefore identical with the data generating structure. The estimated $\widehat{\Lambda}_i$ are presented in Figure B.11 (c). We observe a high agreement between the generated and the estimated risk which is also reflected by a low MSE of 0.381.

Other models have difficulties to model the generated high risk represented by the Gaussian kernel. The maximum of the generated Λ_i is 160.5, for the optimal model Pm1 we estimate a maximum of 166.5. In contrast, the maxima for other applied models are 123.7 (model Pf7), 118.9 (model Po6), 142.8 (model Pz), and 154.4 (model Py). This also results in differences between estimated and generated values in the surrounding regions and therefore increased MSEs and worsened model fit.

For example, a Poisson–Gamma model that assumes benzene to be an excess risk factor such as model Pf7 does not lead to satisfying results. On one hand, the DIC does not improve remarkably after including five to seven latent covariates. Therefore, we do not add further covariates. On the other hand, model fit is worse compared to multiplicative modelling of benzene, compare the DICs of model Pm1 and Pf7 in Table B.12.

We find similar conclusions for Poisson–Gamma models that do not include benzene such as model Po6. Here we estimate a MSE of 6.267. The corresponding DIC of 478.4 is even worse than the one for model Pf7. Beside the problems of estimating high risk appropriately, the model has also difficulties in identifying low risk regions. Here, the minimum of $\widehat{\Lambda}_i$ is 2.0. In contrast, $\min(\Lambda_i) = 0.8$. It is possible that a larger number of latent covariates im-

♯ latent factors	0	1	2	3	4	5	6	7
model f	2055.6 (70.386)	664.8 (23.530)	506.3 (4.761)	494.6 (9.254)	446.4 (8.297)	425.5 (5.339)	425.3 (5.077)	425.0 (4.901)
model m	1785.8 (64.035)	341.1 (0.381)	342.6 (0.394)	344.1 (0.435)	—	—	—	—
model o	—	768.9 (26.783)	627.7 (13.212)	524.5 (5.027)	551.0 (10.212)	479.0 (6.733)	478.4 (6.267)	478.8 (6.058)
model y	482.2 (3.929)							
model v	401.6 (2.481)							
model z	400.9 (2.711)							

Table B.12: DIC (MSE) values for extended models applied to structure P.

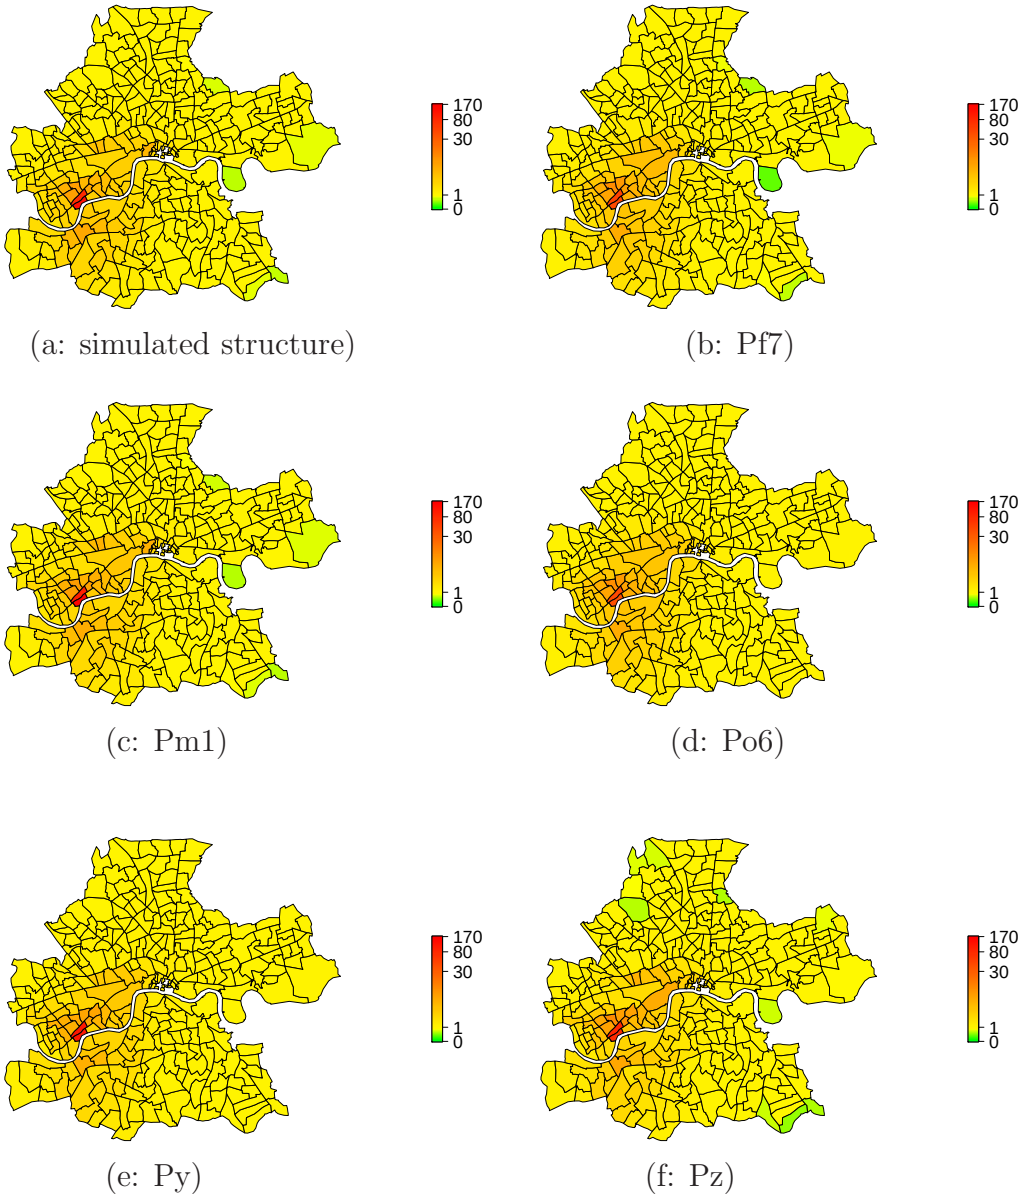

(a: simulated structure) (b: Pf7)

(c: Pm1) (d: Po6)

(e: Py) (f: Pz)

Figure B.11: Simulated pattern Λ_i for structure P (a) and results: Estimated rate $\widehat{\Lambda}_i$ of the Poisson–Gamma model with additive influence of benzene and seven latent risk sources Pf7 (b), with multiplicative influence of benzene and one latent risk sources Pm1 (c) and without benzene and six latent risk sources Po6 (d), the BDCD model Py (e) and the MRF field model where wards across Thames are assumed to be neighbors (f).

proves the model fit, but as DIC increases from model Po6 to Po7, we do not proceed further.

The BDCD model shows similar disadvantages. The model fails to identify low risk regions correctly, though the minimum of 1.2 is lower than the one of model Po6. Given the DIC, the BCDC model is least appropriate for this structure. The risk as plotted in Figure B.11 (e) confirms the discussed findings.

The MRF model estimates DIC values of about 400, see Table B.12. Both neighbourhood structures produce comparable results. This is not surprising as there is no obvious reason why differences in the neighbourhood close to the river should lead to better or worsened fit. The estimated spatial pattern is close to the generated structure which is also reflected by a MSE of 2.711 for model Pz and 2.481 for model Pv, see Table B.12 and the plot of $\widehat{\Lambda}_i$ for model Pz in Figure B.11 (f).

Altogether, we are able to identify the underlying structure by our models correctly. In contrast to other structures we observe the necessity to use benzene as a relative risk factor. Using other model classes, the models have difficulties to estimate the underlying structure in a comparable quality. Nevertheless, if we compare the spatial structures as given in Figure B.11 we recognise the ability of all models to identify the main characteristics of the generated structure.

B.11 Structure Q

Structure Q assumes a multiplicative influence of benzene at a low level that is combined with a covariate representing a linear spatial trend that decreases from north to south. A plot of the spatial structure of generated Λ_i can be found in Figure 6.4.

Again, all models are applied to a data set generated corresponding to this structure. For achieved DICs and MSEs see Table B.13.

As the huge drop in DIC indicates we require to include a minimum of one

♯ latent factors	0	1	2	3	4
model f	517.9 *2.0* (1.253)	339.7 *5.6* (0.100)	340.0 *6.5* (0.096)	343.2 *15.3* (0.061)	—
model m	517.9 *2.0* (1.253)	338.3 *5.5* (0.076)	340.2 *9.8* (0.062)	341.1 *11.9* (0.057)	—
model o	—	343.5 *4.8* (0.158)	346.0 *13.0* (0.101)	342.7 *12.8* (0.097)	343.4 *17.4* (0.079)
model y	345.5, *20.2* (0.163)				
model v	356.7, *37.1* (0.134)				
model z	353.0, *52.0* (0.139)				

Table B.13: DIC, *pD* and (MSE) values for extended models applied to structure Q.

latent covariate into the model. Inclusion of further Gaussian kernels slowly increases the DIC. Hence our favoured models in this group are:

- Poisson–Gamma model with benzene as excess risk factor and one latent covariate (Qf1), DIC = 339.7;
- Poisson–Gamma model with benzene as relative risk factor and one latent covariate (Qm1), DIC = 338.3;

Ignorance of benzenes influence in modelling such as for models Qo3 (Poisson-Gamma model with three latent covariates only) and Qy (BDCD algorithm) increases the DIC by about 10 points for the generated data set. Both models produce similar qualities of fit as confirmed by their DIC values and the scatterplot in Figure B.12 (a).

Fit by MRF models increases DIC by about 15 points compared to Qf2 and Qm2. Different neighbourhood structures do not lead to any remarkable discrepancy.

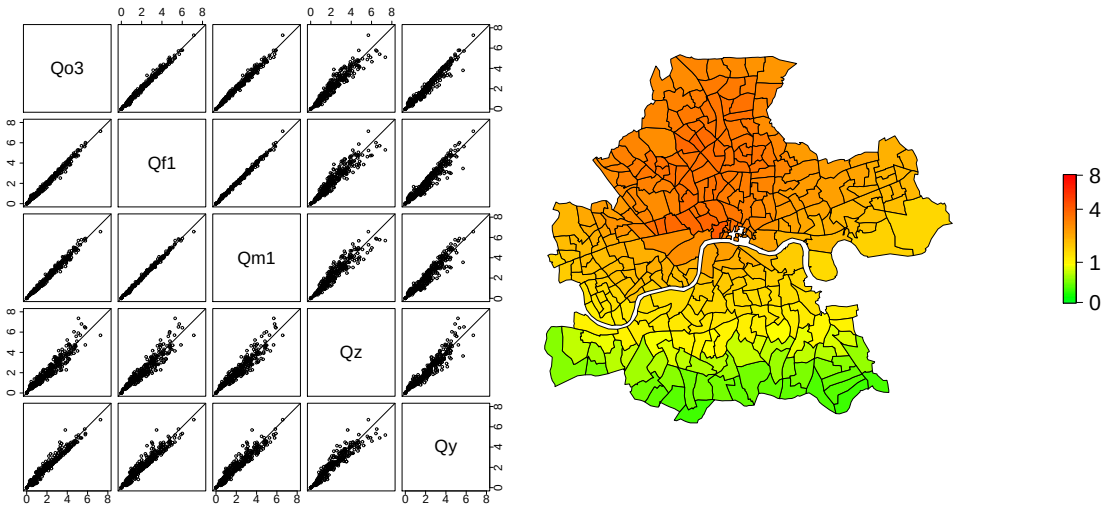

(a) (b)

Figure B.12: Results for structure Q: Scatterplot matrix of the estimated parameters $\widehat{\Lambda}_i\, E_i$ of the Poisson–Gamma model with multiplicative influence of benzene (Qm1), additive influence of benzene (Qf1) and no benzene influence (Qo3) (all models include one latent risk source), and MRF model Qz and BDCD model Qy (a); spatial pattern of $\widehat{\Lambda}_i$ estimated by model Qm1 (b).

We compare the estimated parameters of the Poisson distribution $\widehat{\Lambda}_i\, E_i$ by a scatterplot matrix, see Figure B.12 (a). Best agreement in the chosen models is between Qf1 and Qm1, model Qo3 has some minor deviations, for Qz and Qo deviations are higher. For all models those deviations occur on the whole interval. Both, the MSE and DIC are very similar for all models, we achieve similar model fits for all applied models.

We present the risk surface $\widehat{\Lambda}_i$ for model Qm1 which has the lowest DIC of all models, that is 338.3. There are some small deviations between generated and modelled values, but the risk surface is reproduced reasonably well. This is reflected by the small MSE of 0.076.

For this structure, our favoured model includes a benzene term and incorporates one latent Gaussian kernel. Differences due to the type of influence is negligible.

♯ latent factors	0	1	2	3	4
model f	619.7 (3.910)	441.6 (2.523)	387.7 (1.213)	395.0 (1.483)	—
model m	555.8 (2.261)	327.7 (0.200)	329.9 (0.191)	331.1 (0.188)	—
model o	—	518.4 (3.804)	481.1 (2.517)	419.7 (1.602)	424.7 (1.766)
model y	443.4 (1.537)				
model v	360.8 (0.454)				
model z	363.0 (0.537)				

Table B.14: DIC (MSE) values for extended models applied to structure R.

B.12 Structure R

Data generated according to structure R are characterised by a multiplicative influence of benzene combined with a covariate that has a linear spatial trend. In contrast to structure Q, benzene accounts for about 770 cases. The trend covariate accounts again for 330 cases. The corresponding risk surface of Λ_i is given in Figure B.13 (a).

We employ Poisson–Gamma models with various settings on the generated data set, resulting DICs are given in Table B.14. Best results are achieved when benzene is included multiplicatively in a Poisson–Gamma model in combination with one latent risk factor. The calculated DIC for model Rm1 is 327.7, the corresponding risk surface is plotted in Figure B.13 (b). Similarity between Rm1 and the generated structure is high, which is reflected by an MSE of 0.200 as well. The main difference between generated and estimated values is found in the high risk regions in the center, where estimations do not reproduce the very high rates. The assumption of an additive influence of benzene leads to an inferior model fit. The best fit among model class Rf is given by model Rf2 with a DIC of 387.7 which is an increase of 60 points.

If we do not include benzene in our model, the best surface corresponds to model Ro3 plotted in Figure B.13 (c). The DIC is 419.7 and the corre-

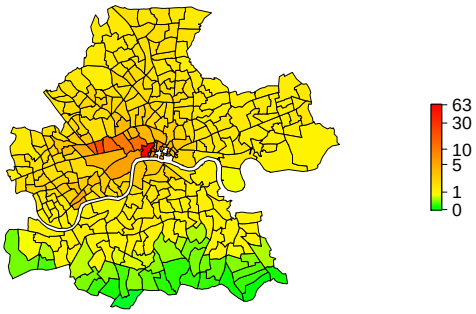

(a: simulated structure)

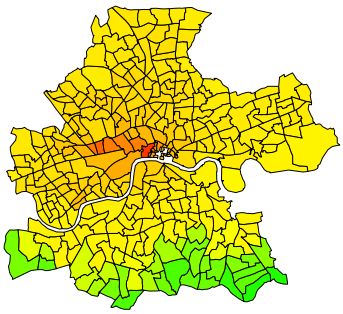

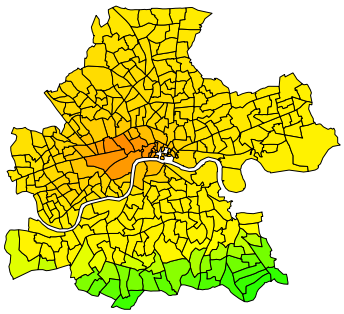

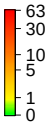

(b: Rm1) (c: Ro3)

Figure B.13: Simulated pattern Λ_i for structure R (a) and results $\widehat{\Lambda}_i$: Estimated spatial pattern of the Poisson–Gamma model Rm1 with multiplicative influence of benzene and one latent risk source (b), and model Ro3 without benzene and three latent risk sources (c).

sponding MSE 1.602. When comparing the estimated and generated spatial surfaces a high agreement is observable. There are deviations in the southern low-risk regions as well as in the north-western part of Inner London. Additionally, the range of $\hat{\Lambda}_i$ is in the interval of $[0.267, 27.612]$ while Λ_i has values in $[0, 62.839]$. The difference of the upper bound values is remarkable. Nevertheless, the fit of model Ro3 reproduces the main characteristics of the generated structure and presents a satisfying estimation. Although DIC of models Rm1 and Ro3 differ in 90 points, main characteristics of the risk surface are similar. We point out that Ro3 especially underestimates high risk as it tends to oversmooth.

For the MRF model the fit is inferior to that of model Rm1. Nevertheless, results are more satisfying than those of Poisson–Gamma models which do not include benzene multiplicatively, see Table B.14. We see no differences due to the different neighbourhood structures applied. The clustering approach of Knorr-Held and Raßer (2000) is not able to model this structure appropriately as the clustering assumption is too restrictive to model a smoothly decreasing risk.

For this structure we get a clear indication that inclusion of benzene is required. All models including a benzene term lead to a lower DIC. Multiplicative modelling as for Rm1 and the MRF models is to favour, best results are achieved for models corresponding to the underlying pattern.

B.13 Structure T

For this structure T, we generate 770 observations due to benzene that is assumed to be a relative risk factor, in combination to 330 further cases that built a plateau of increased risk on the southern river bank of the Thames. Figure B.14 (b) gives an impression of the generated risk surface.

As for other structures we apply Poisson–Gamma models with different settings as well as MRF model and the BDCD algorithm on our generated data. Achieved MSEs and DICs are given in Table B.15. The lowest DIC of 386.3 is achieved for the BDCD algorithm. Obviously, the clustering algorithm is the most appropriate model to identify a sharp decrease in the risk surface as present in our data. This structure also corresponds to the neighbourhood applied for model v. The decrease in risk is supported by the model, leading to a DIC of 392.8. Although benzene is considered in both data generation and applied model, DIC is increased by 4.5 points compared to model y not including benzene. We conclude that although the MRF model leads to an acceptable model fit after identifying an appropriate neighbourhood the BDCD model is more appropriate when clustering is present.

Model z does not include such a corresponding neighbourhood. This de-

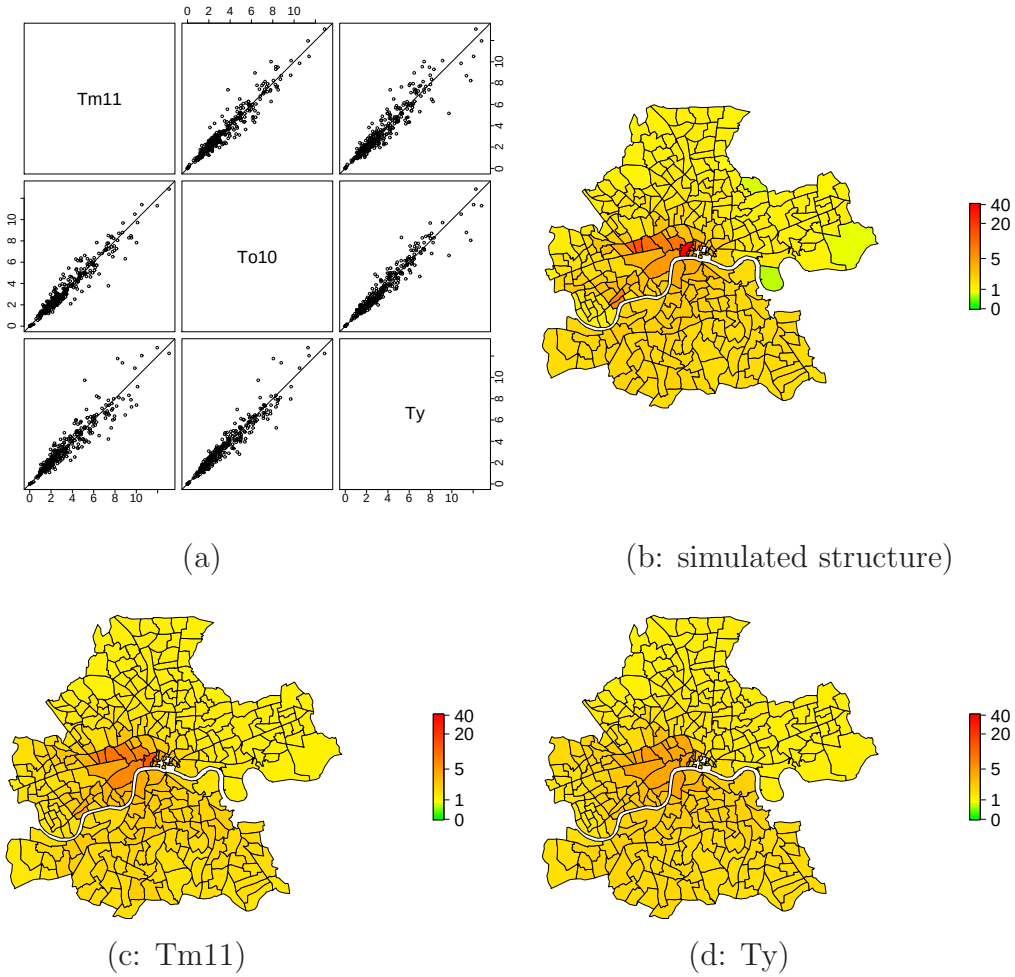

(a) (b: simulated structure)

(c: Tm11) (d: Ty)

Figure B.14: Results for structure T: Scatterplot matrix of the estimated parameters $\widehat{\Lambda}_i\, E_i$ of Poisson–Gamma models Tm11, To10 and Ty; simulated Λ_i of structure T (b); spatial pattern of $\widehat{\Lambda}_i$ estimated by model Tm11 (c) and by model Ty (d).

♯ latent factors	0	1	2	3	4	5
model f	576.3 (2.777)	456.5 (0.991)	437.1 (0.804)	425.5 (0.659)	430.1 (0.687)	—
model m	554.1 (2.425)	425.0 (0.498)	422.0 (0.371)	419.3 (0.350)	412.9 (0.364)	410.9 (0.382)
	6	7	8	9	10	11
	409.5 (0.399)	408.2 (0.415)	406.8 (0.431)	405.5 (0.444)	404.7 (0.454)	404.1 (0.461)
model o	—	512.9 (1.854)	465.6 (1.314)	448.7 (1.064)	448.1 (1.040)	447.2 (1.029)
	6	7	8	9	10	11
	446.4 (1.023)	445.8 (1.018)	439.9 (1.110)	438.9 (1.109)	438.4 (1.108)	444.6 (1.016)
model y	386.3 (0.487)					
model v	392.8 (0.631)					
model z	440.9 (1.078)					

Table B.15: DIC (MSE) values for extended models applied to structure T.

creases the DIC by 54.6 points compared to model Ty. We use this model to compare with the results of Poisson–Gamma models that — in our case — use Gaussian kernels to model latent risk. An overall neighbourhood is therefore more appropriate for a comparison.

While Poisson–Gamma models not accounting for the covariate lead to a similar model fit represented by a DIC of 438.4, those Poisson–Gamma models including benzene satisfy by lower values. If benzene is considered multiplicatively as in data generation, DIC drops to a value of 404.1. Even if this corresponds to an increase of 18 points compared to BDCD, the result is satisfying. Inclusion of benzene as an excess risk factor leads to higher DIC values, although better results than model To10.

In Figure B.14 (a) we give a scatterplot matrix of $\widehat{\Lambda}_i\, E_i$ as estimated by model Ty, To10 and Tm11. Although deviations occur, we rather observe an overall agreement of all models, instead of systematic differences. We therefore take a closer look on risk surfaces $\widehat{\Lambda}_i$ of models Tm11 and Ty. Both show very similar pattern and represent an approximation to the generated structure. Nevertheless, both models oversmooth the high risk in central wards. While we generate values in the interval $[0.750, 40.840]$, estimates of model Ty are in a subset of this interval only, namely in the interval $[1.536, 15.498]$. For model Tm11, we get $[1.349, 26.517]$. This results not only in the disability to estimate high risk regions correctly, but also to overestimate low risk regions.

As DIC of the clustering model as well as of model Tv are lower than those compared to Poisson–Gamma models with Gaussian latent risk structures, we expect our results to be highly improvable by alternative latent risk modelling, for example via Uniform kernels or half–Gaussian ones.

B.14 Structure U

Data containing 330 cases due to benzene and 330 cases in three clusters located all over the are of Inner London is denoted as structure U. The corresponding risk surface is given in Figure 6.6 on page 70.

We apply all selected models on the generated sat set leading to the values presented in Table B.16. For Poisson–Gamma models applied on this structure, a minimum DIC is not reached after inclusion of 15 Gaussian kernels. Nevertheless we stop our procedure here as model selection gets too time-consuming if further kernels should be considered.

Poisson–Gamma models profit from the inclusion of benzene as comparison of the calculated DIC values show. Model Uo7 reaches the lowest DIC of 417.6 which is about 24 points higher than the DICs of model Um15 (DIC = 393.1) and model Uf15 (DIC=393.1). We do not notice any difference in those values concerning DIC or MSE.

Lower DIC values are estimated for the clustering approach. Here, we get a value of 358.9. For MRF models we get higher values again. Although considering benzene, they are at a similar level as those for Poisson–Gamma models without benzene. As a comparison of model Uv and Uz reveals, a neighbourhood structure parted at the Thames is not optimal. Other settings are possible, but not practical as we have to many potential options.

For Poisson–Gamma models a special choice of a neighbourhood is not required. From our results we conclude that Gaussian kernels are not optimal for this structure. Better results for BDCD model Uy confirm this idea. As modelling of sharp risk is required we recommend the usage of according kernels, for example Uniform ones.

To get an impression on the actual estimated risk surface and the differences in the modelling approaches, we give a scatterplot matrix of $\widehat{\Lambda}_i\, E_i$ of the two Poisson–Gamma models including benzene as well as the BDCD model in Figure B.15. Additionally, we present the corresponding risk surfaces $\widehat{\Lambda}_i$.

As we already conclude from the scatterplot of model f15 against m15 both

♯ latent factors	0	1	2	3	4	5	6	7
model f	1194.5	913.4	527.3	446.0	435.0	434.3	426.1	418.6
	2.0	*6.2*	*10.6*	*15.9*	*20.2*	*33.3*	*37.0*	*40.4*
	(11.278)	(7.462)	(3.136)	(2.360)	(2.061)	(1.611)	(1.390)	(1.215)
	8	9	10	11	12	13	14	15
	413.2	404.9	403.8	399.3	397.1	396.4	394.7	393.1
	42.0	*41.4*	*41.9*	*41.5*	*41.5*	*41.8*	*41.7*	*41.7*
	(1.164)	(1.116)	(1.103)	(1.087)	(1.078)	(1.079)	(1.072)	(1.067)
model m	1186.7	1071.5	529.4	445.1	444.5	436.6	430.5	421.3
	2.0	*212.8*	*9.8*	*14.8*	*28.4*	*32.6*	*41.6*	*42.6*
	(11.233)	(5.334)	(3.237)	(2.457)	(1.947)	(1.737)	(1.372)	(1.274)
	8	9	10	11	12	13	14	15
	416.2	408.9	403.2	401.2	397.3	395.9	393.3	393.1
	43.4	*42.8*	*42.2*	*42.2*	*41.6*	*41.6*	*41.2*	*41.4*
	(1.225)	(1.168)	(1.140)	(42.184)	(1.111)	(1.109)	(1.904)	(1.089)
model o	—	1210.0	528.5	450.0	438.0	433.8	426.3	417.6
		296.2	*9.7*	*13.9*	*22.2*	*28.6*	*35.4*	*39.8*
		(5.584)	(3.234)	(2.443)	(2.172)	(1.707)	(1.370)	(1.198)
	8	9	10	11	12	13	14	15
	419.0	—	—	—	—	—	—	—
	35.7							
	(1.890)							
model y	358.9, *50.9* (0.402)							
model v	431.0, *120.6* (1.259)							
model z	427.4, *120.5* (1.222)							

Table B.16: DIC, p_D and (MSE) values for extended models applied to structure U.

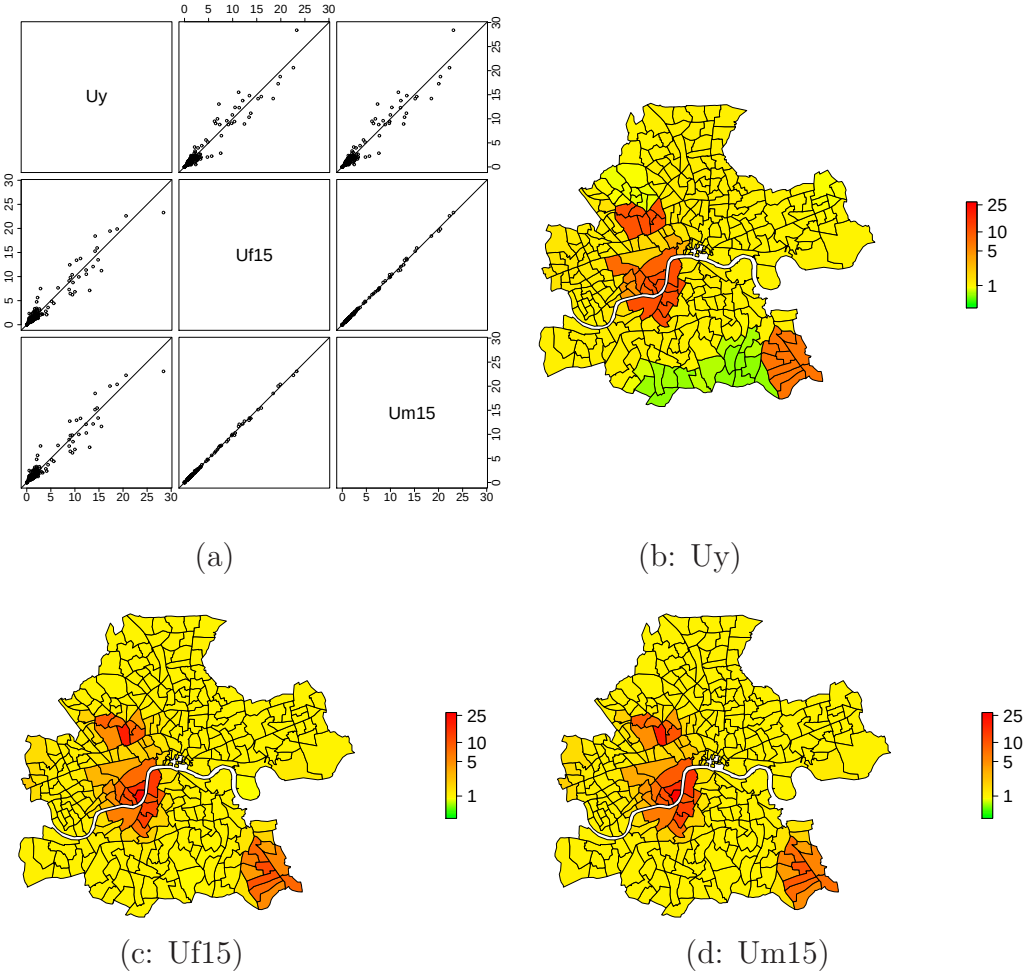

(a) (b: Uy)

(c: Uf15) (d: Um15)

Figure B.15: Results for structure U: Scatterplot matrix of the estimated parameters $\widehat{\Lambda}_i\, E_i$ of the BDCD model Uy and Poisson–Gamma models assuming additive (Uf15) and multiplicative influence of benzene (Um15), both include 15 Gaussian kernels; and spatial risk surfaces $\widehat{\Lambda}_i$ of model Uy (b), model Uf15 (c) and model Um15 (d).

Poisson–Gamma models give very similar risk surfaces. We detect only minor differences. In contrast, the comparison with BDCD reveals larger discrepancies. Most obvious, BDCD estimates a low risk region for some very southern wards that is not present from data generation. On the other hand, this model is able to estimate the almost constant high level risk of the three clusters correctly, Poisson–Gamma models fail to do so. High risk regions are identified correctly, but due to Gaussian kernels we cannot estimate a constant risk here. As generated Λ_i is large here, differences $\widehat{\Lambda}_i\, E_i - \Lambda_i\, E_i$ are more influenced by such high risk regions than underestimated regions by BDCD.

We expect Poisson–Gamma models to perform much better when allowing for alternative kernels as the models already lead to a satisfying fit of lower risk regions.

B.15 Structure V

Data generated by structure V is characterised by a multiplicative influence of benzene accounting for 770 cases. Additional 330 cases are generated in three cluster centers. A plot of the corresponding risk surface is given in Figure B.16 (a).

We model those data by Poisson–Gamma models with different settings. When assuming a multiplicative influence of benzene in our model best results are achieved when three latent covariates are included (model Vm3). The calculated DIC value equals 410.5, see Table B.17. The number of required Gaussian kernels therefore corresponds to the number of clusters in data generation. The estimated pattern of $\widehat{\Lambda}_i$ is given in Figure B.16 (c). It is very similar to the generated pattern. Nevertheless, comparison of the range of the values shows differences. While Λ_i is in the interval of $[0.75, 40.8]$, we estimate $\widehat{\Lambda}_i$ to be in $[0.83, 34.7]$. Inclusion of a fourth kernel leads to a worse model fit.

Poisson–Gamma models that assume an additive influence of benzene re-

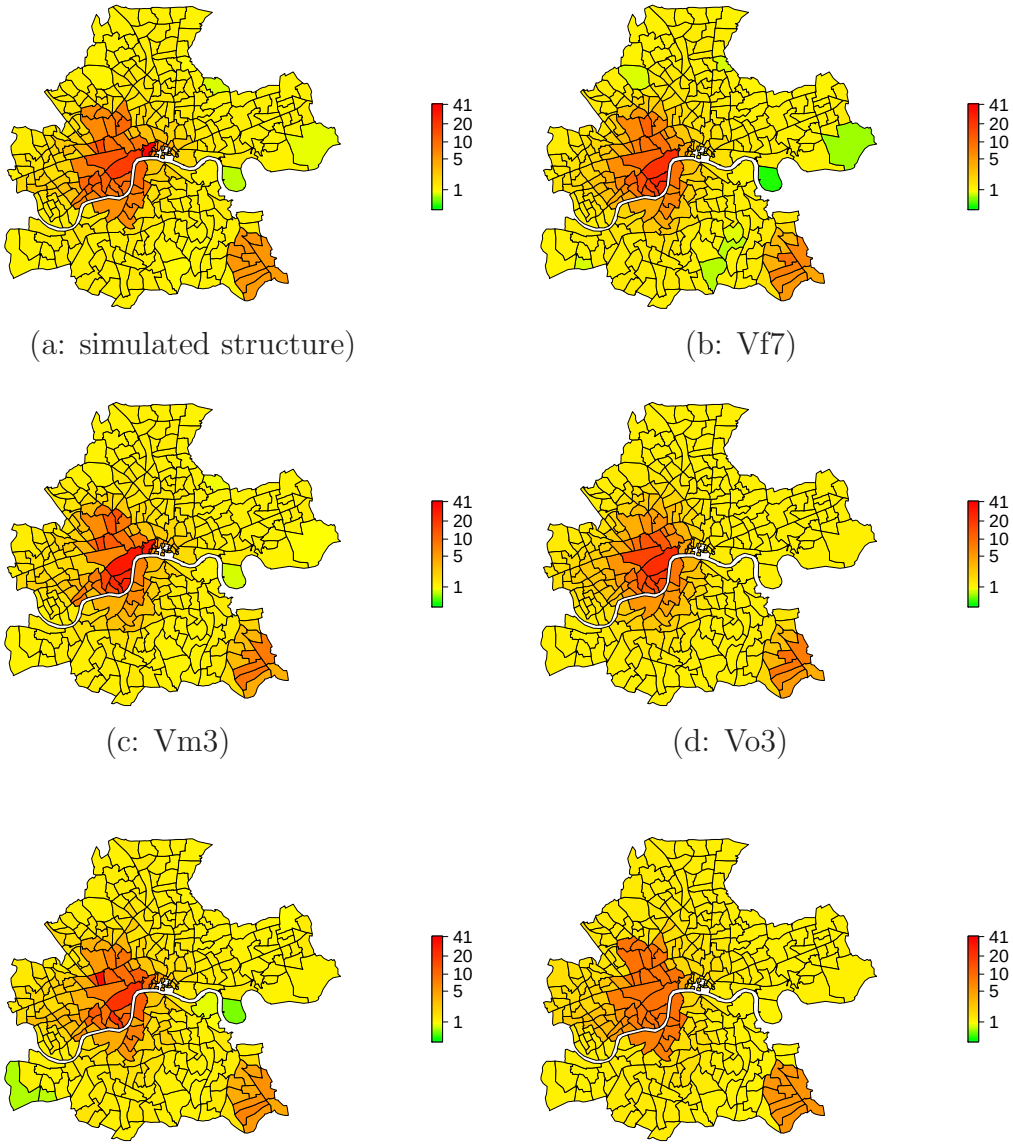

(a: simulated structure) (b: Vf7)

(c: Vm3) (d: Vo3)

(e: Vz) (f: Vy)

Figure B.16: Simulated pattern Λ_i for structure V (a) and results: estimated spatial pattern $\widehat{\Lambda}_i$ of the Poisson–Gamma model with additive influence of benzene and seven latent covariates (Vf7, b), multiplicative influence of benzene and three latent covariates (Vm3, c), without the influence of benzene and three latent covariates (Vo3, d), the MRF model where wards across the Thames are assumed to be neighbours (Vz, e), and the BDCD model (Vy, f).

♯ latent factors	0	1	2	3	4
model f	1077.5 (12.431)	963.5 (5.448)	445.1 (3.352)	427.0 (3.123)	421.8 (2.410)
	5	6	7	8	9
	418.9 (2.300)	418.7 (2.150)	418.3 (2.093)	418.1 (2.044)	—
model m	984.7 (11.210)	683.4 (7.841)	474.1 (3.921)	410.5 (2.838)	421.3 (2.902)
model o	—	1291.2 (7.398)	505.5 (4.321)	476.5 (3.790)	476.8 (3.474)
model y	389.4 (0.982)				
model v	434.9 (1.469)				
model z	433.0 (1.468)				

Table B.17: DIC (MSE) values for extended models applied to structure V.

quire a higher number of latent covariates to achieve similar results. Here, the number is more difficult to determine. For model Vf6 the DIC is 418.9. Inclusion of one additional covariate leads to a value of 418.3 (Vf7), two additional kernels lead to 418.1 (Vf8). The DIC values decrease very slowly here. Furthermore, the number of latent covariates is much higher compared to the multiplicative Poisson–Gamma model which has a lower DIC. Therefore, we do not include more latent covariates. In Figure B.16 (b) we plot the risk surface $\widehat{\Lambda}_i$ for model Vf7. Main differences to the underlying structure are in the low risk regions where the risk is underestimated. The maximum of the estimated values $\widehat{\Lambda}_i$ is 34.7. Higher DIC values are calculated for model Vo3 (Poisson–Gamma model with three latent covariates and no benzene), see Table B.16. Again, a model that includes three covariates corresponding to the number of clusters achieves the best results. Here, it is easy to determine that at least three kernels are required due to the huge drop in the DIC values, but we observe only a small increase in DIC afterwards.

MRF models show a highly acceptable model fit with a DIC of 433.0 (Vz) and 434.9 (Vv) points. The risk surface of model Vz is given exemplarily in Figure B.16 (e). We plot a similar surface when the neighbourhood structure

	V	Vf7	Vm3	Vo3	Vz	Vy
V	1	0.928	0.903	0.866	0.950	0.968
Vf7	0.928	1	0.971	0.963	0.950	0.915
Vm3	0.903	0.971	1	0.933	0.933	0.876
Vo3	0.866	0.963	0.933	1	0.910	0.863
Vz	0.950	0.950	0.933	0.910	1	0.950
Vy	0.968	0.915	0.876	0.863	0.950	1

Table B.18: Structure V: Pearsons correlation coefficient of the estimated parameters $\widehat{\Lambda}_i\, E_i$ of selected models.

of model Vv is used.

Lowest DIC values for this structure are achieved for the clustering algorithm where the DIC is 389.4, see Table B.17. The corresponding MSE equals 0.982. Differences occur in the maximal risk which is 22.41 here. Although the risk is underestimated, the cluster regions are well identified as high risk regions by the model, see Figure B.16 (f).

All models are suitable to identify the underlying risk structure correctly. They are able to reproduce the high risk in the selected wards. Sharpest distinction between high and low risk regions is possible by the BDCD model which has also the lowest DIC value. Recall that BDCD does not involve any covariate information. Table B.18 gives Pearsons correlation coefficient for the discussed models. Those are close to one. Lowest correlation between generated and modelled values is 0.866 (model Vo3) corresponding to the DIC values. Nevertheless this model still brings us to a suitable fit.

Inclusion of benzene leads to better estimates. Here multiplicative modelling uses three latent covariates only. The additive approach compensates disadvantages of treating benzene as an excess risk factor by the inclusion of seven latent covariates. This leads to a similar fit.