Statistics Netherlands

2016 Scientific Paper

Trend changes in tremor rates in Groningen

update April 2016

The views expressed in this paper are those of the author and do not necessarily reflect the policies of Statistics Netherlands.

Frank P. Pijpers

Contents

Nederlands

Deze rapportage behelst een voortzetting van onderzoek dat is uitgevoerd sinds midden 2014 in het kader een onderzoeksproject door het CBS in opdracht van Staatstoezicht op de Mijnen (SodM). Dit onderzoek is ten behoeve van een statistische onderbouwing van het meet- en regelprotocol voor gasexploitatie in de provincie Groningen.

Uit de analyse beschreven in eerdere halfjaarlijkse rapportages sinds December 2014, blijkt dat ongeveer 9 weken nadat de productie sterk was gereduceerd er een statistisch significante trendwijziging is opgetreden in de lineaire component van de bodemdaling gemeten met behulp van GPS stations. De dalingssnelheid is een factor 2.8 lager na de trendwijziging, die optreedt tussen ongeveer 15 Maart en 4 April 2014, vergeleken bij de periode daar direct aan voorafgaand. De trendwijziging kan zich geleidelijk gemanifesteerd hebben over een periode van enkele weken, en er is daarom een onzekerheid van ruwweg een week of twee over de centrale datum van deze overgang (Pijpers, 2014), (Pijpers and van der Laan, 2015a), (Pijpers and van der Laan, 2015b).

In deze update ligt de aandacht op een analyse van de tijden en locaties van aardbevingen die worden gerapporteerd door het Koninklijk Nederlands Meteorologisch Instituut (KNMI) gebaseerd op hun analyses van de gegevens verzameld door het netwerk van seismometers dat het KNMI beheert. Met behulp van een Monte Carlo analyse kan worden bepaald dat het aantal aardbevingen in het centrale deel van het veld na 23 Maart statistisch significant lager is dan het zou zijn geweest wanneer de trend van de periode daarvoor zou zijn voortgezet. Deze analyse is een uitbreiding van het onderzoek eerst gerapporteerd in 2014 met updates in 2015 in de zin dat aardbevingen tot eind Februari 2016 in de catalogus van het KNMI zijn meegenomen in de analyse.

Het uitgangspunt voor deze analyse is om zoveel mogelijk data gedreven te zijn en onafhankelijk van modellen. Dit betekent dat op basis van uitsluitend de analyse die hier wordt gepresenteerd een direct causaal verband tussen productievariaties en frequentie van bevingen noch aangetoond, noch verworpen kan worden.

English

This report is a continuation of research, commenced in 2014, which is part of a research project being carried out by Statistics Netherlands and commissioned by State Supervision of Mines (SodM). This research is part of the underpinning of the statistical methods employed to support the protocol for measurement and regulation of the production of natural gas in the province of Groningen.

From the analyses described in semi-annual updates since December 2014 in which GPS time series are used to measure subsidence, it can be determined that some 9 weeks after the strong reduction in production, there has been a significant change in the linear component of the rate of ground subsidence. The speed of subsidence is a factor of 2.8 lower after the break in the trend, between about March 15 and April 4 2014, compared to the period directly previous to that date. The changeover in the trend of ground subsidence can have become manifest gradually over a period of several weeks, which implies that the central date of the transition is also uncertain by roughly a week or two (Pijpers, 2014), (Pijpers and van der Laan, 2015a), (Pijpers and van der Laan, 2015b).

In this update, the focus is on an analysis of the times and locations of earthquakes as reported by the Royal Netherlands Meteorological Institute (KNMI) based on their processing of the network of seismometers that they manage. A Monte Carlo analysis is employed to demonstrate that the rate at which earthquakes occur in the central production field, after March 23 2014 is significantly lower than would be expected under a null hypothesis that the rate follows the same trend as before that date. This analysis is an update of the reports of December 2014 and 2015 in the sense that it uses earthquake data recorded by the KNMI up to the end of February 2016.

This analysis was purposely set up to be data driven and as much as feasible to remain model-independent. The consequence of this is that on the basis of the analysis presented here by itself, a direct causal relationship between production variations and the frequency of tremors can neither be proved nor disproved.

1 Introduction

For some decades earthquakes of modest magnitudes have occurred in the Groningen gas field. It is recognized that these events are induced by the production of gas from the field. Following an $M_L = 3.6$ event near Huizinge, and the public concern that this raised, an extensive study program has started into the understanding of the hazard and risk due to gas production-induced earthquakes.

A protocol needs to be established with the aim of mitigating these hazards and risks by adjusting the production strategy in time and space. In order to implement this regulation protocol and adaptively control production it is necessary also to measure the effects on subsidence and earthquakes in order to provide the necessary feedback.

The causality of the earthquakes induced by gas production is likely to be through the interaction of compaction of the reservoir rock with existing faults and differentiated geology of the subsurface layers. The ground subsidence occurs because with the extraction of gas, pressure support decreases in the layer from which the gas is extracted. The weight of overlying layers then compacts that extraction layer until a new pressure equilibrium can be established. The technical addendum to the winningsplan Groningen 2013 ”Subsidence, Induced Earthquakes and Seismic Hazard Analysis in the Groningen Field” (Nederlandse Aardolie Maatschappij BV, 2013) discusses all of these aspects in much more detail.

The seismic network of the KNMI has been in operation for some decades, and detailed reporting on and (complete) data for earthquakes in the Groningen region are available from 1991 onwards. The locations of all earthquakes in the region are shown in fig. 1.1, together with the locations of the gas production clusters. Also indicated are the boundaries of the regions for which the earthquake rates are determined in this report.

In this technical report, the available earthquake data are examined for a signature of changes in rates. The analysis procedure is unchanged from previous reports, and is presented as well as the conclusions one can draw from this phase of the research project.

Figure 1.1 The locations of earthquakes as reported by the KNMI. The red squares are locations of the production clusters, some of which are identified by name. The purple ellipse `zone large' demarks the reference area for earthquake rates. The red and green smaller ellipses (central and SW respectively) mark the two regions of interest also reported on in previous reports. The yellow ellipse (SE) is an additional region first considered in the report of Nov 2015. The production field is also shown in dark gray, overplotted on a map of the region

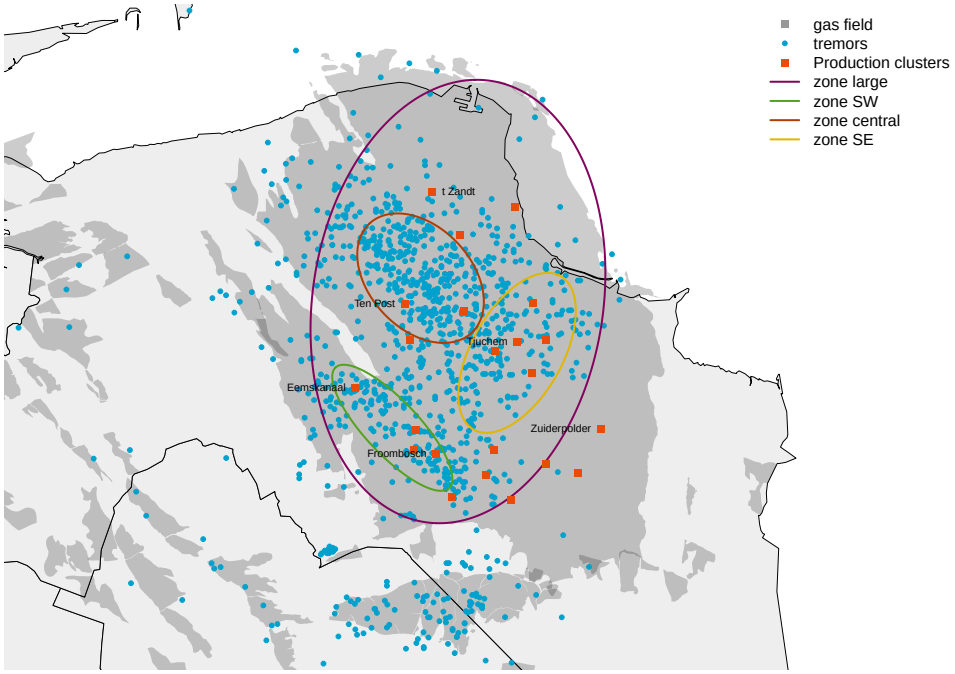

2 Background

2.1 The earthquake data

The available earthquake dataset contains in total 1237 events recorded after 1 Jan. 1991 up to 1 Mar. 2016. Of these, there are 848 that are located within the zone indicated as ‘zone large’ in fig. 1.1. An earthquake magnitude and time of event as well as the KNMI’s present best estimate of the longitude-latitude position is available for each of the earthquakes. The KNMI has indicated that the network of seismometers was designed to be complete in terms of both detection and localisation of earthquakes in the Groningen region above magnitudes of 1.5. Above magnitude 1.0 the detection is fairly secure, but the localisation may be more problematic. An upgrade to the network of seismometers is in the process of being implemented which should push down these limiting magnitudes for completeness and localisation. This may also have consequences for analyses such as these, because a better detection rate will imply that more will be recorded in the catalog which is only an apparent increase in the rate of tremors. The current network was only fully operational from 1995 onwards. This means that in order to ensure a reasonably uniform quality of the catalog, it is preferable to exclude all data from before Jan. 1 1995. With this restriction in time, there are data on 1182 earthquakes available. Of these, there are 811 within the large ellipse shown in fig. 1.1.

It is evident from fig. 1.1 that the distribution of events is not uniform over the area under consideration. It is also known that the distribution function of earthquakes is not uniform as a

Figure 2.1 The logarithmic cumulative magnitude distribution of earthquakes for all earthquakes in the set. The red line is a linear function with a slope of -0.9 similar to values reported elsewhere (Dost, B., Goutbeek, F., van Eck, T., Kraaijpoel, D., 2012). 95% confidence intervals are indicated under the assumption that the underlying process obeys Poisson statistics.

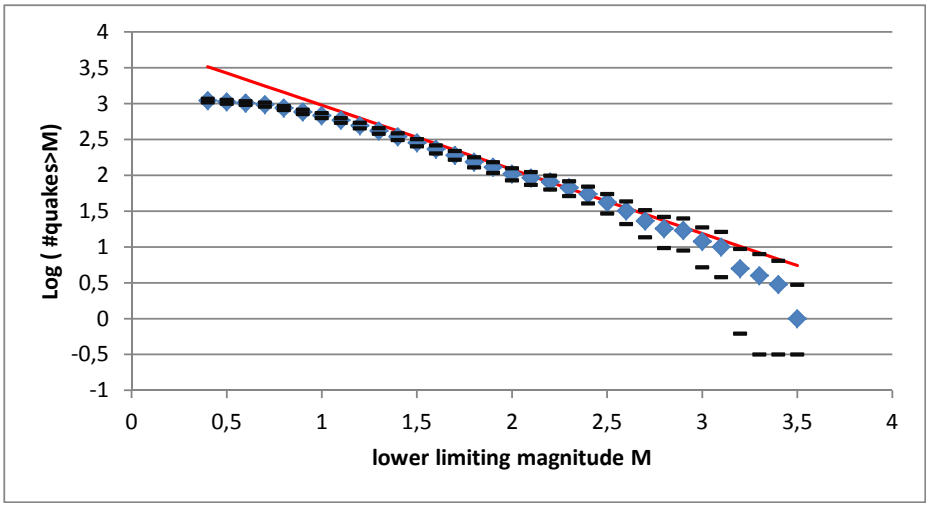

function of magnitude. For all 1182 quakes the distribution is shown in fig. 2.1. The way in which this is plotted is in a cumulative form: all earthquakes with a magnitude above a lower limit are counted and the base-10 logarithm of that count is shown as a function of the lower limiting magnitude. As this limiting magnitude increases there are fewer and fewer earthquakes with magnitudes above that limit, so this is a cumulative distribution function (or cdf) when reading the figure from right to left. This is a commonly used way to represent earthquakes in the field, known as frequency-magnitude or Gutenberg-Richter plot. The horizontal lines indicate margins of 95% confidence under the assumption that within each interval of the cumulative distribution in quake magnitude the value obeys Poisson statistics. The statistics of induced earthquakes is not well known, which implies that using margins of confidence from a particular probability distribution function such as the Poisson distribution may well be inappropriate. Towards higher values of the lower limiting magnitude, the margins of uncertainty become larger because there are fewer events on which to build the statistics.

Also shown in fig. 2.1 is a linear function with a slope of -0.9, ie. very close in value to the results of (Dost, B., Goutbeek, F., van Eck, T., Kraaijpoel, D., 2012) and an offset selected to match the range $1.1 < M_L < 3.1$. This shows that the slope of the distribution function appears to be constant over this range. For lower limiting magnitudes the distribution function is systematically lower than the straight line. The apparent 'deficit' of earthquakes with very low magnitudes is known to be indicative of the limitations of the sensitivity of the seismometer network. If tremors of such small magnitudes occur too far away from any of the seismometers in the network the signal becomes indistinguishable from noise or cannot be located with sufficient accuracy. For tremors with magnitudes below about 1.0 the 'missing' smaller earthquakes or tremors probably do occur but the detection of such events is no longer complete. The KNMI may use a higher value than this lower limit, such as 1.5, when taking into account not only the magnitude as is done here but also an accurate localisation of the events, which requires that a positive detection is available from at least 3 seismic wells in order to carry out the triangulation.

The catalog of quake events is likely also to contain events that are aftershocks. This means that some fraction of events has not occurred completely independently from preceding ones which

Figure 2.2 The logarithmic cdf-s of earthquakes for a selection of the 8 consecutive subsets, only every second set is shown for clarity. Upper panel: before normalization, lower panel: after normalization.

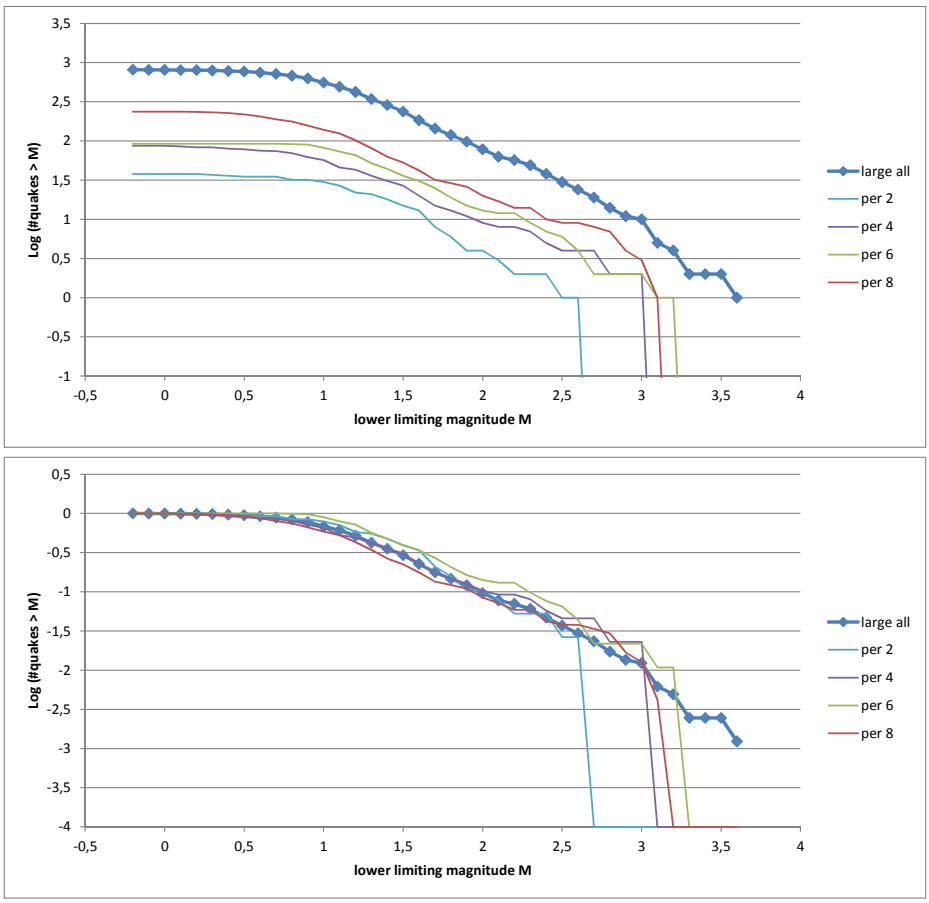

implies that it is inappropriate to assume Poissonian statistics. This is considered in more detail in section 4.

Further, in the strict sense a Poissonian process should be stationary in time, although if the change in the time parameter of the process is very slow a separation of time scales may allow modelling as a succession of Poisson processes with different time parameters. If the quakes were due to a stationary stochastic process this implies that if the dataset were to be divided into subsets that are equal in length of time, the number of quakes in each subset should not show any significant difference. To test the assumption, the dataset is divided into eight equal sections, and the logarithmic cdf is determined separately for each subset, numbered consecutively from 1 to 8. The resulting cdf-s are shown only for all the even numbered subsets to avoid overcrowding the plot. A similar plot for the odd numbered subsets would be much the same. Subdividing the dataset into a larger or smaller numbers of subsets has also been tested, which confirms that this behaviour does not depend on the precise boundaries between the divisions or the number of divisions used.

From fig. 2.2 it is clear that the number of earthquakes increases with time, and therefore the underlying process cannot be a stationary stochastic Poissonian process. In principle this might have been due to a detection effect: if over time the seismometer network has expanded, or the individual seismometers have become more sensitive due to upgrades, or the processing has improved to reduce noise levels, the dataset would contain more earthquakes towards later

Figure 2.3 The total number of earthquakes for each of the 8 consecutive subsets and an exponential fit to these points.

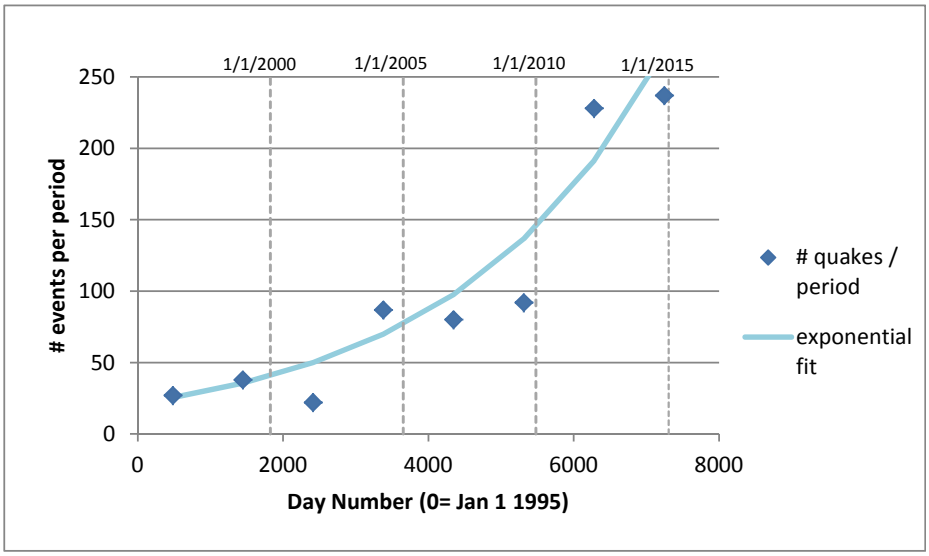

times, because more of the smaller earthquakes are being detected. However, if this were the case, the magnitude at which the cdf bends over should have moved progressively to lower magnitudes. Also, at magnitudes well above the completeness limit of around 0.8 the cdf-s should not show any trends with time if the process were stationary.

It is evident from fig. 2.2 that the shape of the distribution functions is very similar between the subsets, and that they appear to simply shift upwards from every period to the next. If the distribution functions are normalized, by dividing by the total number of quakes in each period, the lower panel of fig. 2.2 is obtained. From this it can be seen that there is very little change in the shape of the distribution functions. At the high end, with lower limiting magnitudes of around 2.5 and higher, the number of events is so small that there may simply be no such events in a given period. Hence the distribution functions for the different periods have different cut-offs at the high end. The near invariance of the shape of the distribution functions means that the increased rate is unlikely to be due to improved sensitivity of the network since 1995, and more likely to be due to a genuinely increasing rate of quake events. While further upgrades are currently in progress the effect of an increased sensitivity, in particular for the lower range of magnitudes is not yet very evident in the catalog, which is as would be expected since the upgrades are quite recent changes.

Fig. 2.3 shows the total number of events in each period and also a fit to these points of the form $A\exp(t/\tau)$. The characteristic timescale τ that is determined from a fit of the function to these data, indicates that the rate of quake events doubles roughly every 5.4 years. Both a least-squares and a maximum likelihood fitting has been performed, with the same result, within the uncertainty of 0.2 years, of the doubling time.

A straightforward method to analyse the behaviour of rate changes of tremor events would be to divide the time axis into sections of several hundred days (eg. half a year or a year), and for each section to count the number of events, with magnitudes above a fixed threshold. This is similar to what is done in fig. 2.3 but more fine-grained. Such a time series would have more sampling points than fig. 2.3 allowing applying standard time series analysis techniques. However, when this is done it becomes clear that the number of events per section is no higher than a few tens at best. This has the consequence that assessing the statistical significance of trend changes

becomes so sensitive to the unknown properties of the underlying distribution function, produced by the process that generates the tremors, that no meaningful conclusions can be drawn. For this reason such a straightforward approach was abandoned, and a Monte Carlo technique was adopted.

2.2 Monte Carlo simulations

The data indicate that the process by which the earthquakes arise is neither stationary in time, nor homogeneous in spatial distribution over the area. This prevents applying the statistics of Poissonian processes to asses whether in particular subregions the rate of earthquakes has altered, following the reduction in production. However, it is possible to use the dataset itself to test various hypotheses. This is done by means of a technique referred to in the literature as bootstrapping or Monte Carlo simulation. Extensive descriptions and applications of this technique can be found eg. in textbooks by (Robert and Casella, 2004), (Tarantola, 2004).

Since in each simulation all the 811 earthquakes are assigned, the same limitations apply to the simulations as apply to the real data: eg. as regards completeness, there are fewer events with low tremor magnitudes than would be expected on the basis of an extrapolation in the Gutenberg-Richter plot fig. 2.1. A close similarity in this aspect, between the artificial and real data, is an essential requirement for the method to function.

In the present case the technique is applied in order to test several hypotheses. The way one proceeds is to use the magnitude of the 811 events as recorded and reported by the KNMI. For the simulations, the location and timing of each event are not used. Instead locations and timings are assigned stochastically, using a random number generator and a pre-set likelihood for an event to belong to a certain group. In the present case there are eight relevant groupings, constructed by a subdivision in time and subdivisions in space :

1. A grouping in time : the event *either* occurs in the period from Jan. 1 1995 up to March 23 2014, *or* it occurs in the period from March 23 2014 to March 1 2016.
2. A grouping in space : the event occurs either within the contours of the area marked ‘zone SW’ in fig. 1.1, or within the area marked ‘zone central’, or within ‘zone SE’, or not in any of these regions, but within the area marked as ‘zone large’ in fig. 1.1.

The eight groups are obtained by events within each zone occurring in either the first or the second time range. There are several null hypotheses that are tested within the scope of this research. The most simple hypothesis is that, despite appearances, the probability for an event to occur is constant over the entire domain ‘zone large’ and also constant in time. Under this null hypothesis the likelihood for an event to occur within each of the three spatial groups is simply proportional to the area of each zone. Also, the likelihood for an event to occur in the first or the second of the two time ranges is proportional to the length of each range. The combined likelihoods are obtained assuming independence ie. by straightfoward multiplication of the likelihoods for the spatial divisions and for the division in time.

The next step is to assign each event (quake magnitude) to one of the eight groups using a random number generator twice: once to decide which of the spatial groups to assign the event to, and once to decide which period. After all 811 events are assigned, a cdf can be constructed for each group. This assigment process is repeated a large number of times, for the present case

10000 repetitions was considered sufficient, since there does not appear to be a need determine the simulated number of quakes N_{sim} and the standard deviation σ to more than 3 significant digits. Using these 10000 simulations an average distribution function for quake magnitudes can be constructed for each group, as well as 95% and 99% confidence limits, because each of the 10000 simulations will produce a different realisation from the stochastic assignment.

Other likelihoods than the ones described above can be assigned as well, giving rise to different null-hypotheses for testing. The measured / true distribution in space and in time of all 811 events can then be used in each case to test wether the null-hypothesis can be rejected or not. The total number of events for each group is shown in table 2.1. The following sections present the results for 5 separate null-hypotheses.

Note that by proceeding in this way, the only assumption that is made about the stochastic properties of the physical processes underlying the generation of earthquakes, is that the events are independent. The presence of considerable numbers of aftershocks in the catalog would violate this assumption. The consequences of that are explored in section 4. By using the bootstrapping technique it is possible to circumvent the necessity of having a spatiotemporal model for the generation of tremors and aftershocks. In particular, by using the earthquake magnitudes of the 811 actual events the distribution functions for magnitudes can be simulated. The detected total number of tremors in each zone can thus be compared directly with the percentiles of the Monte Carlo distributions for the total numbers which directly translates to whether given percentile confidence limits are exceeded, for each group, without requiring a model for the rate at which quakes with magnitudes of any particular strength will be produced. In all cases the proper limits (percentiles) for the probability distribution function as determined from the Monte Carlo simulations are used as (non-)rejection criterion.

2.3 Null-hypothesis I: homogeneous and stationary process

Table 2.1 The measured total number of quake events since Jan. 1 1995, for each group and the probabilities for assignment to each group for homogeneous and stationary test case.

region	period	Number of events	probability
zone SW	before 23-03-2014	88	0.0619
	after 23-03-2014	35	0.0062
zone central	before 23-03-2014	217	0.1024
	after 23-03-2014	39	0.0103
zone SE	before 23-03-2014	99	0.1123
	after 23-03-2014	31	0.0113
zone large	before 23-03-2014	239	0.6322
(not SW, SE or central)	after 23-03-2014	63	0.0635

From section 2.1 it does not appear very probable a-priori that quake events are spread uniformly over the area of interest and that there is no time dependence in the rate at which quakes occur. Nevertheless it is useful to present these results as a measure of the capability of the Monte Carlo approach to test hypotheses. Also, the relevant probabilities are a useful reference to asses by how much quake rates are enhanced or lowered in the other models.

Using the probabilities shown in table 2.1, the cdf-s of quake magnitudes are determined. The total number of events in each group can be compared directly, and tested for significance, with the true numbers shown in table 2.1. From the simulations a mean value and a standard

Table 2.2 Simulated number of quake events for each group for homogeneous and stationary test case, and standardised difference. Entries 'bef' and 'aft' in the column period refer to before and after the date of March 23, 2014. The columns 99%l and 99%u refer to the lower and upper 99% confidence levels, rounded to the nearest integer, and the column 95%l and 95%u refer to the same for the 95% confidence levels.

region	period	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	bef	50.2	6.89	5.5	33	37	64	69
	aft	5.0	2.22	13.5	0	1	10	12
zone central	bef	83.0	8.53	15.7	62	66	100	106
	aft	8.3	2.86	10.7	2	3	14	16
zone SE	bef	91.1	8.94	0.9	69	74	109	115
	aft	9.1	3.00	7.3	2	4	15	18
zone large (not SW, SE or central)	bef	512.7	13.78	−19.9	477	485	540	548
	aft	51.5	7.04	1.6	34	38	66	70

deviation can be determined, and also the 1%, 5%, 95%, and 99% percentiles of the distributions of total numbers of events.

The mean and standard deviation for the 10000 simulations are shown in table 2.2, as well as the standardised difference between the measured and simulated total number of quake events for each group. Any value larger than ~ 2 is statistically significant. While the distribution of the simulated data does not conform exactly to a normal distribution, the standardised differences are nevertheless a good enough indicator to see directly that, apart from the result for the group 'zone SE' before March 23 2014, this null hypothesis is strongly rejected. In all cases the proper limits (percentiles) for the probability distribution function as determined from the Monte Carlo simulations are used as (non-)rejection criterion. Here the null hypothesis is rejected at a confidence level of 99%.

2.4 Null-hypothesis II: non-homogeneous and stationary process

More of interest for the problem at hand is to test the null-hypothesis that the rate at which quakes occur has not changed with time, but that the spatial distribution of that rate is not homogeneous: there is an enhanced likelihood in the various regions of interest. Geophysical modelling of the subsurface and the response of existing fractures to pressure changes might in future enable predicting a rate, but at present the true probability is not known with high precision. For this operational reason in the Monte Carlo simulation the probability is assigned according to the proportions of the true total number of events in each region, combined for both before and after March 23 2014.

Comparing table 2.3 with table 2.1, the probability for quakes to occur within zone SW is now enhanced by a factor of ~ 2.2 over the homogeneous value, and for the zone central the probability is enhanced by a factor of roughly ~ 2.8. For zone SE there is a more modest enhancement of a factor of 1.3. Using these probabilities, shown in table 2.3, the cdf-s of quake magnitudes are again determined, following the same procedures as in section 2.3. The total number of events in each group which can be compared directly, and tested for significance, with the true numbers also shown in table 2.3. From the simulations the mean value and the standard deviation is shown in table 2.4.

Table 2.3 Probabilities for assignment to each group for non-homogeneous and stationary test case. For convenience the numbers of true events are repeated.

region	period	probability	Number of events
zone SW	before 23-03-2014	0.1378	88
	after 23-03-2014	0.0138	35
zone central	before 23-03-2014	0.2869	217
	after 23-03-2014	0.0288	39
zone SE	before 23-03-2014	0.1457	99
	after 23-03-2014	0.0146	31
zone large	before 23-03-2014	0.3384	239
(not SW, SE or central)	after 23-03-2014	0.0340	63

Table 2.4 Simulated number of quake events for each group for non-homogeneous and stationary test case, and standardised difference. The columns are as in table 2.3

region	period	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	bef	111.8	9.86	−2.4	87	93	131	137
	aft	11.2	3.27	7.3	4	5	18	20
zone central	bef	232.6	12.85	−1.2	201	208	258	266
	aft	23.3	4.71	3.3	12	15	33	36
zone SE	bef	118.1	10.02	−1.9	93	99	138	145
	aft	11.8	3.42	5.6	4	6	19	22
zone large	bef	274.5	13.47	−2.6	240	248	301	309
(not SW, SE or central)	aft	27.5	5.17	6.9	15	18	38	42

The mean and standard deviation for the 10000 simulations are shown, as well as the standardised difference between the measured and simulated total number of tremor events for each group. As one would expect this null-hypothesis is better in the sense that it is not rejected for more groups. However, there is still strong rejection of this hypothesis for the zones SW and SE in the period after March 23, as well as the region within zone large, outside of the zones SW and central. For all of these zones the simulation results indicate total numbers of simulated events that are somewhat too high before March 23 and then substantially too low after March 23.

2.5 Null-hypothesis III: non-homogeneous and exponentially increasing process

From the discussion in section 2.1 it is clear that the stationary null hypothesis also does not appear very realistic. Using the number of earthquakes recorded in each of the 8 successive periods discussed in section 2.1, one can re-assess the likelihood for earthquakes to occur after March 23 2014 by extending the trend over the past years. Using the fit shown in fig. 2.3, the likelihoods can be re-determined for each of the 6 groups and Monte Carlo simulations produced to test whether this time dependence, together with the same enhanced likelihoods in the regions of interest is consistent with the data. Comparing the probabilities for a quake to occur after March 23 from table 2.5 with the probabilities of the previous section (table 2.3), shows that this probability is now higher by a factor of roughly 2.7.

From the final five columns in table 2.6 it can be seen that for most regions this null hypothesis cannot be rejected, with one exception. The zone central, in the period after March 23 when the GPS data indicate that the subsidence rate was reduced by a statistically significant amount, appears not to follow the increasing trend in the tremor rate. The actual number of quakes is

Table 2.5 Probabilities for assignment to each group for non-homogeneous and exponentially increasing test case. For convenience the numbers of true events are repeated.

region	period	probability	Number of events
zone SW	before 23-03-2014	0.1149	88
	after 23-03-2014	0.0367	35
zone central	before 23-03-2014	0.2392	217
	after 23-03-2014	0.0764	39
zone SE	before 23-03-2014	0.1215	99
	after 23-03-2014	0.0388	31
zone large	before 23-03-2014	0.2822	239
(not SW, SE or central)	after 23-03-2014	0.0902	63

Table 2.6 Simulated number of quake events for each group for non-homogeneous and exponentially increasing test case, and standardised difference. The columns are as in table 2.3

region	period	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	bef	93.2	9.08	-0.6	71	76	111	117
	aft	29.8	5.36	1.0	17	20	41	44
zone central	bef	194.0	12.17	1.9	163	170	218	226
	aft	62.0	7.50	-3.1	43	48	77	82
zone SE	bef	98.5	9.32	0.1	75	81	117	123
	aft	31.5	5.46	-0.1	18	21	42	46
zone large	bef	228.9	12.82	0.8	197	204	254	262
(not SW, SE or central)	aft	73.1	8.17	-1.2	52	57	89	94

lower by a statistically significant amount compared to the general increasing trend. The null hypothesis at least for this group is rejected at a confidence level of 99%.

2.6 Null-hypothesis IV: non-homogeneous and non-exponential time increasing process

While the non-homogeneous and exponentially increasing rate of events from section 2.5 appears to be a reasonable representation of the true rate of events, in the sense that it is not rejected, it does produce a total number of quake events, after March 23, of 196 which is higher than the 168 that are counted in the four spatial regions combined. Since the predicted probability is in some sense an extrapolation of the fitting function shown in fig. 2.3, it could be argued that this over-estimates the true actual rate. One can therefore instead take the view that a better estimate of the true probability is obtained by taking the proportions of the actual number of events, analogously to what is done in section 2.4 for the different spatial regions.

The probabilities and simulated events are shown in tables 2.7 and 2.8. Comparing with the probabilities from the previous sections it is clear that the probability for an event to fall after March 23 is still enhanced by a factor of roughly 2.3 over the stationary case, and slightly lower than is predicted by the exponentially increasing rate of section 2.5.

By construction, the mean value of the simulated total number of events after March 23 is now virtually equal to the number of actually recorded events. There is one region where the difference between the true number of events and the simulated value is large enough to doubt the validity of the null hypothesis. For zone central, in the period after March 23 2014, the null

Table 2.7 Probabilities for assignment to each group for non-homogeneous and non-exponentially increasing test case. For convenience the numbers of true events are repeated.

region	period	probability	Number of events
zone SW	before 23-03-2014	0.1202	88
	after 23-03-2014	0.0314	35
zone central	before 23-03-2014	0.2503	217
	after 23-03-2014	0.0654	39
zone SE	before 23-03-2014	0.1271	99
	after 23-03-2014	0.0332	31
zone large	before 23-03-2014	0.2952	239
(not SW, SE or central)	after 23-03-2014	0.0771	63

Table 2.8 Simulated number of quake events for each group for non-homogeneous and non-exponentially increasing test case, and standardised difference. The columns are as in table 2.3

region	period	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	bef	97.5	9.24	−1.0	75	80	116	122
	aft	25.5	4.95	1.9	14	16	36	39
zone central	bef	202.9	12.31	1.1	172	179	227	236
	aft	53.0	6.97	−2.0	36	40	67	72
zone SE	bef	103.1	9.52	−0.4	79	85	122	128
	aft	26.9	5.08	0.8	15	17	37	41
zone large	bef	239.5	12.97	−0.0	207	214	265	273
(not SW, SE or central)	aft	62.6	7.71	0.1	44	48	78	83

hypothesis is rejected at a 95% confidence level. For the zone SW the actual number lies just within the 95% confidence interval. For this region the exponential increase produces values that are very slightly better centred around the detected number of tremors but formally it is not possible to distinguish whether an exponential rate increase is a better model than the slightly lower increase with time used in this section. Even compared to this somewhat lower trend increase in rate of quake events, the zone central appears to have experienced fewer quakes than the statistical model predicts.

2.7 Null-hypothesis V: non-homogeneous and exponential time increasing process except for zone central

Instead of opting for an overall slightly lower increase with time as discussed in section 2.6, one can instead choose to leave the overall increase as in section 2.5 except for the central zone after March 23. If for this zone, and this time frame, one chooses a rate that is stabilised at the same level that it was on March 23 without increasing further, the probability for earthquakes to happen for the entire period after March 23 in that zone is lowered.

Comparing with the probabilities from the previous sections it is clear that in the central zone the probability for an event to fall after March 23 is still slightly enhanced over the stationary case, but it is not as high as is predicted by the exponentially increasing rate of section 2.5.

The mean value of the simulated total number of events after March 23 is slightly larger than the number of actually recorded events. Once again the zones SW and central remain where the

Table 2.9 Probabilities for assignment to each group for non-homogeneous and exponentially increasing test case, except for zone central where the rate stabilises after March 23 2014. For convenience the numbers of true events are repeated.

region	period	probability	Number of events
zone SW	before 23-03-2014	0.1202	88
	after 23-03-2014	0.0314	35
zone central	before 23-03-2014	0.2484	217
	after 23-03-2014	0.0673	39
zone SE	before 23-03-2014	0.1271	99
	after 23-03-2014	0.0332	31
zone large	before 23-03-2014	0.2952	239
(not SW, SE or central)	after 23-03-2014	0.0771	63

Table 2.10 Simulated number of quake events for each group for non-homogeneous and exponentially increasing test case, except for zone central where the rate stabilises after March 23 2014, and standardised difference. The columns are as in table 2.3

region	period	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	bef	97.5	9.24	-1.0	75	80	116	122
	aft	25.5	4.94	1.9	13	16	35	39
zone central	bef	201.4	12.29	1.3	171	178	226	233
	aft	54.5	7.13	-2.2	37	41	69	74
zone SE	bef	103.1	9.50	-0.4	79	85	122	128
	aft	26.9	5.07	0.8	15	18	37	40
zone large	bef	239.4	12.97	-0.0	206	214	266	273
(not SW, SE or central)	aft	62.6	7.63	0.1	44	48	78	83

difference, between the true number of events and the simulated value, is large enough to reject the null hypothesis at the 95% confidence level, for the period after March 23.

This means that the zone central appears to have experienced fewer quakes than the statistical model predicts, even if the rate is assumed to have stabilised at the level of March 23. For zone SW the present rate of tremors may indicate that the doubling time scale for this region is perhaps slightly shorter than the 5.4 years deduced in sect 2.1. Given that the number of tremors for the period after March 23 2014 is right on the 95% confidence level border, it appears premature to draw such a conclusion.

To assess whether perhaps any increase in zone SW is of recent origin, the same analysis has been repeated, but rather than counting events before and after March 23 2014, the events are grouped by having occurred before or after January 1 2015. However, the total number of events is very small, and the same hypothesis for the tremor rate behaviour used in this section is in this case not rejected for any zone.

3 The influence of incompleteness

3.1 excluding tremors with magnitudes below 1

In the previous sections the full catalog of events is used, including a range of low magnitudes where it is likely that not all events have been detected. It is argued above that the bootstrapping procedure does not require that the catalog be complete. This is correct as long as the detection likelihood for tremors with small magnitudes is spatially and temporally sufficiently uniform. In this case the averaged sensitivity over the 4 regions and/or the two epochs is sufficiently similar that the likelihoods as quoted in the tables of the previous sections correspond to the likelihoods of occurrence. If the detection likelihood is spatially very inhomogeneous but over very small scales, averaging over larger regions could result in the same as for a uniform sensitivity. The current distribution of deep wells is likely to produce modest variations in detection likelihood over the region for tremors below magnitudes of 1. The upgrading of the seismic network implies addition of extra wells with measurement equipment which improves sensitivity and therefore affect in particular the detection rate of lower magnitude tremors. For these reasons it is worthwhile to redo the analysis for a subset of tremors taking into account a lower limiting magnitude that excludes the range where (past) incompleteness can influence the statistics.

Table 3.1 The measured number of tremor events since Jan. 1 1995 with magnitude 1.0 or higher, for each spatial group and also tremors with magnitudes 1.5 or higher.

region	period	Number of events	
		$M > 1$	$M > 1.5$
zone SW	before 23-03-2014	57	22
	after 23-03-2014	21	8
zone central	before 23-03-2014	180	94
	after 23-03-2014	17	6
zone SE	before 23-03-2014	68	20
	after 23-03-2014	16	5
zone large	before 23-03-2014	164	68
(not SW, SE or central)	after 23-03-2014	31	14

In section 2.1 (fig. 2.3 and the pertaining discussion) a characteristic timescale τ is determined, indicating that the rate of quake events doubles roughly every 5.4 years, with an uncertainty of 0.2 years, of the doubling time. The same analysis when performed after selecting only the tremors with magnitudes 1 or higher, produces a doubling time of 6.2 years. Further restricting

Table 3.2 Probabilities for assignment to each group for non-homogeneous and exponentially increasing test case, except for zone central where the rate stabilises after March 23 2014. Only tremors with magnitudes > 1, or with magnitudes greater than 1.5 .

region	period	probability	
		$M > 1$	$M > 1.5$
zone SW	before 23-03-2014	0.1197	0.1094
	after 23-03-2014	0.0211	0.0171
zone central	before 23-03-2014	0.3080	0.3708
	after 23-03-2014	0.0476	0.0511
zone SE	before 23-03-2014	0.1289	0.0912
	after 23-03-2014	0.0227	0.0143
zone large	before 23-03-2014	0.2993	0.2991
(not SW, SE or central)	after 23-03-2014	0.0527	0.0469

Table 3.3 Simulated number of quake events, with magnitudes > 1 (top block), or > 1.5 (bottom block), for each group for non-homogeneous and exponentially increasing test case, except for zone central where the rate stabilises after March 23 2014, and standardised difference. The columns are as in table 2.3

region	period	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	bef	66.3	7.68	−1.2	47	52	82	86
	aft	11.7	3.38	2.8	4	6	19	21
zone central	bef	170.6	10.89	0.9	143	149	192	199
	aft	26.4	4.98	−1.9	15	17	36	40
zone SE	bef	71.5	7.98	−0.4	51	56	87	92
	aft	12.6	3.48	1.0	5	6	20	22
zone large (not SW, SE or central)	bef	165.8	10.78	−0.2	138	145	187	195
	aft	29.2	5.27	0.3	17	19	40	44
zone SW	bef	26.0	4.89	−0.8	14	17	36	39
	aft	4.1	1.99	2.0	0	1	8	10
zone central	bef	87.9	7.48	0.8	69	74	103	107
	aft	12.1	3.34	−1.8	4	6	19	21
zone SE	bef	21.6	4.43	−0.4	11	13	30	34
	aft	3.4	1.83	0.9	0	0	7	9
zone large (not SW, SE or central)	bef	70.9	7.07	−0.4	53	57	85	90
	aft	11.1	3.26	0.9	4	5	18	20

the analysis to include only those tremors with magnitudes 1,5 or higher produces a value for the doubling time of 7.3 years. The uncertainty for these latter two slopes is higher, so that the difference between them is not statistically significant. The shorter timescale deduced when the limiting magnitude is set to lower values may in part be due to improvements in detection of smaller tremors. At face value it cannot be excluded that the number of smaller tremors has increased more quickly which would imply that slope of the distribution in the Gutenberg-Richter plot would also genuinely have become slightly steeper over this time.

First, the hypothesis is tested corresponding to the hypothesis II discussed in section 2.4, where the spatial inhomogeneity is determined from the relative numbers of tremors with magnitudes of 1.0 or higher in each region. This spatial enhancement factor is 2.1 for the SW zone, and for the central zone becomes 3.2, whereas for zone SE it is 1.2. This hypothesis is strongly excluded, at 99% confidence, because the number of detected events after March 23 is higher than this hypothesis produces. As a second option the equivalent of hypothesis V of section 2.7 is tested, for which the probabilities and results are shown in tables 3.2 and 3.3. The time scale used is the time scale appropriate for the subset of tremors: 6.2 years for tremors with $M > 1$. This hypothesis is rejected at the 99% confidence level for zone SW and at the 95% confidence level for the central zone after March 23 2014.

3.2 excluding tremors with magnitudes below 1.5

While fig. 2.1 appears to indicate that incompleteness becomes a serious issue only below magnitudes of 1, it is known that tremors with magnitudes between 1 and 1.5, although they are detected, are often difficult to localise because the signal exceeds the noise at only 1 or 2 seismic wells which means standard triangulation is impossible. The lower resulting spatial accuracy of the catalog at these magnitudes might also influence the statistics. For this reason the analysis is repeated, excluding all tremors with magnitudes below 1.5. If only the best-localised tremors

with magnitudes above 1.5 are taken into account, the spatial enhancements in the zones SW and central become 1.9 and 3.7 respectively, but for zone SE it is now 0.9 (ie. a lowering rather than an enhancement). Now, the equivalent of hypothesis II can be rejected at a 99% confidence level for zone SW, and at the 95% confidence level for the larger region, (without zones SW, SE, and central). Hypothesis V, using a doubling time of 6.2 years, is rejected at the 95% confidence level for zones SW and central after March 23 (see tables 3.2 and 3.3).

4 The influence of aftershocks

It is possible that some of the tremors in the catalog, even at magnitudes higher than 1 or 1.5, are events that are triggered by preceding tremors. This means that there is some finite correlation, both in time and in space, in the likelihood for a tremor to occur. This likelihood for a tremor to occur close in time and space to a previous tremor is then slightly in excess of what it would be if each event occurred completely independently from all previous events. This would mean that the fluctuations around a mean trend or inhomogeneities in spatial distribution are somewhat higher than a random assignment simulation produces. Conversely, the confidence limits used to determine whether a particular deviation is statistically significant must then be appropriately enlarged, from what is obtained from simulations that do not take correlations into account.

The Monte Carlo simulations used for this paper do not have such an excess of correlation. In principle it would be possible to introduce this, for instance through adding a Markov chain process to the simulations, with a finite probability for a tremor to be flagged as an aftershock in the simulations, and then assigned an appropriate location and time relatively close to the preceding tremor rather than completely at random. However, this would require a knowledge of the likelihood for an earthquake of a given strength to produce an aftershock, and distribution functions for the distances and times between progenitor and aftershocks. Relevant methods of analysis reported in the literature are (Huc, M., Main, I.G., 2003) and (Naylor, M., Main, I.G., Touati, S., 2009), or a modelling approach for aftershock generation (Kumazawa and Ogata, 2014) to simulate data. Progress on the analysis using these methods is to be reported at a later stage.

An alternative approach is to exclude from the catalog any event that is sufficiently close in space and in time to a preceding event, so that one might reasonably suppose that it could be an aftershock. There is some arbitrariness in the possible choice of parameters but to operationalise this exclusion criterion, shocks are designated as aftershocks, and therefore excluded from the analysis, if they occur within 10 days of an earlier entry in the catalog, and are located within a distance of 5 km of that entry. The numbers of events remaining after applying the exclusion criteria are shown in table 4.1.

The probabilities for an increasing rate in accordance with the exponential trend in numbers of events are shown in table 4.1, and the simulation results in table 4.2. This model is similar to the one discussed in section 2.7, but excluding potential aftershocks, and also excluding tremors with magnitudes below 1. This model is rejected at the 95% confidence level, because in zone SW after March 23 2014 more tremors have been detected than the hypothesis produces. This means that excluding potential aftershocks in this way still leads to rejection of hypothesis V but at a lower level of confidence. Just as is described in section 2.7, to assess whether perhaps any increase in zone SW is of recent origin, the same analysis has been repeated but grouping the

Table 4.1 Probabilities for assignment to each group for non-homogeneous and exponentially increasing test case, except for zone central where the rate stabilises after March 23 2014. Only tremors with magnitudes > 1, and excluding all events within 10 days and 5 km of a previous event recorded in the catalog. The relevant recorded number of events are given in the final column.

region	period	probability	Number of events
zone SW	before 23-03-2014	0.1181	41
	after 23-03-2014	0.0208	14
zone central	before 23-03-2014	0.2800	117
	after 23-03-2014	0.0433	11
zone SE	before 23-03-2014	0.1202	44
	after 23-03-2014	0.0212	12
zone large	before 23-03-2014	0.3371	138
(not SW or central)	after 23-03-2014	0.0594	19

Table 4.2 Simulated number of quake events, with magnitudes > 1 and excluding all events within 10 days and 5 km of a previous event recorded in the catalog, for each group for non-homogeneous and exponentially increasing test case and standardised difference. The columns are as in table 2.3

region	period	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	bef	46.7	6.37	-0.9	31	35	59	64
	aft	8.2	2.84	2.0	2	3	14	17
zone central	bef	110.8	8.93	0.7	88	93	128	134
	aft	17.2	4.02	-1.5	8	10	25	28
zone SE	bef	47.7	6.63	-0.6	31	35	61	65
	aft	8.4	2.85	1.3	2	3	14	16
zone large	bef	133.6	9.43	0.5	110	115	152	158
(not SW, SE or central)	aft	23.5	4.71	-1.0	12	15	33	36

events by having occurred before or after January 1 2015. The total number of events is small, and the same hypothesis for the tremor rate behaviour used in this section is in this case not rejected for any zone.

In sum there is no strong evidence that the existence of aftershocks, and the correlation between events that this produces, affects the data to such a large extent that the confidence limits produced by the Monte Carlo simulations are a severe underestimate. Thus the comparison of the results from hypothesis III and hypothesis V appear to point to a genuine change in the rate of generation of tremors.

5 Conclusions

From the analysis presented in this report, it can be concluded that spatially there is a statistically significant enhancement of the earthquake rate in two zones, SW and central, where gas production takes place, compared to the surrounding region, by factors of around 2.2 and 2.8 respectively. If only the best-localised tremors with magnitudes above 1.5 are taken into account the enhancements become 1.8 and 3.8 respectively. The region SE does not show such enhancements. For all regions SW, SE, and central, as well as the area directly surrounding these regions of particular interest, there is an increasing trend in the earthquake rate with time since Jan. 1 1995, which can be fit with an exponential increase with a doubling time of ~ 5.4 years.

Including only tremors with magnitudes $M > 1$, where completeness is more assured, leads to a slightly longer doubling time of 6.2 years.

Also in the most recent 23 months, since March 23 2014, the data are consistent with this spatial and temporal behaviour except for the central zone, where production was reduced substantially within weeks from the beginning of January 2014, and where the subsidence measured using GPS data appears to indicate a break in the trend after around March 23 2014. For this region the hypothesis of a continued increasing rate is rejected at the 99% confidence level. In this region the number of earthquakes is significantly lower than such a trend would produce, and is consistent with the hypothesis that the rate has reduced below the level of March 23 2014. While causality can neither be proved nor disproved on the basis of this research on its own, at present the data are consistent with the hypothesis that the reduced production has had a localised effect of making this earthquake rate not just rise less quickly than in the surrounding region but even drop somewhat. In the regions to the south (SE and SW) where production has not been reduced substantially, the exponential increase appears to continue. In region SW there are now some tentative indications that the rate increase is even slightly above the exponential trend with the doubling time quoted above. In other words the time scale for increase might now be somewhat shorter in the zone SW than elsewhere. Re-doing the analysis, but using a more recent date of Januari 1 2015 as dividing point in time does not confirm such a shortening of timescale, however. A more firm conclusion must await more tremor data, as well as a better assessment of the effects of an improved sensitivity due to recent and ongoing upgrades of the seismic measurement network.

References

Dost, B., Goutbeek, F., van Eck, T., Kraaijpoel, D. (2012). Monitoring induced seismicity in the north of the netherlands: status report 2010; wr 2012-03. Technical report, KNMI.

Huc, M., Main, I.G. (2003). Anomalous stress diffusion in earthquake triggering: correlation length, time dependence, and directionality. *J. Geophys. Res. 108 (B7)*, 2324.

Kumazawa, T. and Y. Ogata (2014). Nonstationary etas models for nonstandard earthquakes. *Ann. Appl. Stat. 8*, 1825--1852.

Naylor, M., Main, I.G., Touati, S. (2009). Quantifying uncertainty on mean earthquake inter-event times for a finite sample. *J. Geophys. Res. 114 (B0)*, 1316.

Nederlandse Aardolie Maatschappij BV (2013). A technical addendum to the winningsplan Groningen 2013 subsidence, induced earthquakes and seismic hazard analysis in the Groningen field. Technical report.

Pijpers, F. (2014). Phase 0 report 1 : significance of trend changes in ground subsidence in Groningen. Technical report, Statistics Netherlands.

Pijpers, F. and D. van der Laan (2015a). Phase 1 update may 2015: trend changes in ground subsidence in Groningen. Technical report, Statistics Netherlands.

Pijpers, F. and D. van der Laan (2015b). Trend changes in ground subsidence in Groningen update november 2015. Technical report, Statistics Netherlands.

Robert, C. and G. Casella (2004). *Monte Carlo Statistical Methods*. Springer.

Tarantola, A. (2004). *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM.

Publisher
Statistics Netherlands
Henri Faasdreef 312, 2492 JP The Hague
www.cbs.nl

Prepress: Statistics Netherlands, Grafimedia
Design: Edenspiekermann

Information
Telephone +31 88 570 70 70, fax +31 70 337 59 94
Via contact form: www.cbs.nl/information

Where to order
verkoop@cbs.nl
Fax +31 45 570 62 68
ISSN 1572-0314