



Statistics
Netherlands

2015 Scientific Paper

Phase 1 update May 2015 : significance of trend changes in tremor rates in Groningen

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

1	Introduction	4
2	Background	5
2.1	The earthquake data	5
2.2	Monte Carlo simulations	9
2.3	Null-hypothesis I: homogeneous and stationary process	10
2.4	Null-hypothesis II: non-homogeneous and stationary process	11
2.5	Null-hypothesis III: non-homogeneous and exponentially increasing process	12
2.6	Null-hypothesis IV: non-homogeneous and non-exponential time increasing process	13
2.7	Null-hypothesis V: non-homogeneous and exponential time increasing process except for zone central	14
3	The influence of incompleteness	15
3.1	excluding tremors with magnitudes below 1	15
3.2	excluding tremors with magnitudes below 1.5	16
4	The influence of aftershocks	16
5	Conclusions	18

Nederlands

Deze rapportage behelst een uitbreiding van onderzoek dat is uitgevoerd in het kader van fase 0 van 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, met in fase 0 in het bijzonder de aandacht gericht op de meetbaarheid van het effect dat het sterk reduceren van productie in delen van het gasveld gehad kan hebben op bodemdaling en aardbevingen in het betreffende gebied.

Uit de analyse beschreven in rapport 1, en bevestigd in een recente update gebruikmakend van een langere reeks GPS gegevens, blijkt dat ongeveer 2 maanden nadat de productie sterk was gereduceerd er een statistisch significante trendwijziging is opgetreden in de lineaire component van de bodemdaling. 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. Deze factor is vergelijkbaar met de reductie in gasproductie, gezamenlijk over de locaties waar productie sterk is gereduceerd met locaties daar direct aan aangrenzend. 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 (zie rapport 1, Pijpers 2014).

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 gerapporteerd in 2014 in de zin dat alle aardbevingen tot April 2015 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 extends earlier research, reported in December 2014, that has been carried out within phase 0 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 phase 0 the particular focus was on the measurability of the effect that the strong reduction of production in part of the field may have had on the ground subsidence and earthquakes in and around that region.

From the analysis described in report 1, and confirmed in a recent update in which a longer GPS time series is used, it can be determined that some 2 months 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 and April 4 2014, compared to the period directly previous to that date. This factor is comparable to the factor by which production has been reduced, in the combined production levels of the locations where a strong reduction was implemented together with well locations directly adjacent to this where production continued. 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.

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 report of Dember 2014 in the sense that it uses earthquake data recorded by the KNMI up to April 2015.

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 to also measure the effects on subsidence and earthquakes 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" 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

Figure 1 The locations of earthquakes as reported by the KNMI. The red dots are locations of the GPS stations, some of which are identified by name. The purple ellipse 'zone large' demarks the reference area for earthquake rates. The two smaller ellipses mark two regions of interest at which gas production takes place. The production field designated here as 'zone central', is the region where production has been reduced. The indicated scales are in km.



1991 onwards. The locations of all earthquakes in the region are shown in figure 1, together with the locations of the more recently established GPS stations. Also indicated are the boundaries of three 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 presented as well as the conclusions one can draw from this phase of the research project.

2 Background

2.1 The earthquake data

The available earthquake dataset contains in total 1131 events recorded after Jan. 1 1991 up to 22 apr. 2015. Of these, there are 763 that are located within the zone indicated as 'zone large' in figure 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. The 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 727 earthquakes available.

Figure 2 The logarithmic cumulative magnitude distribution of earthquakes for the 686 earthquakes in the set. The red line is a linear function with a slope of -1 consistent with the value reported elsewhere (Dost et al., 2012). 95% confidence intervals are indicated under the assumption that the underlying process obeys Poisson statistics.






Figure 3 The logarithmic cdf-s of earthquakes for each of the 8 consecutive subsets. Upper panel: before normalization, lower panel: after normalization.



It is evident from figure 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 727 quakes the distribution is shown in figure 2. 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 large because there are few events on which to build the statistics.

Also shown in figure 2 is a linear function with a slope of -1 in accordance with 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 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.

Figure 4 The total number of earthquakes for each of the 8 consecutive subsets and an exponential fit to these points.



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. 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 have also been tested, which confirm that this behaviour does not depend on the precise boundaries between the divisions or the number of divisions used.

From figure 3 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 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 figure 3 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 figure 3 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.

Figure 4 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 figure 4 but more fine-grained. Such a time series would have more sampling points than figure 4 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

Table 1 The measured total number of quake events since Jan. 1 1995, for each group.

region	period	Number of events
zone SW	before 23-03-2014	87
	after 23-03-2014	21
zone central	before 23-03-2014	216
	after 23-03-2014	11
zone large (not SW or central)	before 23-03-2014	342
	after 23-03-2014	50

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 assess 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) or Tarantola (2005).

Since in each simulation all the 727 earthquakes are assigned the same limitations apply to the simulations as apply to the real data, which is the 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 727 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 six 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 April 22 2015.
2. A grouping in space : the event occurs either within the contours of the area marked 'zone SW' in figure 1, or within the area marked 'zone central', or not in either of these regions, but within the area marked as 'zone large' in figure 1.

The six 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 straightforward multiplication of the likelihoods for the spatial divisions and for the division in time.

Table 2 Probabilities for assignment to each group for homogeneous and stationary test case.

region	period	probability
zone SW	before 23-03-2014	0.0650
	after 23-03-2014	0.0037
zone central	before 23-03-2014	0.1114
	after 23-03-2014	0.0063
zone large (not SW or central)	before 23-03-2014	0.7703
	after 23-03-2014	0.0433

Table 3 Simulated number of quake events for each group for homogeneous and stationary test case, and standardised difference.

region	period	N_{sim}	σ	$(N_{true} - N_{sim}) / \sigma$
zone SW	before 23-03-2014	47.2	6.61	6.0
	after 23-03-2014	2.7	1.61	11.4
zone central	before 23-03-2014	81.0	8.46	16.0
	after 23-03-2014	4.6	2.09	3.1
zone large (not SW or central)	before 23-03-2014	560.8	11.34	-19.2
	after 23-03-2014	31.5	5.56	3.3

The next step is to assign each event (quake magnitude) to one of the six 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 686 events are assigned, a cdf can be constructed for each group. This assignment process is repeated a large number of times, for the present case 1000 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 1000 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 1000 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 727 events can then be used in each case to test whether the null-hypothesis can be rejected or not. The total number of events for each group is shown in table 1. The following sections present the results for 5 separate null-hypotheses.

Note that by proceeding in this way, no assumption is made about the properties of the stochastic processes underlying the generation of earthquakes. 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 727 actual events the distribution functions for magnitudes can be simulated, and confidence limits obtained, for each group, without requiring a model for the rate at which quakes with magnitudes of any particular strength will be produced.

2.3 Null-hypothesis I: homogeneous and stationary process

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

Table 4 Probabilities for assignment to each group for non-homogeneous and stationary test case.

region	period	probability
zone SW	before 23-03-2014	0.1407
	after 23-03-2014	0.0079
zone central	before 23-03-2014	0.2956
	after 23-03-2014	0.0166
zone large (not SW or central)	before 23-03-2014	0.5105
	after 23-03-2014	0.0287

Table 5 Simulated number of quake events for each group for non-homogeneous and stationary test case, and standardised difference.

region	period	N_{sim}	σ	$(N_{true} - N_{sim}) / \sigma$
zone SW	before 23-03-2014	102.3	9.43	-1.6
	after 23-03-2014	5.7	2.38	6.4
zone central	before 23-03-2014	214.9	12.11	0.1
	after 23-03-2014	12.1	3.38	-0.3
zone large (not SW or central)	before 23-03-2014	371.1	13.43	-2.2
	after 23-03-2014	20.8	4.64	6.3

the Monte Carlo approach to test hypotheses. Also, the relevant probabilities are a useful reference to assess by how much quake rates are enhanced or lowered in the other models.

Using the probabilities shown in table 2, 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 1. From the simulations a mean value and a standard deviation can be determined, and this is shown in table 3.

The mean and standard deviation for the 1000 simulations are shown, 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, the standardised differences are nevertheless a good enough indicator to see directly that, apart from the result for the group 'zone central' after March 23, this null hypothesis is strongly rejected. Also when using the proper limits for the probability distribution function as determined from the Monte Carlo simulations, 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 two 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 4 with table 2, 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

Table 6 Probabilities for assignment to each group for non-homogeneous and exponentially increasing test case.

region	period	probability
zone SW	before 23-03-2014	0.1270
	after 23-03-2014	0.0215
zone central	before 23-03-2014	0.2670
	after 23-03-2014	0.0452
zone large (not SW or central)	before 23-03-2014	0.4611
	after 23-03-2014	0.0781

Table 7 Simulated number of quake events for each group for non-homogeneous and exponentially increasing test case, and standardised difference.

region	period	N_{sim}	σ	$(N_{true} - N_{sim}) / \sigma$
zone SW	before 23-03-2014	92.4	8.98	-0.6
	after 23-03-2014	15.6	3.85	1.4
zone central	before 23-03-2014	194.1	11.95	1.8
	after 23-03-2014	32.9	5.60	-3.9
zone large (not SW or central)	before 23-03-2014	335.2	13.39	0.5
	after 23-03-2014	56.8	7.25	-0.9

probability is enhanced by a factor of roughly ~ 2.7 . Using these probabilities, shown in table 4, 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 shown in table 1. From the simulations the mean value and the standard deviation is shown in table 5.

The mean and standard deviation for the 1000 simulations are shown, as well as the standardised difference between the measured and simulated total number of quake 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 zone SW in the period after March 23, as well as the region within zone large, outside of the zones SW and central. For both 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 figure 4, 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 6 with the probabilities of the previous section (table 4), shows that this probability is now higher by a factor of roughly 2.7.

From the final column in table 7 it can be seen that for most regions this null hypothesis cannot be rejected, with one exception. The region zone central, in the period after March 23 when the

Table 8 Probabilities for assignment to each group for non-homogeneous and non-exponentially increasing test case.

region	period	probability
zone SW	before 23-03-2014	0.1318
	after 23-03-2014	0.0168
zone central	before 23-03-2014	0.2770
	after 23-03-2014	0.0352
zone large (not SW or central)	before 23-03-2014	0.4784
	after 23-03-2014	0.0608

Table 9 Simulated number of quake events for each group for non-homogeneous and non-exponentially increasing test case, and standardised difference.

region	period	N_{sim}	σ	$(N_{true} - N_{sim}) / \sigma$
zone SW	before 23-03-2014	95.8	9.14	-1.0
	after 23-03-2014	12.2	3.43	2.6
zone central	before 23-03-2014	201.4	12.18	1.2
	after 23-03-2014	25.6	5.01	-2.9
zone large (not SW or central)	before 23-03-2014	347.8	13.47	-0.4
	after 23-03-2014	44.2	6.37	0.9

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 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 105 which is higher than the 82 that are counted in the three spatial regions combined. Since the predicted probability is in some sense an extrapolation of the fitting function shown in figure 4, 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.

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 2.5 over the stationary case, and slightly higher 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. However, there are now two regions where the difference between the true number of events and the simulated value is large enough to reject the null hypothesis. For the zone SW, in the period after March 23, the null hypothesis is rejected at a 95% confidence level. For the central zone, in that period, the null hypothesis is still rejected at a 99% confidence level.

Table 10 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.

region	period	probability
zone SW	before 23-03-2014	0.1270
	after 23-03-2014	0.0215
zone central	before 23-03-2014	0.2701
	after 23-03-2014	0.0421
zone large (not SW or central)	before 23-03-2014	0.4611
	after 23-03-2014	0.0781

Table 11 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.

region	period	N_{sim}	σ	$(N_{true} - N_{sim}) / \sigma$
zone SW	before 23-03-2014	92.4	9.02	-0.6
	after 23-03-2014	15.6	3.89	1.4
zone central	before 23-03-2014	196.3	11.93	1.6
	after 23-03-2014	30.7	5.47	-3.6
zone large (not SW or central)	before 23-03-2014	335.2	13.39	0.5
	after 23-03-2014	56.8	7.30	-0.9

This means that 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, and the zone SW appears to have experienced more.

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 that date, the probability for earthquakes to happen 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 central zone remains where the difference, between the true number of events and the simulated value, is large enough to reject the null hypothesis at the 99% 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.

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 3 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 small scales, averaging over larger regions could result in the same. The current distribution of deep wells is likely to produce modest variations in detection likelihood over the region for tremors below magnitudes of 1. Nevertheless, it is worthwhile to redo the analysis for a subset of tremors taking into account a lower limiting magnitude that excludes the range where incompleteness can influence the statistics.

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 remains 2 for the SW zone, and for the central zone becomes close to 3. 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 13 and 14. This hypothesis is rejected at the 99% and 95% confidence level respectively for the SW and central zones after March 23.

The precise date at which the subsidence rate of the central region changed, as reported by Pijpers (2014) and confirmed using more recent GPS data (Pijpers & van der Laan, 2015), is uncertain by a few weeks forward or backwards. It is therefore of interest also to assess the influence of this date on the simulations. For this reason hypothesis V is also tested using only tremors with magnitudes > 1 but with a transition date of March 9 2014, and again but with a transition date of April 6 2014. The conclusions regarding rejection of this hypothesis V with these different before/after dates are identical to those already reported. The precise date, within a few weeks of March 23 2014, therefore appears not to affect the conclusions in any significant way.

Table 12 The measured number of quake events since Jan. 1 1995 with magnitude 1.0 or higher, for each group and also with magnitudes 1.5 or higher.

region	period	Number of events	
		$M > 1$	$M > 1.5$
zone SW	before 23-03-2014	59	22
	after 23-03-2014	16	6
zone central	before 23-03-2014	181	95
	after 23-03-2014	9	3
zone large (not SW or central)	before 23-03-2014	233	89
	after 23-03-2014	31	9

Table 13 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	0.1277	0.1203
	after 23-03-2014	0.0141	0.0047
zone central	before 23-03-2014	0.3257	0.4220
	after 23-03-2014	0.0335	0.0155
zone large (not SW or central)	before 23-03-2014	0.4493	0.4211
	after 23-03-2014	0.0497	0.0164

Table 14 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.

region	period	N_{sim}	σ	$(N_{true} - N_{sim}) / \sigma$
zone SW	before 23-03-2014	67.6	7.67	-1.1
	after 23-03-2014	7.5	2.73	3.1
zone central	before 23-03-2014	172.3	10.72	0.8
	after 23-03-2014	17.7	4.21	-2.1
zone large (not SW or central)	before 23-03-2014	237.8	11.41	-0.4
	after 23-03-2014	26.2	4.98	1.0
zone SW	before 23-03-2014	26.9	4.84	-1.0
	after 23-03-2014	1.0	1.0	4.9
zone central	before 23-03-2014	94.4	7.37	0.1
	after 23-03-2014	3.5	1.83	-0.3
zone large (not SW or central)	before 23-03-2014	94.4	7.37	-0.7
	after 23-03-2014	3.7	1.96	2.7

3.2 excluding tremors with magnitudes below 1.5

While figure 2 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 well-localised tremors with magnitudes above 1.5 are taken into account, the spatial enhancements in the zones SW and central become 1.7 and 3.8 respectively. Now, the equivalent of hypothesis II can not be rejected at a 99% confidence level, but there is still rejection at 95% confidence. Hypothesis V is rejected at the 99% confidence level for zone SW after March 23 (see tables 13 and 14).

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 correlation,

both in time and in space, in the likelihood for a tremor to occur, which is in excess of what it would be if each event occurred 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. A relevant method of analysis reported in the literature are (Huc & Main, 2003) and (Naylor et al., 2009), or a modelling approach for aftershock generation (Kumazawa & Ogata, 2014) to simulate data. Progress on the analysis using these methods is to be reported at a later stage.

The model-independent approach also used in the previous report is to compare the results for the most recent period discussed in this paper (March 23 2014- Apr 22, 2015) with other periods of the same length in the same area. If one assumes that the likelihoods and properties of the aftershock generation process have not changed substantially in the most recent years, the same Monte Carlo analysis carried out on other epochs provides some measurement of the extent to which the Monte Carlo simulations capture this aspect of tremor generation. If at other epochs similar deviations from the trends are seen, it becomes more likely that the Monte Carlo simulations underestimate the spatiotemporal clustering of tremor events. If that were the case the lower rate of tremors in the central zone in the most recent period might still be a statistical fluctuation

To this end the simulations corresponding to hypothesis IV (non-homogeneous, with an appropriate total number within the time interval) are repeated for a time interval of the same length, but 1 year and 2 months earlier (to avoid overlap with the current period) and also for that length of time interval 2 and 2 months years earlier. In both cases the same parameter for the exponential time dependence is used as is used in section 2.5. These simulations show that both epochs fit and cannot be rejected. For every region, both within the time interval of interest and outside of that interval (before or after) the simulations are within 1.5σ of the true counts of events. This holds true for both the period taken 1 year and 2 months earlier and for the period two years and two months earlier

There is therefore 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 the two zones, SW and central, where gas production takes place, compared to the surrounding region, by factors of around 2.1 and 2.7 respectively. If only the well-localised tremors with magnitudes above 1.5 are taken into account the enhancements become 1.8 and 3.7 respectively. For all regions 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.

Also in the most recent 13 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 around the middle 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. Tests when moving this date forward or backward by two weeks, ie. the uncertainty allowed by the GPS subsidence measurements, show that this result does not depend on the precise date of March 23. 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 the effect of making this earthquake rate not just rise less quickly than in the surrounding region but even drop somewhat.

References

Nederlandse Aardolie Maatschappij BV, 2013, A technical addendum to the winningsplan Groningen 2013 Subsidence, Induced Earthquakes and Seismic Hazard Analysis in the Groningen Field.

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

Huc, M., Main, I.G. (2003), Anomalous stress diffusion in earthquake triggering: correlation length, time dependence, and directionality, *J. Geophys. Res.* 108, (B7), 2324, doi: 10.1029/2001JB001645

Kumazawa, T. Ogata, Y. (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, doi: 10.1029/2008JB005870

Pijpers, F.P., 2014, Phase 0 report 1: significance of trend changes in ground subsidence in Groningen, CBS.

Pijpers, F.P., 2014, Phase 0 report 2: significance of trend changes in tremor rates in Groningen, CBS.

Pijpers, F.P., van der Laan, D.J., 2015, Phase 1 update May 2015: significance of trend changes in ground subsidence in Groningen, CBS.

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

Tarantola, A., 2005, *Inverse problem Theory and Methods for Model Parameter Estimation*, ch. 2, SIAM.