Scientific Paper

Trend changes in tremor rates in Groningen

Frank P. Pijpers, Vincent van Straalen

May 2018

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.

In dit rapport 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. Deze analyse is een uitbreiding van het onderzoek eerst gerapporteerd in 2014, met actualisaties in 2015, 2016 en meest recent in november 2017. Voor deze rapportage zijn aardbevingen tot eind april 2018 in de catalogus van het KNMI meegenomen in de analyse. In deze actualisatie ligt de focus op de verschillen in aardbevingsfrequentie voor en na de datum van 1 januari 2015. Deze datum is gekozen omdat de totale maandelijkse productie door gaswinning vanaf begin 2015 minder in de tijd is gevarieerd dan in de periode daarvoor. Voor een aantal gaswinningsclusters is de productie al langer vlak en laag. Met behulp van een Monte Carlo analyse kan worden bepaald dat het aantal aardbevingen na 1 januari 2015 statistisch significant lager is dan het zou zijn geweest wanneer de trend van de periode daarvoor zou zijn voortgezet.

Het uitgangspunt voor deze analyse is om zoveel mogelijk data gedreven te zijn en onafhankelijk van modellen. In combinatie met de eerder gepubliceerde analyses is een direct causaal verband tussen productievariaties en frequentie van bevingen voldoende plausibel als werkhypothese, maar de analysestappen worden uitgevoerd zonder gebruikmaking van deze hypothese.

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.

In this report, 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. This analysis is an update of the reports of 2014, 2015, 2016 and most recently November 2017, where it uses earthquake data recorded by the KNMI up to the end of April 2018. In this update the focus lies on differences in tremor frequencies before and after January 1 2015. This date is chosen because the total monthly gas production in the epoch starting in 2015 has varied much less in time than in the previous epoch. For a few specific clusters the production has been flat and low for longer. A Monte Carlo analysis is employed to demonstrate that the rate at which earthquakes occur after January 1 2015 is significantly lower than would be expected under a null hypothesis that the rate follows the same trend as before that date.

This analysis was purposely set up to be data driven and as much as feasible to remain model-independent. In combination with previously published analyses, a direct causal connection between production variations and tremor frequencies is sufficiently plausible that

it may be used as a working hypothesis. However, the various steps in the analysis reported here do not require making use of this hypothesis.

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.

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

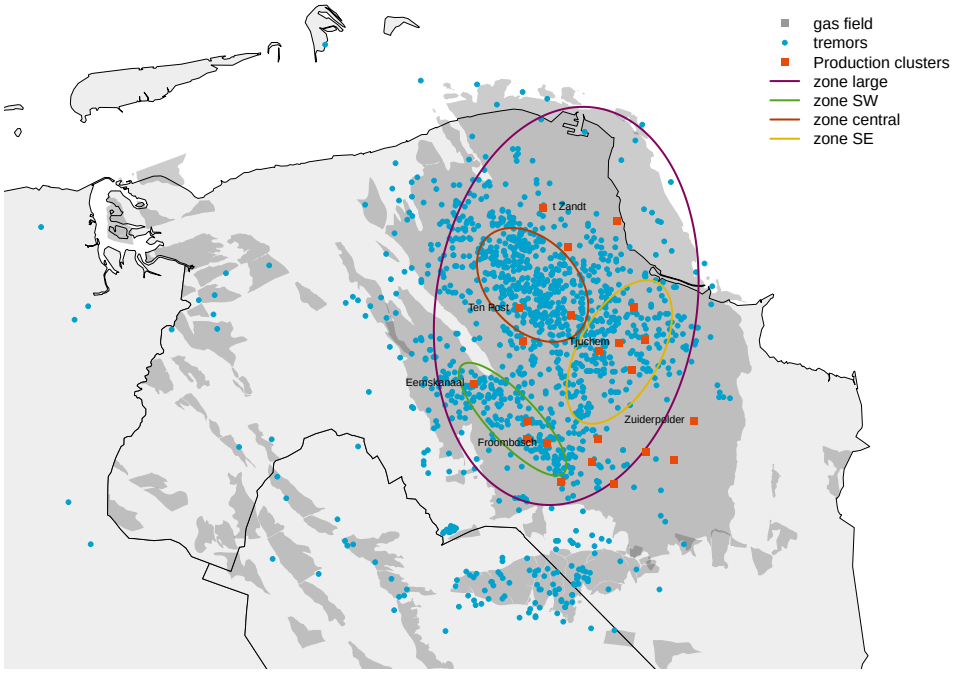

The causality of the earthquakes induced by gas production is likely to be through the interaction of compaction of the reservoir rock with existing faults and differentiated geology of the subsurface layers. The ground subsidence occurs because with the extraction of gas, pressure support decreases 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, cf. Dake

(1978); Doornhof et al. (2006). 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 the context of the Groningen reservoir 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, which are the same as in the previous semi-annual updates (cf. Pijpers (2014, 2015a,b, 2016a,b); Pijpers and van Straalen (2017a))

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, cf. Pijpers (2014, 2015a,b, 2016a,b); Pijpers and van Straalen (2017a), and is presented as well as the conclusions one can draw from this phase of the research project. In a sense this is therefore a classical approach, as opposed to for instance the Bayesian approach reported in Nepveu et al. (2016).

2 Background

2.1 The earthquake data

The available earthquake dataset contains in total 1502 events recorded after 1 Jan. 1991 up to 1 May 2018. Of these, there are 1066 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 in the '90s 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 likelihood is near 100%, but the localisation may be more problematic. The elliptical contour of the localisation uncertainty progressively increases in size towards lower magnitudes and also is not uniform in orientation or size for different subregions, depending on the distances to the nearest seismic stations.

Starting towards the end of 2014 an upgrade to the network of seismometers has been implemented which has pushed down these limiting magnitudes for completeness and localisation. That may also have consequences for analyses such as reported here, 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. In previous analyses carried out by CBS cf. Pijpers (2014, 2015a,b, 2016a), all data collected since 1995 has been used, but there might be some additional issues with completeness for the earlier years. While it is unlikely that such issues, even if present, materially affect those analyses, for the analysis reported here all data from before Jan. 1 2003 are excluded. This leaves a total of 1218 tremors of which 943 occurred within the ellipse 'zone large' that can be included in the analysis. Furthermore, comparisons between epochs, and also the tests of the various hypotheses discussed here, will be done taking only tremors with magnitudes greater than $M = 1$ into account. While this is therefore slightly lower than the completeness limit quoted, the data itself suggest that incompleteness becomes

severe only at lower magnitudes than $M = 1$. With this restriction in time and magnitude, there are data on 739 $M \geq 1$ earthquakes available. Of these, there are 564 within the large ellipse shown in fig. 1.1.

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 et al., 2012). 95% confidence intervals are indicated under the assumption that the underlying process obeys Poisson statistics.

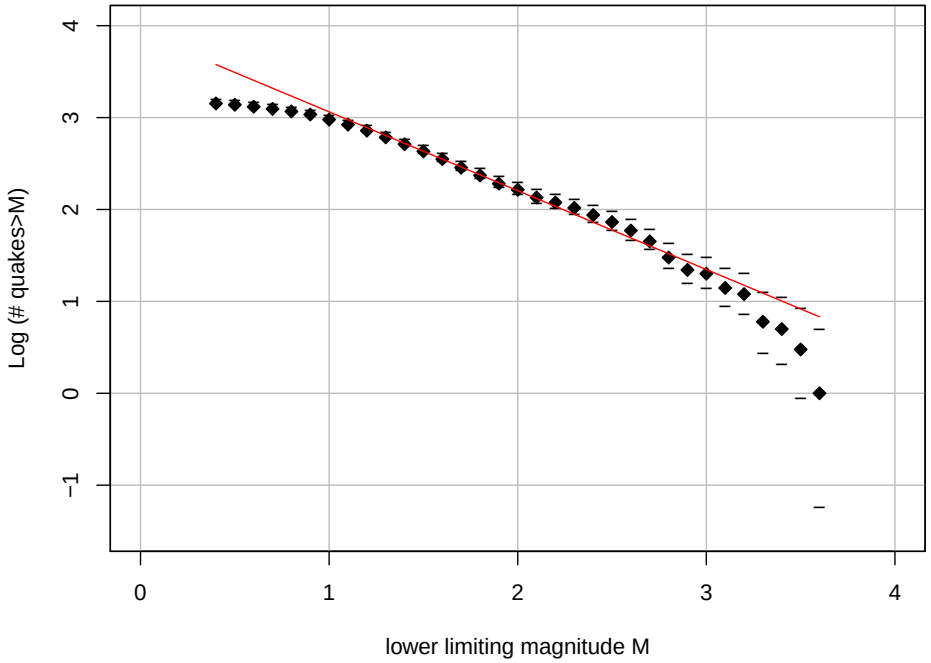

It is evident from fig. 1.1 that the distribution of events is not uniform over the area under consideration. It is also known that the distribution function of earthquakes is not uniform as a function of magnitude. For all 1066 quakes in the catalog since 1991 that occurred within zone large the distribution 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 (e.g. Garwood (1936)). 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, i.e. very close in value to the results of Dost et al. (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

Figure 2.2 The total number of earthquakes for 20 consecutive subsets of equal length time intervals, covering the period from 1 Jan. 2003 to 1 May 2018.

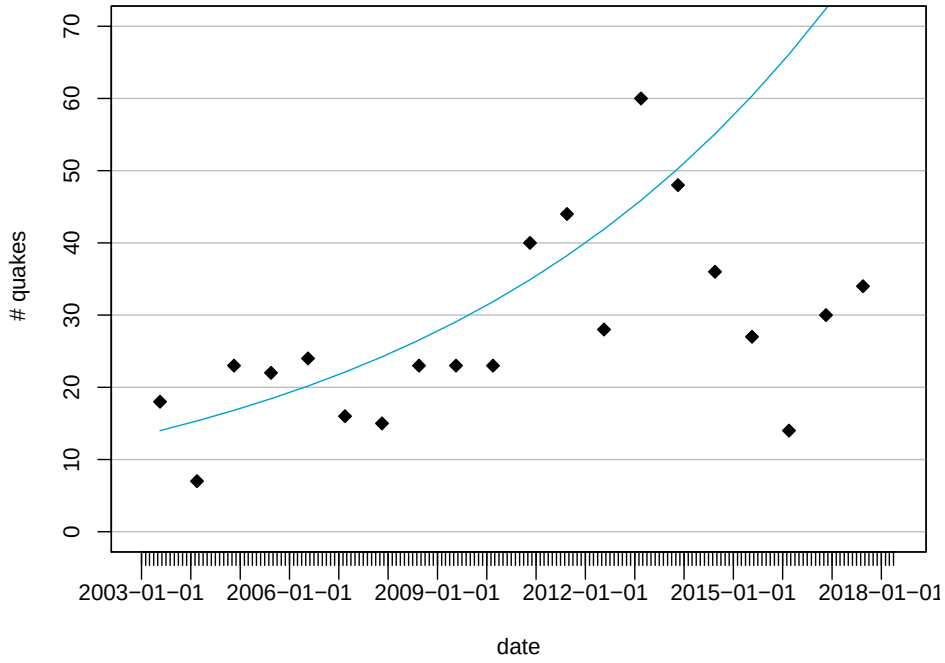

occur too far away from any of the seismometers in the network the signal becomes indistinguishable from noise or cannot be located with sufficient accuracy. For tremors with magnitudes below about 1.0 the `missing' smaller earthquakes or tremors probably do occur but the detection of such events is 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. This is true in particular for data collected before major upgrades in the detector network in the course of 2014 and 2015. That upgrade has pushed down the completeness limits to lower values

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

Fig. 2.2 shows the total number of events for 20 consecutive subsets of the data of about 9 months each. The first 19 taken together cover the same time period as used in the previous CBS report Pijpers and van Straalen (2017b). However each subinterval starts and ends at a different date since 7 months of new data have been added since the previous report. Further detail concerning the choice of length of these subintervals can be found in the supplementary paper Pijpers (2017). Also a fit to these points is shown in fig. 2.2 of the form $A\exp(t/\tau)$. Fit parameters are determined using all points except the final four, the cut-off being set at February 2014 when the first measures were taken to reduce production. The characteristic timescale τ that is determined from the fit of the function to these data, indicates that the rate of quake events doubles roughly every 5.5 years. Both a least-squares and a maximum likelihood fitting has been performed, with the same result, within the uncertainty of 0.2 years, for the doubling time. This value determined for τ has remained essentially the same over all the reports so the precise choice of where each subinterval begins and ends has no substantial influence on that value.

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 (e.g. half a year or less), 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. 2.2 but more fine-grained. Such a time series would have more sampling points than fig. 2.2 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 meaningful conclusions are difficult to obtain. For this reason from the first report (Pijpers (2014)) onwards 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 e.g. in textbooks by Robert and Casella (2004), Tarantola (2004).

Since in each simulation all the 564 earthquakes with $M \geq 1$ in 'zone large' are assigned, the same limitations apply to the simulations as apply to the real data. A close similarity in this aspect, between the synthetic 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 564 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 probability for an event to belong to a certain group. In the present case there are twelve relevant groupings, constructed by a subdivision in time and subdivisions in space :

1. A grouping in time : the event *either* occurs in the period epoch I from Jan. 1 2003 up to Jan. 1 2010, in epoch II from Jan. 1 2010 to Jan. 1 2015, *or* it occurs in epoch III from Jan. 1 2015 to May 1 2018.
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 twelve groups are obtained by events within each zone occurring in either the first, second, or third 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 probability for an event to occur within each of the three spatial groups is simply proportional to the area of each zone. Also, the probability for an event to occur in the first, second or third of the two time ranges is proportional to the length of each range. The combined probabilities are obtained assuming independence i.e. by straightfoward multiplication of the probabilities for the spatial divisions and for the division in time.

The next step is to assign each event (quake magnitude) to one of the twelve 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 564 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 5000 repetitions was considered sufficient, since there does not appear to be a need to determine the simulated number of quakes N_{sim} and the standard deviation σ to more than 3 significant digits. Using these 5000 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 5000 simulations will produce a different realisation from the stochastic assignment. Some further details concerning the use of a bootstrapping approach, rather than using a multinomial distribution for the hypothesis tests are given in Pijpers (2017).

Other probabilities 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 564 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 4 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 564 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 2003, for each group and the probabilities for assignment to each group for homogeneous and stationary test case.

region	epoch	Number of events	probability
zone SW	I	22	0.0312
	II	40	0.0223
	III	17	0.0147
zone central	I	79	0.0515
	II	81	0.0368
	III	40	0.0243
zone SE	I	12	0.0565
	II	52	0.0404
	III	25	0.0266
zone large	I	65	0.3184
(not SW, SE or central)	II	91	0.2273
	III	40	0.15

Table 2.2 Simulated number of quake events for each group for homogeneous and stationary test case, and standardised difference. 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	epoch	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	I	17.5	4.11	1.1	8	10	26	29
	II	12.5	3.50	7.9	4	6	20	22
	III	8.3	2.83	3.1	2	3	14	16
zone central	I	29.2	5.16	9.7	16	20	40	43
	II	20.8	4.49	13.4	10	12	30	34
	III	13.7	3.61	7.3	6	7	21	24
zone SE	I	32.0	5.44	−3.7	18	22	43	46
	II	22.9	4.68	6.2	12	14	32	36
	III	15.1	3.79	2.6	6	8	23	26
zone large	I	179.5	11.09	−10.3	150	157	201	207
(not SW, SE	II	128.1	10.02	−3.7	103	109	148	155
or central)	III	84.5	8.60	−5.2	63	68	102	107

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. Under this null hypothesis the probability for an event to occur within each of the three spatial groups is simply proportional to the area of each zone. Also, the probability for an event to occur in the first, second or third time ranges is proportional to the length of each range. The combined probabilities are obtained assuming independence i.e. by straightfoward multiplication of the probabilities for the spatial divisions and for the division in time.

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 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 5000 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. While the distribution of the simulated data does not conform exactly to a normal distribution, a value larger than ~ 2 for the standardised difference implies a statistically significant deviation at the 95% level at least in most cases, except for very small total counts. The standardised differences are a good enough indicator here to see directly that, apart from the result for the group `zone SW ' in epoch I, this null hypothesis is strongly rejected. The standardised differences are shown for illustrative purposes; proper limits (percentiles) for the probability distribution function as determined from the Monte Carlo simulations are used as (non-)rejection criterion. 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 probability 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 all three periods.

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	epoch	Number of events	probability
zone SW	I	22	0.0641
	II	40	0.0458
	III	17	0.0302
zone central	I	79	0.1623
	II	81	0.1159
	III	40	0.0764
zone SE	I	12	0.0722
	II	52	0.0516
	III	25	0.034
zone large	I	65	0.159
(not SW, SE or central)	II	91	0.1136
	III	40	0.0749

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 2.2

region	epoch	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	I	36.1	5.76	−2.5	22	25	48	52
	II	25.7	4.97	2.9	14	16	36	39
	III	16.9	4.01	0	8	9	25	28
zone central	I	91.7	8.64	−1.5	69	75	109	115
	II	65.3	7.62	2.1	46	51	81	86
	III	43.2	6.37	−0.5	28	31	56	60
zone SE	I	40.8	6.24	−4.6	26	29	53	58
	II	29.1	5.23	4.4	16	19	40	43
	III	19.2	4.29	1.4	9	11	28	31
zone large	I	89.8	8.66	−2.9	68	73	108	112
(not SW, SE	II	64.0	7.48	3.6	46	50	79	84
or central)	III	42.2	6.30	−0.4	27	30	55	59

Comparing table 2.3 with table 2.1, the probability for quakes to occur within zone SW is now enhanced by a factor of ~ 2.1 over the homogeneous value, and for the zone central the probability is enhanced by a factor of roughly ~ 3.1. 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.

The mean and standard deviation for the 5000 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. For both the SW and SE zones in epoch I and II this hypothesis is rejected, as well as for the zone "large". The overall number of tremors for all zones combined in epoch III (after Jan. 1 2015) of 122, is exactly the same as the number of simulated events, so the

hypothesis is not rejected for this epoch. Effectively this means that the tremor rate since Jan. 1 2015 for tremors with $M \geq 1$, is about the same as the average rate between Jan. 1 2003 and May 1 2018. For the epoch II total, with 264 recorded events, and 184 synthetic events, rejection is clear.

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

From the discussion in section 2.4 it is clear that the non-homogeneous stationary null hypothesis also does not appear very realistic. Using the number of earthquakes recorded in each of the 20 successive periods discussed in section 2.1 (fig. 2.2) , one can re-assess the probability for earthquakes to occur after Jan. 1 2015 by extending the trend over the past years.

Table 2.5 Probabilities for assignment to each group for non-homogeneous and exponentially increasing test case with $\tau = 5.5$ years. For convenience the numbers of true events are repeated.

region	epoch	Number of events	probability
zone SW	I	22	0.0351
	II	40	0.0506
	III	17	0.0543
zone central	I	79	0.0889
	II	81	0.1281
	III	40	0.1375
zone SE	I	12	0.0396
	II	52	0.057
	III	25	0.0612
zone large	I	65	0.0872
(not SW, SE or central)	II	91	0.1256
	III	40	0.1348

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 2.2

region	epoch	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	I	19.8	4.36	0.5	9	12	29	31
	II	28.4	5.19	2.2	16	18	39	42
	III	30.6	5.43	−2.5	18	21	42	46
zone central	I	50.0	6.75	4.3	34	37	64	68
	II	72.4	7.96	1.1	52	57	88	93
	III	77.5	8.13	−4.6	57	62	93	99
zone SE	I	22.3	4.61	−2.2	11	14	32	35
	II	32.1	5.53	3.6	18	22	43	47
	III	34.6	5.72	−1.7	20	24	46	49
zone large	I	49.2	6.76	2.3	32	36	62	66
(not SW, SE	II	71.0	7.99	2.5	51	56	87	92
or central)	III	76.2	8.14	−4.4	56	61	92	98

Using the fit discussed, with the doubling time $\tau = 5.4$ years, the probabilities can be re-determined for each of the 12 groups and Monte Carlo simulations produced to test whether this time dependence, together with the same enhanced probabilities in the regions of interest is consistent with the data. Comparing the probabilities for a quake to occur after Jan. 1 2015 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. From the final five columns in table 2.6 it can be seen that this null hypothesis is rejected at 99% confidence for region SW in epoch III, zone central in epochs I and III, zone SE in epoch II and the region within the largest ellipse but outside of the subregions in epoch III as well. For zone SW in epoch II, for zone SE in epoch I and for zone large in epoch I and II there is rejection at 95% confidence.

The total number of quakes in epoch II (264) is higher than even the increasing trend produces which is 203.9. On the other hand, the total number of quakes over all regions in epoch III, after Jan. 1 2015, of 122 is lower by a statistically significant amount compared to the general increasing trend for which the synthetic data produce a count of 218.9. In this sense it appears clear (as in the previous reports Pijpers (2014, 2015a,b, 2016a)) that in epoch II the earthquake rate was significantly enhanced, but this is no longer the case in epoch III. The null hypothesis of continuation after Jan. 1 2015 of the increasing earthquake rate is rejected at a confidence level of 99%.

2.6 Null-hypothesis IV: non-homogeneous process with reverted rates

Combining the insights from the hypothesis tests II and III, it appears that a better model might be obtained if one assumes that there is an increasing trend over epochs I and II with a break so that the rate reduces again in epoch III.

Table 2.7 Probabilities for assignment to each group for non-homogeneous test case with rate reverting to long term average after Jan. 1 2015. For convenience the numbers of true events are repeated.

region	epoch	Number of events	probability
zone SW	I	22	0.0351
	II	40	0.0747
	III	17	0.0302
zone central	I	79	0.0889
	II	81	0.1892
	III	40	0.0764
zone SE	I	12	0.0396
	II	52	0.0842
	III	25	0.034
zone large	I	65	0.0872
(not SW, SE or central)	II	91	0.1855
	III	40	0.0749

In the light of the time dependence shown in fig. 2.2 this must be interpreted in the sense that the (exponentially) increasing trend is not continuing. One can reasonably hypothesise that after Jan. 1 2015 the tremor rate has dropped back down to a value, roughly equal to an average value, for instance an average for the entire range from Jan. 1 2003 to May 1 2018. The resulting probabilities (table 2.7) and simulation results (table 2.8) are shown.

This hypothesis is rejected at 99% confidence only for zone central in epoch I and II. The total number of quakes after Jan. 1 2015 for all zones combined is 122 which is not statistically significantly different from the synthetic value of 121.3 Of all the hypotheses tested, this is the least unlikely one, although one perhaps could conclude that for the central zone either the characteristic time scale for increase in epochs I and II was longer than elsewhere since the

Table 2.8 Simulated number of quake events for each group for non-homogeneous and reverted test case, and standardised difference. The columns are as in table 2.2

region	epoch	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	I	19.7	4.32	0.5	10	12	29	32
	II	42.4	6.33	−0.4	27	30	55	59
	III	17.0	4.04	0	8	10	26	28
zone central	I	50.1	6.77	4.3	34	37	64	68
	II	106.9	9.25	−2.8	84	89	125	131
	III	43.0	6.32	−0.5	27	31	56	61
zone SE	I	22.3	4.60	−2.2	11	14	32	35
	II	47.6	6.53	0.7	31	35	61	65
	III	19.2	4.26	1.4	10	12	28	31
zone large	I	49.1	6.65	2.4	33	36	63	67
(not SW, SE	II	104.6	9.25	−1.5	81	87	123	130
or central)	III	42.1	6.24	−0.3	27	30	55	59

increase shown by the simulations is too high, or it could be that the rates dropped earlier there than for the other zones which would be consistent with the earlier (Jan. 2014) reduction of production levels at a central production cluster.

2.6.1 Gutenberg-Richter relationship

In addition to the number or frequency of earthquakes there is an interest in following the development of the magnitude distribution function with time. For epoch III the Gutenberg-Richter plot per region is shown in Fig. 2.3. These plots also relate to the question whether it is possible to determine whether there is a maximum quake magnitude. Such a

Figure 2.3 The Gutenberg-Richter plot separately for each region, including only earthquakes in epoch III. Top left panel: region SW, top right: central, bottom left: SE, bottom right: the rest of zone large. The black symbols denote the earthquake counts, the red line is the average from the simulations (under hypothesis IV) with the horizontal dashes indicating $\pm 1\sigma$ margins above and below.

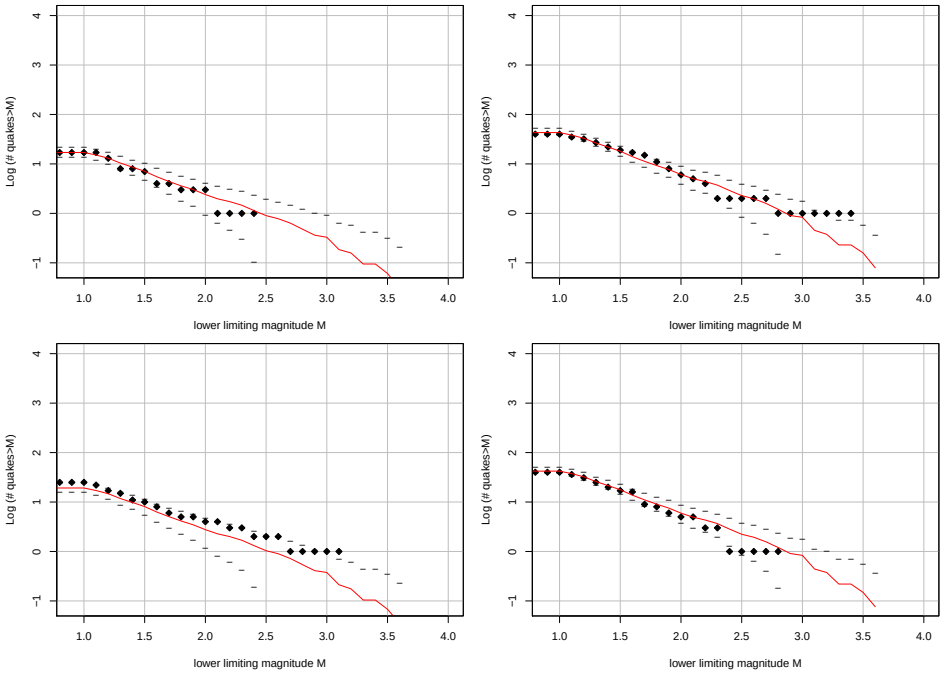

maximum could either be in the sense of a cutoff in the distribution function, or could be more operationally defined as a magnitude above which the likelihood of such an earthquake occurring in a given time frame falls below a certain threshold value.

Given the selection criterion that all earthquakes with $M < 1$ are excluded, all counts and simulations in Fig. 2.3 are by construction flat below $M = 1$. From this figure it is clear that the confidence margins towards higher M become very large, which illustrates the difficulty in establishing for instance whether there is evidence for a cutoff in magnitude. Neither Fig. 2.1 nor Fig. 2.3 provide strong support for or against a cutoff.

An operational definition for a M_{max} might for instance be that there is likelihood of 1% or less of occurrence of an induced earthquake exceeding that magnitude in the next 30 years. This type of definition corresponds to finding the magnitude at which a 99% upper confidence bound in Fig. 2.1 or Fig. 2.3 crosses the level of 1 quake (i.e. 0 on the log-scale used). Since the earthquake rate is demonstrably fluctuating over time, even such an operational definition requires making the assumption that the rates remain within the bounds that have been seen over the past fifteen or so years.

The 99% upper bounds would lie about 0.45 units above the horizontal dashes in 2.3. A period of 30 years is roughly a factor of 10 longer than the period covered by 2.3, or a factor of 2 longer than the period covered by Fig. 2.1. This factor of 10 implies increasing the upper limit another 1 unit on the logarithmic scale of Fig. 2.3. After accounting for these two effects, the dashes for the 99% confidence level all lie above 0 so the trend over magnitude needs to be extrapolated towards higher magnitudes in order to establish at what magnitude the confidence bound would cross a level of 0. The uncertainty in the resulting value appears to be too large to provide any useful value at this stage.

3 The influence of incompleteness

3.1 The exclusion of tremors with magnitudes below 1

In previous reports, cf. Pijpers (2014, 2015a,b, 2016a), the full catalog of events was used, including a range of low magnitudes where it is likely that not all events have been detected. From comparison of the shape of the cdfs before and after Jan. 1 2015, it appears that this is likely to play a role in particular at tremor magnitudes $M < 0.8$. For this reason, in all analyses for this report the simulations have been done for the magnitude range $M \geq 1$ and compared with counts of real tremors in the same range. In the previous reports, attention was paid to this limited range as well as to the full range, so it is straightforward to make comparisons of the previous results with what is reported here.

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 instance, less accurately localised tremors that occur near the borders of the various zones in fig. 1.1 might erroneously be assigned to the wrong area. This type of error will tend to reduce the inferred spatial contrasts in earthquake rate. For this reason the analysis as in hypothesis IV is repeated, excluding all tremors with magnitudes below 1.5. In total there are then 237 tremors left in the catalog since Jan. 1 2003.

Table 3.1 Probabilities for assignment to each group for hypothesis IV test case. Only tremors with magnitudes > 1.5 .

region	epoch	Number of events	probability
zone SW	I	9	0.0317
	II	14	0.0676
	III	7	0.0273
zone central	I	48	0.1037
	II	31	0.2207
	III	19	0.0891
zone SE	I	1	0.0286
	II	16	0.0608
	III	10	0.0246
zone large	I	27	0.0868
(not SW, SE or central)	II	38	0.1846
	III	17	0.0746

Table 3.2 Simulated number of quake events, with magnitudes > 1.5 , for each group for the hypothesis IV test case, with standardised difference. The columns are as in table 2.2

region	epoch	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	I	7.5	2.72	0.5	2	3	13	15
	II	15.9	3.86	−0.5	7	9	24	26
	III	6.5	2.54	0.2	1	2	12	14
zone central	I	24.5	4.80	4.9	13	15	34	37
	II	52.2	6.43	−3.3	36	40	65	69
	III	21.3	4.43	−0.5	11	13	30	34
zone SE	I	6.8	2.56	−2.3	1	2	12	14
	II	14.4	3.61	0.4	6	8	22	25
	III	5.9	2.43	1.7	1	2	11	13
zone large	I	20.5	4.34	1.5	10	13	29	32
(not SW, SE	II	43.8	6.10	−0.9	29	32	56	61
or central)	III	17.6	3.95	−0.1	8	10	26	28

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 (i.e. a lowering rather than an enhancement). Hypothesis IV, (see tables 3.1 and 3.2), has the same issue for the central zone as is noted with the limiting magnitude of 1 because the increase between epochs I and II appears to be too large for the synthetic events: the number of earthquakes actually decreases from epoch I to II in the central zone. It is also more clearly apparent with this limiting magnitude that in epoch I the central zone experienced rather more events than the scenario predicts.

The results for the equivalent of hypothesis II are shown in table 3.3. In particular for zone SE the discrepancies between synthetic and real data are larger than under hypothesis IV. For the

Table 3.3 Simulated number of quake events, with magnitudes > 1.5 , for each group for the hypothesis II test case, with standardised difference. The columns are as in table 2.2

region	epoch	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	I	13.7	3.59	−1.3	6	7	21	24
	II	9.8	3.06	1.4	3	4	16	18
	III	6.4	2.49	0.2	1	2	12	14
zone central	I	44.8	5.99	0.5	30	33	57	61
	II	32	5.24	−0.2	19	22	43	46
	III	21.1	4.32	−0.5	11	13	30	33
zone SE	I	12.3	3.41	−3.3	4	6	19	22
	II	8.8	2.94	2.4	2	4	15	17
	III	5.8	2.38	1.8	1	2	11	13
zone large	I	37.7	5.75	−1.9	24	27	49	53
(not SW, SE	II	26.8	4.9	2.3	15	18	37	40
or central)	III	17.7	4.03	−0.2	8	10	26	29

overall number of events in epoch III, the two hypotheses produce synthetic counts that are very close to the true event count. Hypothesis II produces too few synthetic counts in epoch II, whereas hypothesis IV produces too many. The differences are such that hypothesis II appears more unlikely than hypothesis IV. In Pijpers (2016a) it is already noted that for quakes with magnitudes above 1.5 a slower increase, i.e. a larger τ, appeared more appropriate, and the results presented here are consistent with that.

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 and Main (2003) and Naylor et al. (2009), or a modelling approach for aftershock generation (Kumazawa and Ogata, 2014) to simulate data. It has been proposed by Post (2017) that a Weibull distribution would be appropriate, and a number of results for a square grid of zones is shown. A Weibull fitting using the same zones as in

the present report is carried out in Pijpers (2018). 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 are a number of windowing methods for identifying aftershocks, e.g. as described in Baiesi and Paczuski (2004). Using that method, a set of events are identified as aftershocks and therefore excluded from the analysis. The numbers of events remaining after applying the exclusion criteria are shown in table 4.1.

Table 4.1 Probabilities for assignment to each group for hypothesis IV test case. Only tremors with magnitudes > 1, and excluding all events which are flagged as potential aftershocks (see text). The relevant recorded number of events are given in the third column.

region	epoch	Number of events	probability
zone SW	I	21	0.0345
	II	35	0.0735
	III	17	0.0297
zone central	I	72	0.0871
	II	74	0.1853
	III	38	0.0748
zone SE	I	12	0.0393
	II	46	0.0836
	III	25	0.0338
zone large	I	64	0.0899
(not SW, SE or central)	II	88	0.1913
	III	38	0.0773

Table 4.2 Simulated number of quake events, with magnitudes > 1 and excluding all events identified as potential aftershocks (see text), for each group for non-homogeneous and reverted rates test case and standardised difference. The columns are as in table 2.2

region	epoch	N_{sim}	σ	$\frac{(N_{true}-N_{sim})}{\sigma}$	99% l	95% l	95% u	99% u
zone SW	I	18.4	4.28	0.6	8	11	27	30
	II	38.9	5.95	−0.7	25	28	51	54
	III	15.8	3.88	0.3	7	9	24	26
zone central	I	46.2	6.51	4	30	34	59	64
	II	98.1	9.06	−2.7	75	80	116	121
	III	39.6	6.05	−0.3	25	28	52	56
zone SE	I	20.9	4.51	−2	10	12	30	33
	II	44.4	6.31	0.3	29	33	57	61
	III	17.9	4.12	1.7	8	10	26	29
zone large	I	47.7	6.53	2.5	32	35	61	65
(not SW, SE	II	101.1	9.06	−1.4	79	83	119	124
or central)	III	41.1	6.22	−0.5	26	30	54	59

The probabilities for tremor rates in accordance with hypothesis IV are shown in table 4.1, and the simulation results in table 4.2. This model is similar to the one discussed in section 2.6, but excluding potential aftershocks. The spatial enhancement factors in the zones, SW, central, and SE are 2.0, 3.1 and 1.3 respectively. The synthetic counts from this hypothesis has the same pattern of deviations as hypothesis IV without filtering potential aftershocks.

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 the various hypotheses appear to point to a genuine reduction in the rate of generation of tremors for epoch III compared to epoch II.

5 Conclusions

From the analysis presented in this report, it can be concluded that, averaged over the period from Jan. 1 2003 to May 1 2018, there is a statistically significant spatial enhancement of the earthquake rate in all three zones, SW, central and SE, where gas production takes place, compared to the surrounding region, by factors of around 2.1, 3.1 and 1.3 respectively. If only the best-localised tremors with magnitudes above 1.5 are taken into account the enhancements for SW, central and SE are 1.9, 3.7 and 0.9 respectively.

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 2003, which can be fit with an exponential increase with a doubling time of ~ 5.4 years. A fit for just the central zone has a slightly shorter time scale for this increase but experiences a drop in the rate earlier than the other zones i.e. before Jan. 2015. For a selection of earthquakes with the higher magnitude limit of 1.5, the time scale could have been somewhat longer for all regions. In the most recent 40 months, since Jan. 1 2015, the data are consistent with a reversal in this trend, reducing the tremor rates to a rate consistent with the average from Jan. 1 2003 to May 1 2018. It has to be kept in mind that the tremors in zone central used to show an exponential increase. Now the data shows that the tremor rate is back at the average since Jan. 1 2003. This is therefore a decrease in the rate of tremors in recent years and clearly not a stationary pattern over time. It is not yet clear whether the rate in the most recent 18 months (final two points in Fig. 2.2) has stabilised or rises again.

References

Baiesi, M. and M. Paczuski (2004). Scale-free networks of earthquakes and aftershocks. *Phys. Rev. E 69(6)*, 066106--066114.

Dake, L. (1978). *fundamentals of reservoir engineering*. Elsevier. (17^{th} impr. 1998, chapter 3).

Doornhof, D., T. Kristiansen, N. Nagel, P. Pattillo, and C. Sayers (2006). Compaction and Subsidence. Technical report, Schlumberger.

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

Garwood, F. (1936). Fiducial limits for the Poisson distribution. *Biometrika 28*, 437--442.

Huc, M. and I. Main (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., I. Main, and S. Touati (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.

Nepveu, M., K. van Thienen-Visser, and D. Sijacic (2016). Statistics of seismic events at the Groningen field. *Bull. Earthquake Eng 1*, 1--20.

Pijpers, F. (2014). Phase 0 report 2 : significance of trend changes in tremor rates in Groningen. Technical report, Statistics Netherlands.

Pijpers, F. (2015a). Phase 1 update may 2015 : significance of trend changes in tremor rates in Groningen. Technical report, Statistics Netherlands.

Pijpers, F. (2015b). Trend changes in tremor rates in Groningen : update november 2015. Technical report, Statistics Netherlands.

Pijpers, F. (2016a). Trend changes in tremor rates in Groningen : update may 2016. Technical report, Statistics Netherlands.

Pijpers, F. (2016b). Trend changes in tremor rates in Groningen : update nov 2016. Technical report, Statistics Netherlands.

Pijpers, F. (2017). Supplementary material to cbs reports on earthquake frequencies. Technical report, Statistics Netherlands.

Pijpers, F. (2018). Weibull fitting of earthquakes in Groningen. Technical report, Statistics Netherlands.

Pijpers, F. and V. van Straalen (2017a). Trend changes in tremor rates in Groningen : update june 2017. Technical report, Statistics Netherlands.

Pijpers, F. and V. van Straalen (2017b). Trend changes in tremor rates in Groningen : update nov 2017. Technical report, Statistics Netherlands.

Post, R. (2017). Statistical inference for induced seismicity in the Groningen gas field (MSc. thesis). Technical report, TU Eindhoven.

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

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

Colophon

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