Production Optimisation 2018

Per Valvatne, Leendert Geurtsen, Assaf Mar-Or, Gosia Kaleta, Jan Limbeck and Gerard Joosten

Table of Contents

1 Introduction

1.1 Study background

2014 Instemmingsbesluit

Following Winningsplan 2013, the Minister of Economic Affairs decided to restrict production from the five production clusters in the Loppersum area to 3 Bcm/year. In subsequent decisions by the Minister and upon appeal at the Raad van State, this volume was further reduced until these cluster were fully closed-in. In effect this was the first intervention and optimisation of the distribution of the gas production in relation to seismicity over the Groningen field.

Winningsplan 2016

In the first quarter of 2016, the first systematic approach to optimisation of the distribution of the production from the clusters in the Groningen field to reduce seismicity was carried out. This was reported in the Technical Addendum to the Winningsplan submitted on 1/4/2016, Reference [1]. This initial optimisation was based both on mathematical optimisation and on a practical understanding of the pressure response of the reservoir to the distribution of the production and the impact on the resulting seismicity. The impact of the alternative distribution of production on the seismicity and risk was checked by performing dedicated hazard and risk assessments.

2017 Production Optimisation study

In Article 3.2 of the 30/9/2016 Instemmingsbesluit for NAM's Winningsplan 2016, Reference [2], and of the subsequent 24/5/2017 Wijzigingsbesluit, Reference [3], NAM was tasked by the Minister of Economic Affairs to investigate whether an alternative distribution of production from the Groningen field could reduce the seismic hazard or risk. NAM embarked on a detailed optimisation study, and issued a final report to SodM on 27/11/2017, Reference [4].

> **Artikel 3**
>
> 2. De Nederlandse Aardolie Maatschappij B.V. onderzoekt of een alternatieve verdeling van de productie over alle regio's tot een lagere seismische dreiging of seismisch risico leidt en brengt daarover uiterlijk op 1 december 2017 ten genoegen van de inspecteur-generaal der mijnen, een rapport uit aan de Minister van Economische Zaken.

Fragment from the 30/9/2016 Instemmingsbesluit Article 3.2

SodM approval and request for update

In a letter to NAM on 30/11/2017, Reference [5], SodM informed NAM that the optimisation study had been executed to the satisfaction of the Inspecteur Generaal der Mijnen and could be issued to the Minister. SodM described the scientific quality of this study as "state-of-the-art". However, SodM did not yet make a judgement on the applicability of the study and referred to its earlier reservations as stated in its reaction to the "Plan van Aanpak", Reference [6], that the limitations of the underlying seismological model will affect the optimisation of production.

On 20/2/2018, SodM informed NAM by letter that the optimised production distribution could not be implemented in the field, because the underlying seismological model was deemed inadequate,

Reference [7]. SodM requested for an update of the study within 6 months, using the V5 update of the seismological model which had meanwhile become available. This V5 update presented a better areal match of the historic earthquakes (as compared to the previously used V4 seismological model).

Basispad Kabinet

On 29/3/2018 the Minister of Economic Affairs sent a letter to Parliament, Reference [8], announcing the ambition of the cabinet to reduce the production from the Groningen field as soon as possible, leading to cessation of production around 2030. The letter contained a scenario of annual production volumes for the period 2018-2031 labelled “Basispad Kabinet”.

Following these developments for the Groningen production outlook, NAM agreed with SodM to base the requested update of the production optimisation study on the new production scenario “Basispad Kabinet”.

1.2 Study objective

The Groningen field has 22 production locations (20 clusters and two satellites) by which gas can be produced from the field (Figure 1-1). All wells are typically drilled vertically, and hence extract gas from a small area directly underneath the cluster area. The only exception is the deviated well, EKL-13, that is producing gas from the reservoir at some distance from the other (vertical) wells of the Eemskanaal cluster. All these clusters can be controlled relatively independently within their respective operational constraints. This offers the opportunity to make choices in how production is spatially distributed over the field. It can be observed from Figure 1-1 that the historical earthquakes display some degree of spatial clustering as well.

In line with Article 3.2 of the Instemmingsbesluit, this study aims to establish the optimal distribution of production across the various production clusters to minimise seismic hazard or risk, for a given total field offtake profile.

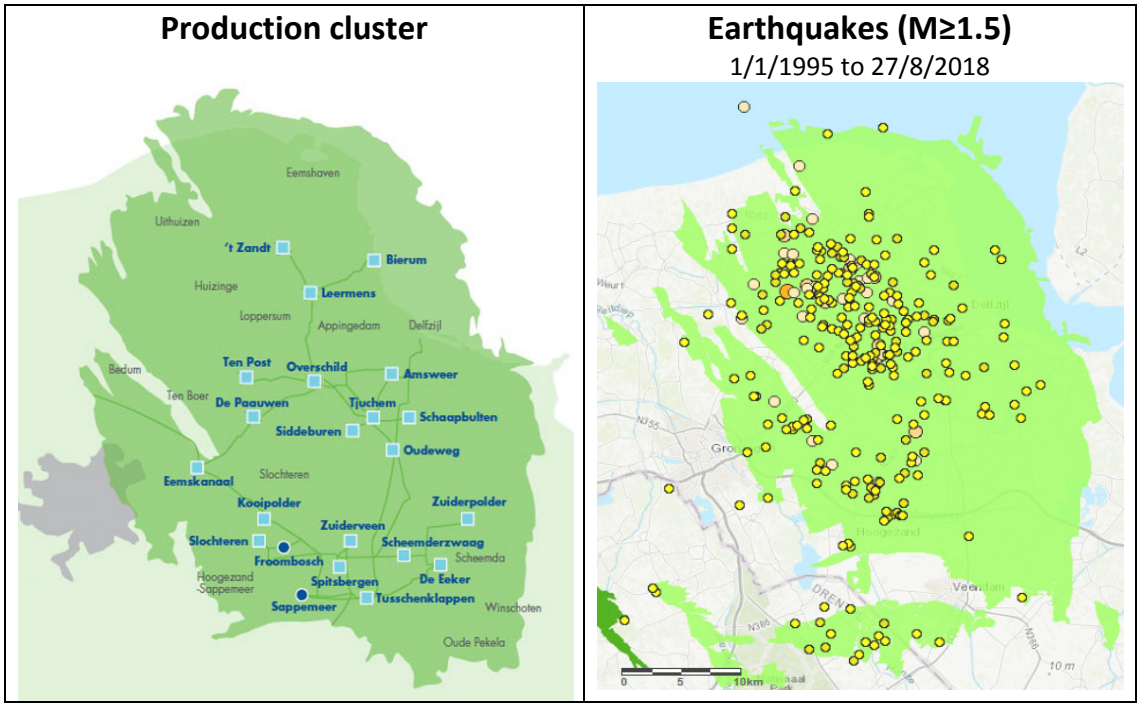

Figure 1-1 Hazard (earthquakes) and controls (production clusters)

1.3 Report structure

The structure of the study is described in chapter 2, and the hazard and risk (HRA) model in chapter 3. In chapter 4 the optimisation is discussed based on simplified assumptions for the production clusters. In this way the theoretical optima are established for different objective functions.

Chapter 5 describes the fitting of a proxy (Random Forest) model to the optimisation results, allowing for further understanding of the behaviour of the HRA model (in addition to the optimal production fractions).

Chapter 6 described how subsequently the theoretical optimum production distributions are rerun using the subsurface model coupled with a high-fidelity surface network model which reflects the operation limitations of the surface facilities, allowing for more realistic implementation of production distributions. These “operationalised” production optimisations are evaluated using the full logic-tree in order to establish the expected numbers of buildings that exceed the Meijdam norm.

2 Study setup

2.1 Production induced seismicity cause and effect chain

Production induced seismicity and its effects at surface are modelled according to the following cause and effect chain (Reference [9]):

⇨ Production from clusters

⇨ Pressure depletion

⇨ Compaction

⇨ Event Rate (Earthquakes)

⇨ Hazard

⇨ Risk

2.2 Model implementation

In the Hazard and Risk Assessments issued by NAM, the first step in the cause and effect chain up to "Pressure depletion" are calculated in the dynamic reservoir simulator (Mores, part of the Dynamo tool suite[1]), as described in Reference [10]. The subsequent calculation steps to predict earthquakes, hazard, and/or risk, are calculated in the seismic Hazard and Risk Assessment (HRA) model as described in References [11] and [12].

This model representation of the cause and effect chain can be implemented within a control loop, which allows for a model-driven optimisation. A high-level overview of the calculation steps is depicted in Figure 2-1 :

1) Optimisation algorithms, which are part of the Dynamo tool-suite, are driving the process. The total required production from the field is distributed across the production clusters (model controls) according to an initial distribution, which is issued from the optimiser to the Mores reservoir simulator.
2) In Mores, the response of the reservoir to the imposed distribution of gas offtake is calculated in terms of reservoir pressures, and provided to the HRA model.
3) Within the HRA model the rock compaction due to the pressure depletion is calculated. As a function of the compaction in time (yearly steps), spatial estimates are generated for events, hazard and risk.
4) This output is subsequently imported back into Dynamo and (a selection thereof is) used during the optimisation.

[1] Shell proprietary software

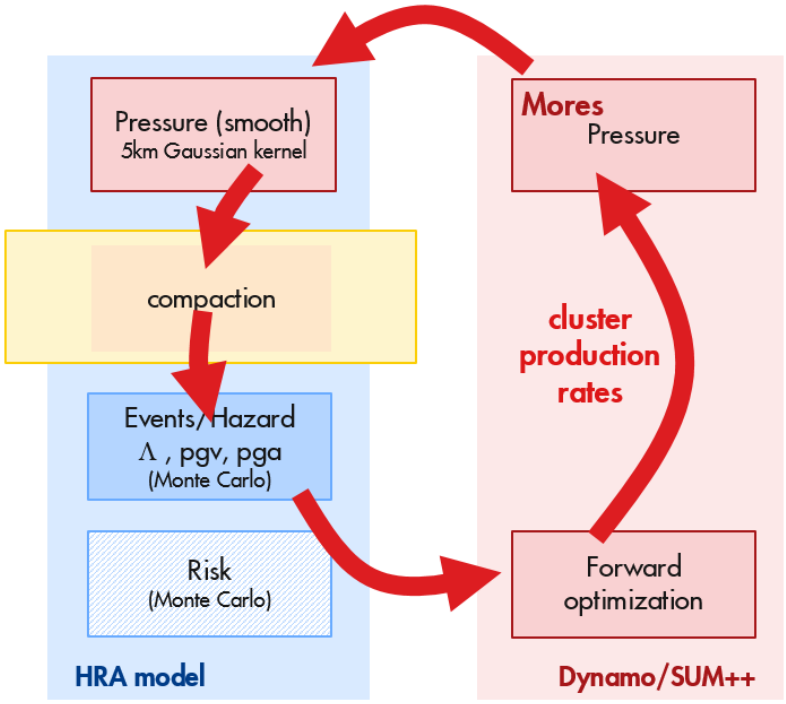

Figure 2-1 Modelling workflow

2.3 Model controls

As described in section 1.2, in principle each one of the 22 production cluster locations across the Groningen field can be controlled relatively independently within their respective operational constraints, and hence act as an independent model control for the optimisation. However, to keep the calculation times of the modelling effort within somewhat reasonable bounds, a certain degree of "lumping" was done of the individual production clusters to reduce the number of model controls. This lumping was done to ensure compatibility of this study with the June 2018 Addendum to the Hazard and Risk Assessment – November 2017, Reference [13], that was requested by the Ministry of Economic Affairs. In the request, Reference [14], NAM was asked to comply with the advice by SodM following the Zeerijp earthquake, Reference [15]. The SodM advice specified amended production regions[2] for which production fluctuations should be limited. To both accommodate this request, and the request to optimise the production distribution, NAM implemented a production cluster start-up list in its modelling effort, Table 2-1.

[2] The "East" region as stipulated in the Instemmingsbesluit, Reference [2], was broken up into 3: Bierum, East-Central (Amsweer, Tjuchem, Schaapbulten, Siddeburen, Oudeweg) and South-East (Zuiderpolder, Scheemderzwaag and De Eeker).

Start-up Priority	Region	Clusters
1	North	BIR (constant rate)
2	Eemskanaal	EKL (constant rate, winter only)
3	South-East	ZPD/SZW/EKR
4	East-Central (1)	OWG/SCB
5	South-West (1)	ZVN/SPI
6	East-Central (2)	AMR/TJM/SDB
7	South-West (2)	SAP/TUS
8	South-West (3)	KPD/SLO/FRB

Table 2-1 *Production start-up list. Starting from the top of this list, groups of clusters are sequentially opened-up by the surface network model until the total required production can be achieved. Source: Reference [13]*

These same regions are used for this report into the optimisation of production distribution. This implementation means a slight update with respect to the 2017 study (depicted in Figure 2-2):

- The former `[East]` and `[Siddeburen]` controls were redefined to improve the resolution around Appingedam.
- The former `[South-Central]` control was split into 2 regions to better honour geology.
- Production clusters 't Zandt and Leermens were previously (2017) dedicated controls, but these are now lumped into a single control `[ZND_LRM]`.

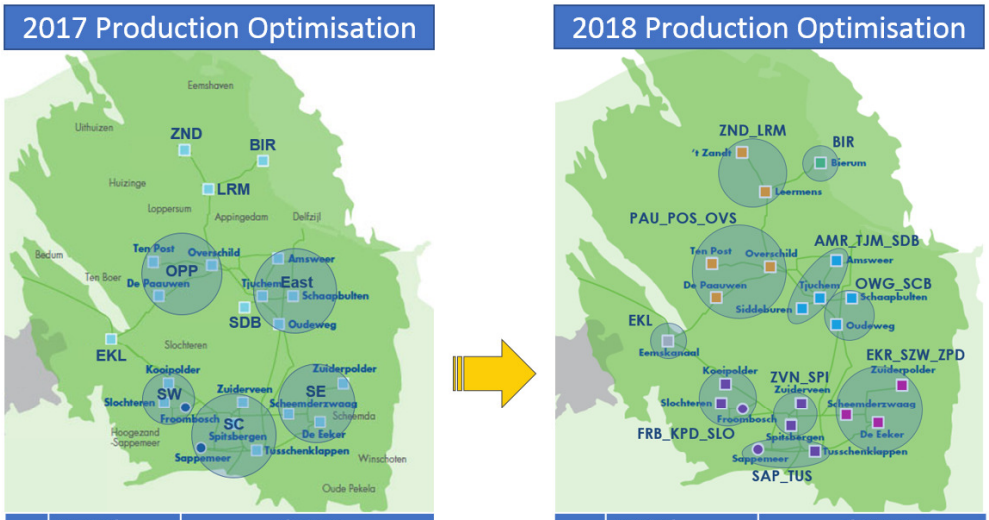

	Control	Clusters			
1	East	AMR	SCB	TJM	OWG
2	OPP	OVS	POS	PAU	
3	EKL	EKL			
4	LRM	LRM			
5	BIR	BIR			
6	ZND	ZND			
7	South-West	SLO	FRB	KPD	
8	South-East	EKR	SZW	ZPD	
9	SDB	SDB			
10	South-Central	ZVN	SPI	TUS	SAP

	Control	Clusters		
1	BIR	BIR		
2	EKL	EKL		
3	AMR_TJM_SDB	AMR	TJM	SDB
4	OWG_SCB	OWG	SCB	
5	EKR_SZW_ZPD	EKR	SZW	ZPD
6	FRB_KPD_SLO	FRB	KPD	SLO
7	SAP_TUS	SAP	TUS	
8	ZVN_SPI	ZVN	SPI	
9	ZND_LRM	ZND	LRM	
10	PAU_POS_OVS	PAU	POS	OVS

Figure 2-2 *Model controls for this study update (2018 Production Optimisation) versus the 2017 Production Optimisation. The production clusters in the right-hand panel are colour-coded in line with the Ministerial regions as defined in Reference [14].*

2.4 Objective functions

2.4.1 Practicability of a risk-based optimisation

From previous experience, it is known that a typical optimisation comprises about 250 individual simulations over 15 outer iterations. Although the dependence of risk on hazard is well understood and modelled, it cannot be distilled into a simple correlation without sacrificing accuracy and applicability. On the other hand, the risk logic-tree is considerably larger than the hazard one (as described in section 3.4.1, the risk logic-tree has 216 branches compared to 24 in the hazard one), and each risk simulation requires significantly more time to run (roughly 50% - 100% increase in run-time).

Taken together, these observations render it impossible to perform a risk-based optimisation within a reasonable time-frame while maintaining a sufficient level of accuracy. For this reason, a hazard-based optimisation approach (with a risk-assessment of the optimised distribution) was followed. The HRA model calculates several hazard-based "HRA metrics", which, separately or together, can be assigned as the objective function for the optimiser.

The objective functions used in this study are based on:

- Nuisance: Event Count
- Hazard: Maximum PGA (Return period of 475 y)[3], Maximum PGV (Return period of 475 y) [4]
- Risk: Population weighted PGV[5]

2.4.2 Event Count

The **Event Count** is a cumulative over the entire field, and, as such, represents the impact of a certain production scenario on the entire field.

2.4.3 Maximum Peak Ground Acceleration (maxPGA) and Peak Ground Velocity (maxPGV)

The Hazard objectives (**maxPGA** and **maxPGV**) are established as the highest recorded value on the hazard map associated with a certain production scenario. As such, these do not provide a response across the entire system, other than that no other place in the field is modelled to experience a hazard that is higher than this value. Focus of the optimisation based on these objective functions is the high seismic region around Loppersum.

2.4.4 Population weighted PGV (pwPGV)

The **population-weighted PGV** was chosen as a proxy for total population risk. This is based on analogy with the PAGER effort in USGS[6], which provides a rapid assessment of the impact of earthquakes (in terms of fatalities and economic loss) and estimates the population exposed to different levels of shaking.

In the PAGER method, the Modified Mercalli Intensity scale (MMI) is used to represent ground shaking combining both observations and measurements. Where measurements are available PGA and PGV are converted to MMI. The paper by Wald et al, Reference [16], shows that PGV correlates with a wider range of macro-seismic intensities, whereas PGA saturates at higher levels of MMI. Once the

[3] PGA hazard is generated for the 475-year return period for all surface locations.
[4] PGV hazard is generated for the 475-year return period for all surface locations.
[5] A weighted average PGV is calculated using the population at each location as the weight
[6] https://earthquake.usgs.gov/data/pager

population within different bands of shaking has been calculated, empirical vulnerability functions are used to estimate the losses, Reference [17]. Furthermore, the paper by Allen et al, Reference [18], states that earthquake mortality appears to be systematically linked to the population exposed to severe ground shaking (MMI VIII+, which would be PGV values above 0.30 m/s according to Wald et al. (1999)). The population weighted PGV again reflects a response of the entire system, it is calculated by multiplying each cell of the hazard map with the population in that location.

The population density is given in Figure 2-3.

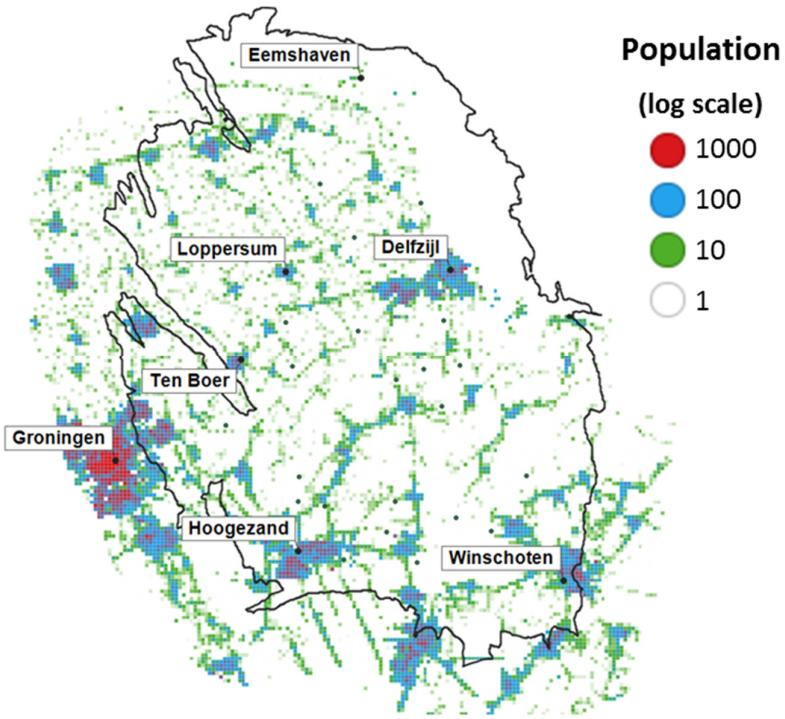

Figure 2-3 *Population map in relation to the Groningen gas field (log scale)*

2.5 Boundary conditions

The following boundary conditions were imposed for the optimisation:

2.5.1 Production profile

The total field production as used for this study was taken as per the production profile "Basispad Kabinet" for average temperature, Figure 2-4, which was received by NAM from the Ministry of Economic Affairs on 2/5/2018, Reference [14]. Note that this production profile is slightly different compared to the 29/3/2018 profile as mentioned in section 1.1, Reference [8].

The optimisation window is for this production scenario limited to 4 years, from 1/10/2018 to 1/10/2022. Hence the (relatively) low production volumes from gas-year 2022/2023 onwards are excluded from the optimisation. These years follow the onstream date of the additional nitrogen plant. For these low production-volumes, the pressure equilibration between the northern and southern areas of the field will dominate the seismic response, masking the effect of the production controls.

For the optimisation and data driven analysis, a flat production profile throughout the production year is assumed (no seasonal fluctuation). For the operationalisation and risk calculations the monthly production variations as specified in Reference [14] are included.

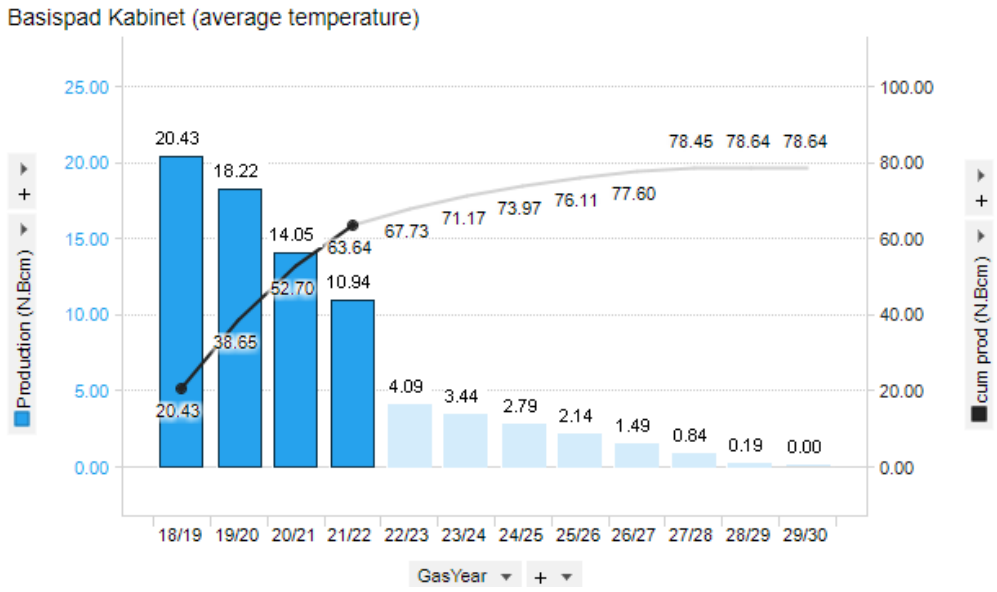

Figure 2-4 *Production profile "Basispad Kabinet" for average temperatures, as received by NAM from the Ministry of Economic Affairs on 2/5/2018. The optimisation window is highlighted (1/10/2018 to 30/9/2022).*

2.5.2 Production fractions

Within the production profile, the production fractions of the various model controls were kept constant, an example is given in Figure 2-5.

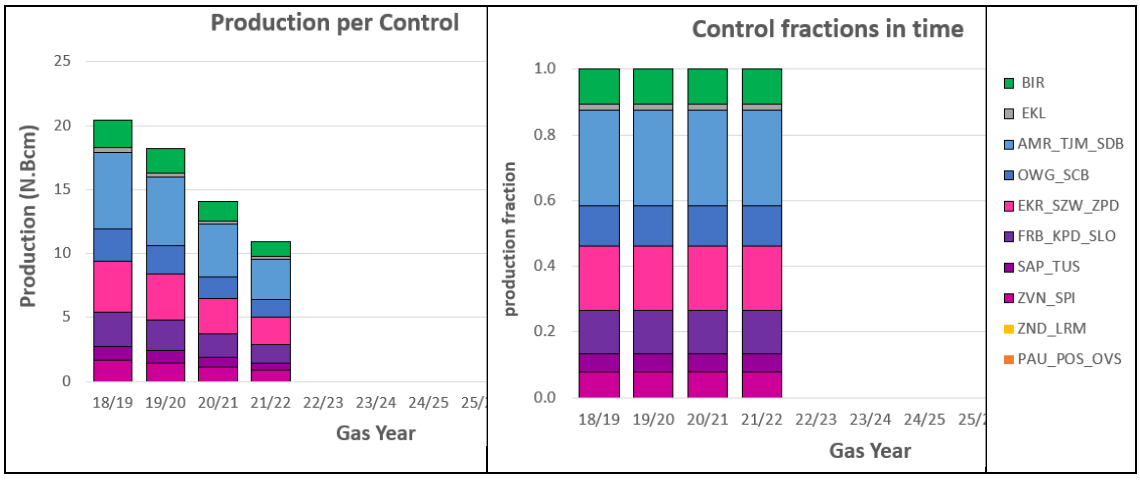

Figure 2-5 *Break-down of optimisation production profile by control*

2.5.3 Optimiser

- The optimiser is allowed full freedom in the requested production from each control, clusters are allowed to produce anything between 0 and their full capacity.
- The optimiser is allowed to vary all controls in order to minimise the objective (although the Loppersum clusters were set to zero during the optimisation runs to limit the calculation effort).

- The maximum cluster capacity is imposed by means of a simple fixed THP constraint, providing a first order reflection of the currently installed compressor capacities.
- There are no additional constraints imposed to reflect the surface facilities (e.g. pressure losses in the Groningen pipeline ring).
- A simple cluster uptime assumption of 90% to reflect planned and un-planned downtime was used.

2.6 Subset of uncertainty tree

The epistemic uncertainty in the hazard and risk model is captured in a logic tree (section 3.4). The 2017 optimisation study showed there is a highly congruous behaviour between the various branches of the Hazard uncertainty tree, Reference [4]. Given the time pressure on this study update and stretched calculation resources, it was decided to limit this study update to 5 branches out of the total 24 branches comprising the hazard logic-tree. The five branches were selected to reflect the full spread of the hazard logic-tree (Figure 2-6). The following branches have been used:

- Branch 22 = U_3_2
- Branch 24 = U_4_2
- Branch 6 = C_3_2
- Branch 9 = L_1_1
- Branch 14 = L_3_2

In chapter 5 the results for these branches are reviewed to establish the representativeness of this subset versus the potential need to evaluate additional branches.

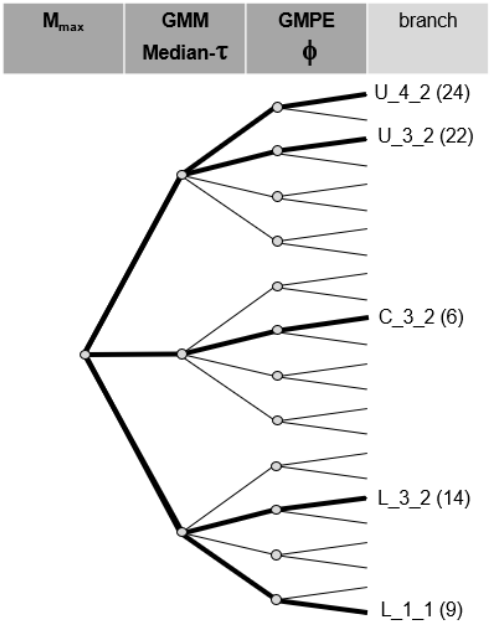

Figure 2-6 Selected subset of branches from the Hazard logic tree

3 Hazard and Risk Model

This study update utilises the V5 Hazard and Risk Model. This is the same HRA models as used in the HRA assessment of the Basispad Kabinet production scenario, Reference [13].

3.1 Dynamic reservoir simulation model

The dynamic response of the reservoir to the choices in production distribution across the various production clusters are modelled with the dynamic reservoir model. For the 2017 optimisation study the V4 Mores model was used, Reference [10], whereas this study update utilises the V5 Mores model. A full description of model V5 and the main changes with respect to V4 are given in Reference [19]. The history matched pressure in 2018 and differences between V4 and V5 models are shown in Figure 3-2. The most prominent changes for the pressure match resulted from a shift in focus towards the (high resolution) CITHP-to-CIBHP data, to better capture production induced transient effects that were induced by the Ministerial production caps.:

- The connectivity between the Ten Post, De Paauwen and Overschildt production clusters was reduced, aligned with the large fault (125 meter throw) in between these production clusters. A fault seal factor multiplier was assigned to this fault in the history matching process.
- From CITHP data it was observed that there is a 3-5 bar pressure lag within the Southwestern area of the field between West (clusters Kooipolder, Slochteren, Froombosch and Sappemeer) and East (clusters Spitsbergen and Tusschenklappen). These two areas are separated by a series of relatively large faults (>100m throw) which to the North are associated with the pop-up blocks. To reproduce this pressure lag in the reservoir model, additional fault seal factors had to be applied.
- The RFT data acquired in 1998 from the Rodewold-1 well (Southwestern periphery) showed depletion. In the V4 history match a depletion path was implemented via the graben separating the northwest from the southwest periphery. The V5 model utilises an alternative depletion path to establish a pressure match. Depletion is implemented along the horst block via Ten Boer towards Eemskanaal.

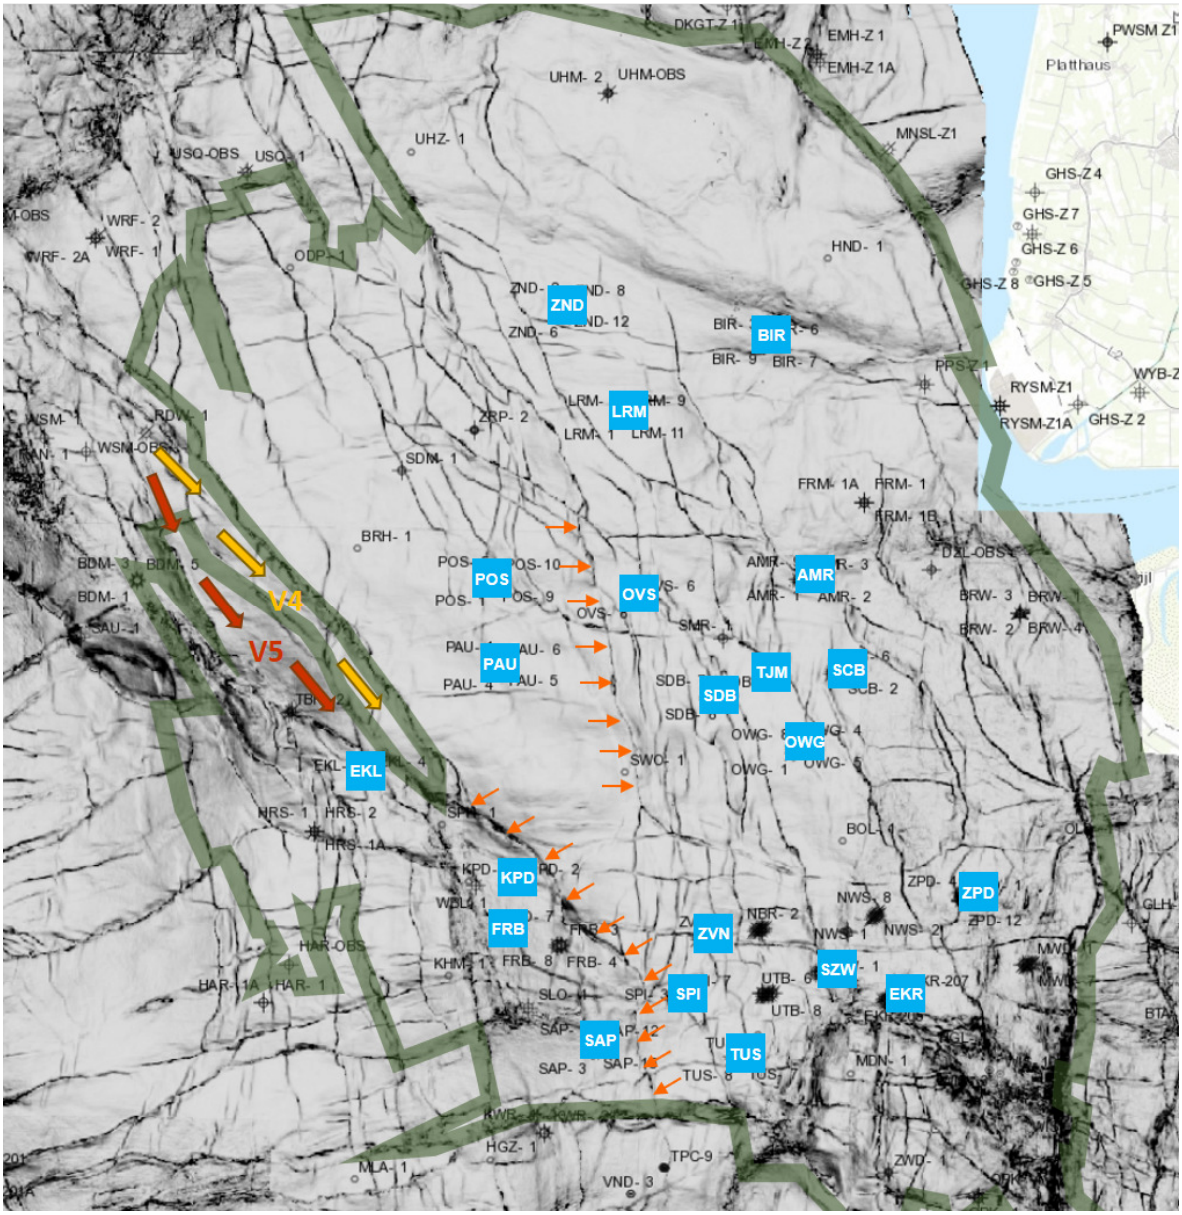

Figure 3-1 *Biggest dynamic changes between Mores model V4 and V5. Faults that were made more sealing indicated by red arrows, alternative depletion paths in southwestern periphery indicated by red and orange arrows.*

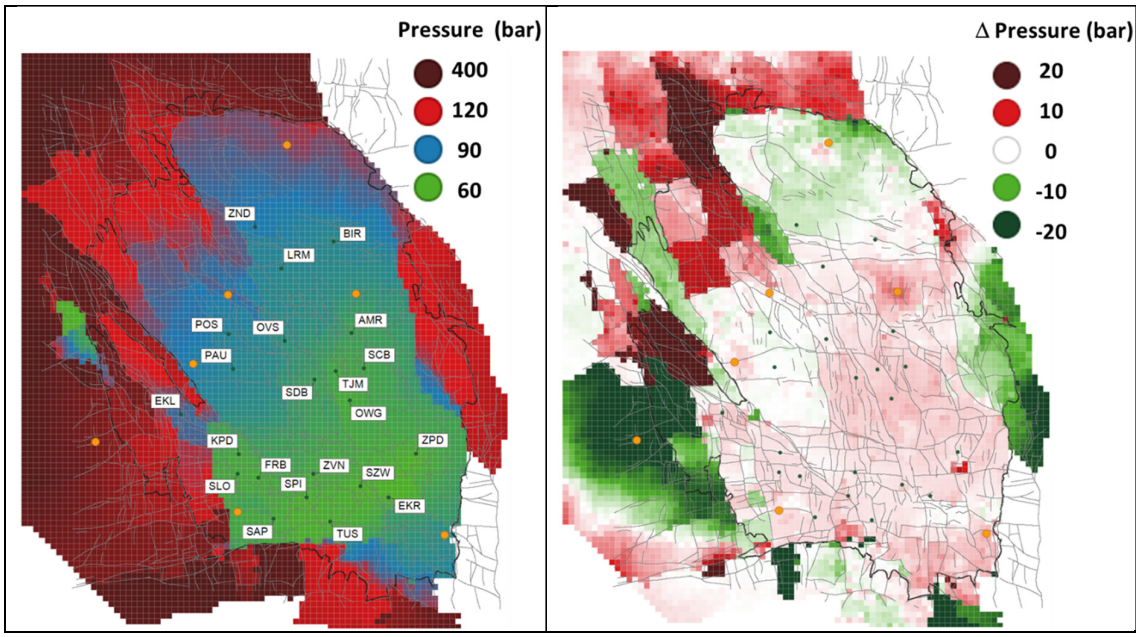

Figure 3-2 History matched pressure using V5 Mores model at 1st January 2018 and differences w.r.t. V4 model at the same time.

3.2 Seismological model

For this study, the V5 seismological model was used as described in References [11] and [12].

3.3 Ground-Motion-Prediction Equation

As part of the V4 to V5 update of the HRA model, the Ground-Motion-Prediction Equation did see some material changes. Up to Base North Sea the GMPE is more or less the same, but the amplification factors governing the propagation of the seismic energy between the Base North Sea and surface were updated.

Figure 3-4 shows the difference between the PGA maps for the Base North Sea and the surface for both the V4 and the V5 model over a period of 5 years for the 21.6 Bcm production scenario. It can be observed that the V4 model calculates a somewhat uniform increase in PGA as the seismic energy travels through the shallower subsurface, whereas the V5 model shows a dampening in the Loppersum and Southwestern area of the field and a strengthening in the East.

Although the GMPE is a highly intricate calculation and as such it is not possible to extrapolate observations from a single production scenario, Figure 3-4 does show that there are significant differences between the V4 and V5 model.

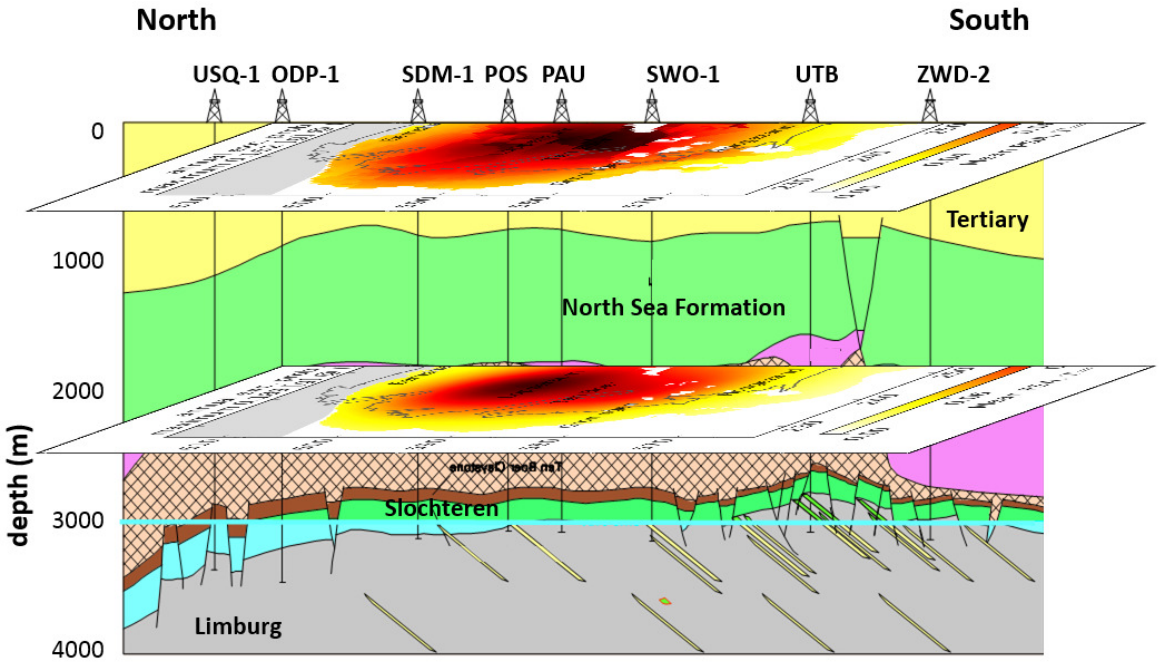

Figure 3-3 *Schematic illustration of the HRA model layers*

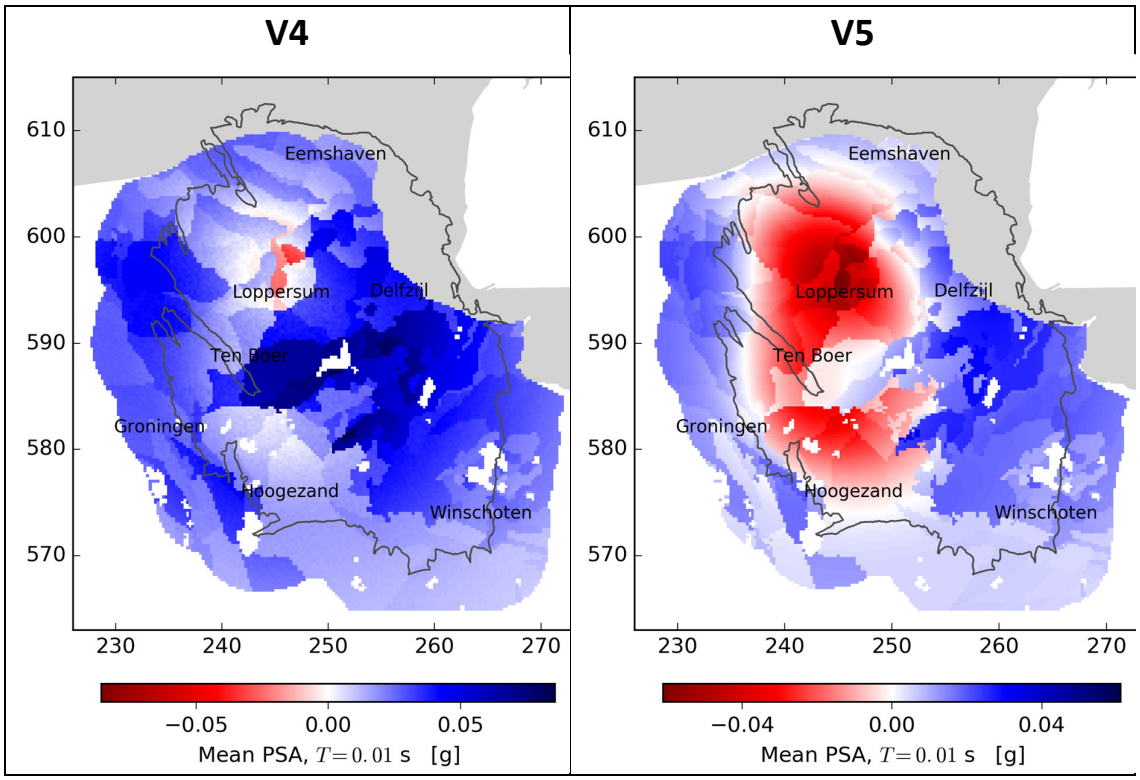

Figure 3-4 *Net difference in PGA between Base North Sea and surface for V4 and V5 HRA model, over the period 1/1/2017 to 1/1/2022 for a 21.6 N.Bcm offtake.*

3.4 Model uncertainty

There are two levels of uncertainty in the model results:

- Aleatory uncertainty
  Aleatory uncertainty is the result of the natural randomness in a process. Given the stochastic nature of the HRA simulation, how big should the reduction in the various metrics be for these to be larger than the stochastic variability of the results? For a given choice of metric and initial condition, this applies to each of the optimisation runs.

- Epistemic uncertainty
  Epistemic uncertainty is the scientific uncertainty in the model of the process. It is due to limited data and knowledge. The epistemic uncertainty is characterised by alternative models, which were captured in a logic tree. Since it is not known which of the 24 logic-tree branches reflects the field reality, how do the gains due to optimisation compare to the alternative HRA model's represented by different branches of the logic tree?

3.4.1 Epistemic uncertainty

In the HRA model, it is assumed that the uncertainty in the reservoir pressure is smaller than the uncertainty in other parameters and hence it was not explicitly included as an uncertainty in the logic tree. To reflect that pressure cannot be measured outside of well control, a spatial smoothing step is applied.

Within the HRA model the epistemic uncertainty is captured by means of an uncertainty tree. The V5 vintage of the HRA model involves a logic-tree for Hazard with 24 branches, and for Risk of 216 branches.

The logic tree for the hazard assessment comprises three sets of branches to capture the uncertainty in the most uncertain elements, Figure 3-5. The first branches cover the uncertainty with respect to M_{max} in the seismological model, and the second and third sets capture the uncertainty related to the GMPE (tau and sigma). Each branch of the logic tree represents a scenario, and by combining all scenarios using the weights in the logic tree, the mean hazard map can be calculated.

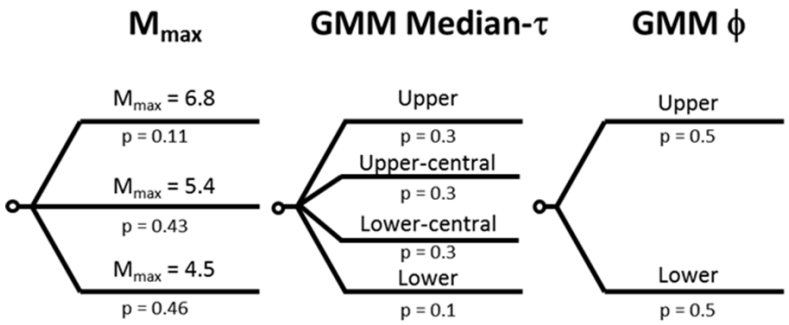

Figure 3-5 Hazard logic tree.

3.4.2 Aleatoric uncertainty

Especially the Monte Carlo setup of the HRA model introduces an aleatoric model variability (also referred to as statistical or stochastic uncertainty).

The approach taken to address the stochastic variability works as follows. For each branch of the logic-tree, the HRA model was run multiple times, using fixed control settings, a fixed grid size, and a fixed number of Monte Carlo draws ('catalogue size'). Each time the seed value for the random-number generator was varied. The resulting set of values was then used to validate stability of the estimated hazard and risk metrics. An example is given in Table 3-1. The aleatoric uncertainty within the HRA model was found to be typically around 1-2%.

	pwPGV	maxPGA	maxPGV	Events	pwPGV005
mean	0.049875	0.181582	0.161139	62.97811	0.080136
var	2.71E-07	3.13E-06	4E-06	0.001903	5.92E-07
std	0.00052	0.00177	0.002001	0.043622	0.00077
std/mean	1.04%	0.97%	1.24%	0.07%	0.96%

Table 3-1 Aleatoric uncertainty in the HRA model as tested for branch 6 using 10 different random seeds

3.5 Model robustness

In line with section 3.4 the model robustness was assessed.

3.5.1 Robustness with respect to aleatoric uncertainty

Each run of the HRA model is performed on an observation grid (with user-supplied spacing) and involves the sampling of probability distributions for the various HRA model parameters. Using finer grids and increased sampling will result in better estimation of the outcome (i.e. smaller confidence intervals), but also consumes more computational resources and consequently slows down the optimisation process. The execution time is roughly linear in the number of draws and quadratic in terms of the grid-size. It is, therefore, necessary to strike a balance between accuracy and overall computational time.

To this end, a number of production scenarios were tested:

- Production fractions based on 2017 actual data
- Production fraction as used in the HRA calculation for Basispad Kabinet (GY 18/19)
- Minimisation of number of Events (earthquakes)
- Minimisation of maximum Peak Ground Acceleration
- Minimisation of population weighted Peak Ground Velocity

Each of these production scenarios were simulated for varying resolutions of the sampling grid:

- 250m
- 500m
- 1000m
- 1500m

And for a varying number of Monte Carlo draws (catalogue size):

- 50,000
- 75,000
- 100,000
- 150,000
- 200,000

The associated model results are summarised in Figure 3-6.

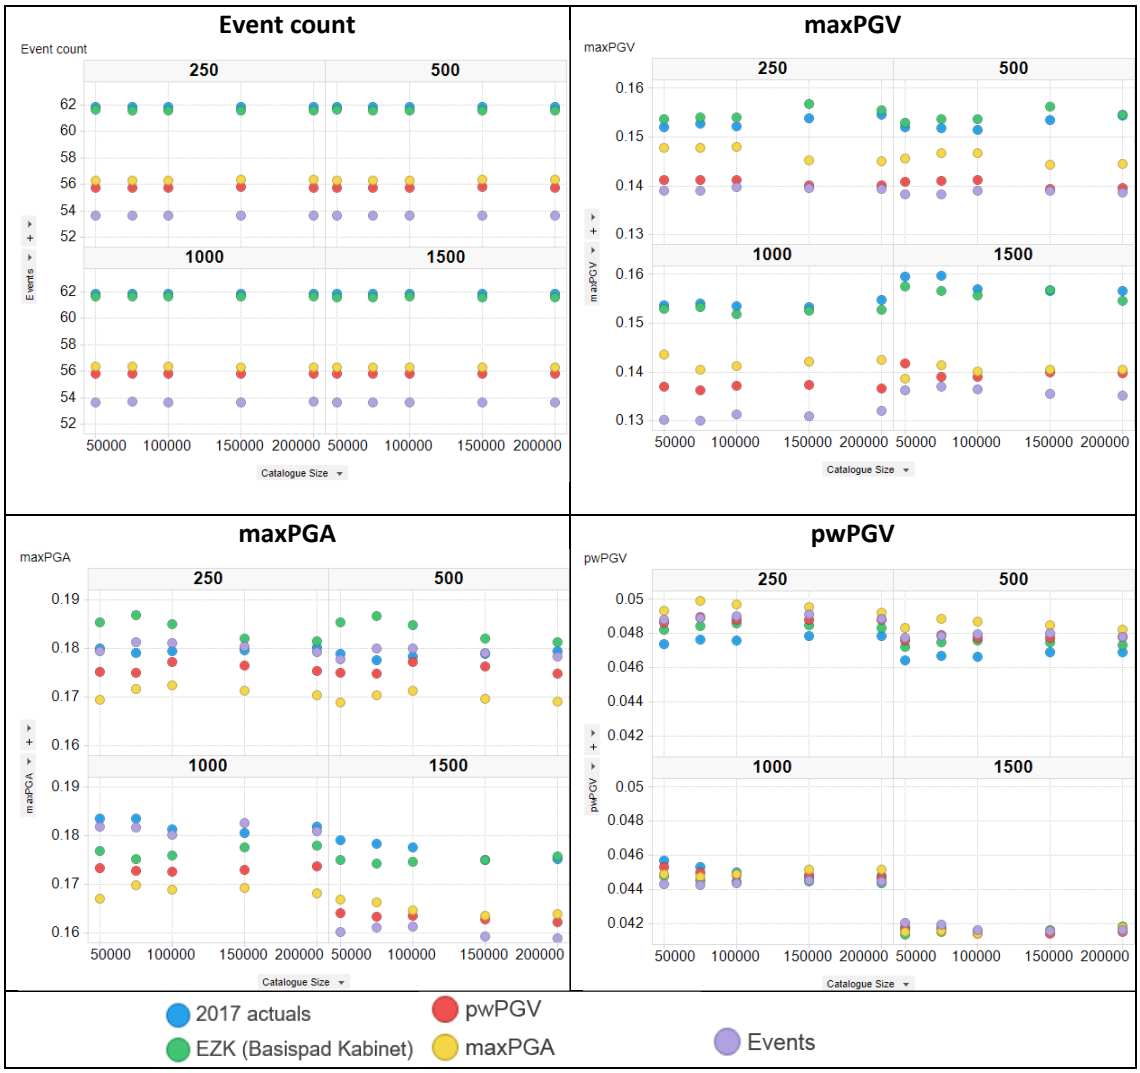

Figure 3-6 Robustness assessment for varying grid resolution and catalogue size

The **Event count** is hardly affected by the choices in grid resolution and catalogue size.

The **pwPGV** metric is obviously affected by the gridsize (number of people per gridblock varies), but between the various gridsizes there are no differences in the order of the scenarios.

From Figure 3-6, it can be concluded that there are no big differences between 250 and 500m resolution, but at 1000 and 1500m differences start showing with respect to the 250m reference. Consequently, the 500m grid was selected. At this grid resolution, a catalogue size of 75,000 yields

stable results. An overview of the numerical performance is given in Table 3-2. For the selected numerical settings the associated runtime for the HRA model is approximately 1:45 hr[7].

		Resolution (grid size)			
		1500m	1000m	500m	250m
Catalogue size	50,000	0:10 hr		1 hr	4 hrs
	75,000		0:30 hr	1:45 hr	
	100,000		0:45 hr	2:15 hr	
	150,000			4 hr	
	200,000	0:45 hr			17 hr

Table 3-2 *Model calculation times as a function of grid resolution and catalogue size*

3.5.2 Robustness with respect to epistemic uncertainty

Qualitative assessment of the robustness with respect to epistemic uncertainty is somewhat more complicated and was done by running the HRA simulation for the full logic-tree and comparing it to a reference case. A somewhat more rigorous approach to the assurance of robustness with respect to the epistemic uncertainty is to calculate, for each of the optimisation results, the expected value of optimisation metric. To ensure that the results are optimised with respect to the entire uncertainty range as captured in the logic tree, a dedicated optimisation is run for each of the uncertainty-tree branches, using the following process (Figure 3-8):

- First the objective function and initial production distribution is selected
- The following steps are executed in parallel for each branch of the logic tree
  - Given the selected set of objective function and initial rate, a dedicated SPMI optimisation is run,
  - An optimum production distribution is obtained,
  - The optimisations' end results are each evaluated on the entire logic tree (24 branches) and the value of the logic-tree mean is calculated for each metric using the weights of the various branches

After running the full uncertainty-tree on the optimal distributions obtained for each branch and weighting the resulting metric values by the probabilities associated with each branch, the optimal distribution (lowest expected value for the optimisation metric) is selected.

This process setup ensures a consistent hazard assessment (i.e. taking into account the epistemic uncertainties) and enables a like-for-like comparison of the optimisation outcomes' performance for each branch. The comparison is done for both PGA and PGV differences. Figure 3-7 shows an example of the areal PGA, along with uniform hazard spectra at specific locations (e.g. Groningen, Loppersum, Ten Boer, etc.), where two scenarios are compared. Although the epistemic confidence intervals do still overlap, there is a significant reduction of hazard in Delfzil and Loppersum.

[7] Note that solely for the optimisation purpose these settings are excessively strict. The optimisation would only require the model response to production changes to be stable within limits. So the model used for optimisation does not need to delivery accurate hazard and risk metrics only changes to these metrics with the correct sign under production changes. Having these tight settings does however allow for a more quantitative comparison with previous HRA analyses.

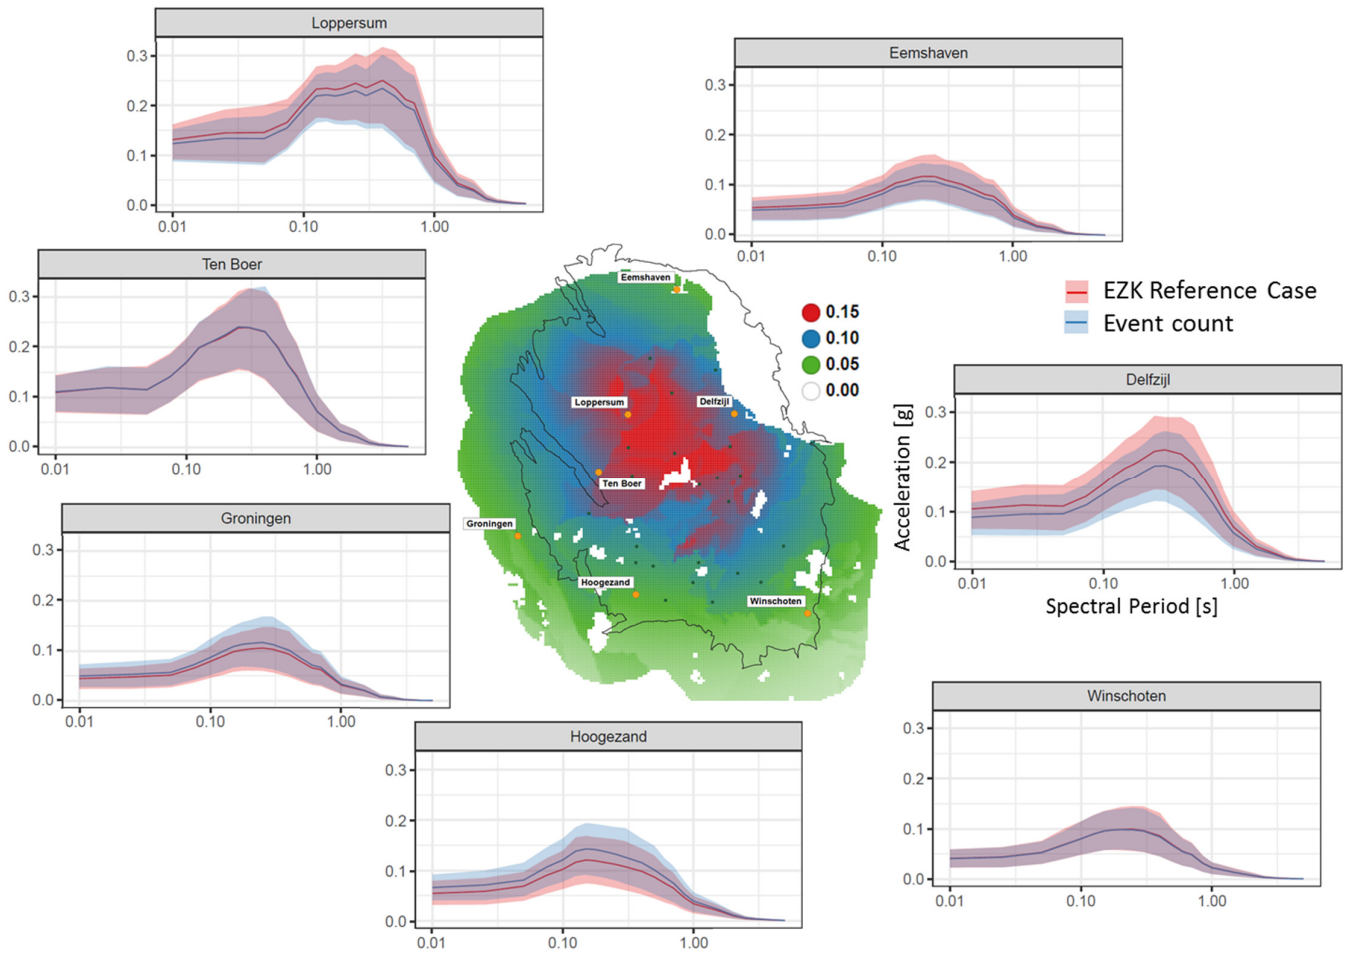

Figure 3-7 *Sample HRA output visualisation showing the PGA map for the EZK reference case (see section 6.3.2) along with uniform hazard spectra for the reference case with a comparison to an optimised case (event count) for selected locations. The reference and optimised scenarios are plotted in red and blue respectively, with shaded area representing the epistemic uncertainty The PGA map shows the acceleration at 0.01s.*

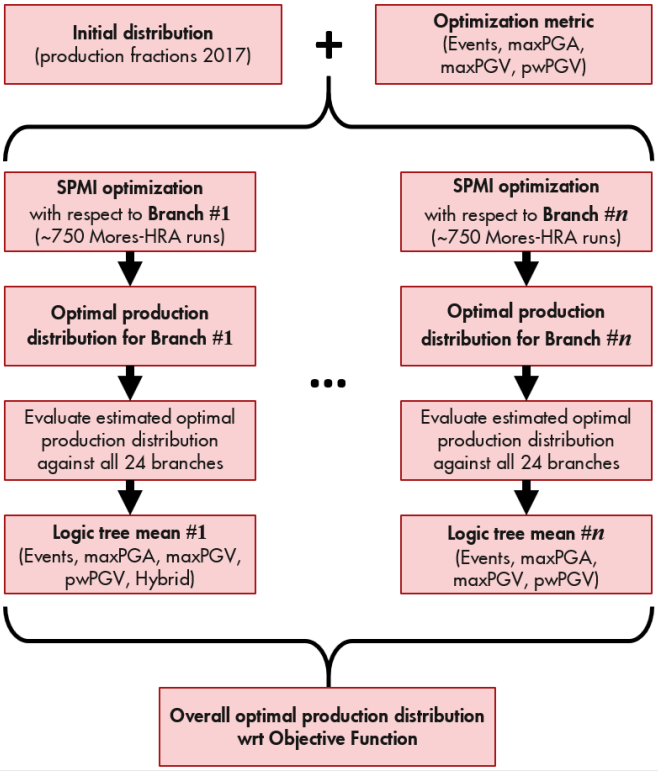

Figure 3-8 *Optimisation process-flow for an optimisation metric/initial distribution pair*

4 Optimisation

4.1 Setup

4.1.1 Controls

To reduce the calculation time, only 8 out of the 10 controls (section 2.1) were used in the optimisation runs. The two controls within the Loppersum region were excluded, and their production fractions were set to zero in all runs:

- [ZND_LRM] = 0
- [PAU_POS_OVS] = 0

The validity of this assumption is cross-checked and validated in chapter 5.

4.1.2 Reference Case

The actual production split as of calendar year 2017 was used as the reference case and as the initial starting point for the optimisation.

	Control	Production fraction
1	[BIR]	0.11
2	[EKL]	0.02
3	[AMR_TJM_SDB]	0.29
4	[OWG_SCB]	0.12
5	[EKR_SZW_ZPD]	0.20
6	[FRB_KPD_SLO]	0.13
7	[SAP_TUS]	0.05
8	[ZVN_SPI]	0.08
	Sub total	1.00
9	[ZND_LRM]	0.00
10	[PAU_POS_OVS]	0.00

Table 4-1 Production and production fractions of the reference case (calendar year 2017) by model control

4.1.3 Optimisation window

A four-year optimisation window was used, to avoid “noise” from pressure equilibration effects once field offtake drops significantly (after nitrogen plant comes on-stream).

4.2 Results

In Table 4-2 to Table 4-5 the optimised production fractions are given for the various optimisation objectives. Per objective function, a dedicated optimisation was executed on each of the selected branches of the uncertainty tree. Each set of optimised fractions were then evaluated across the entire logic tree (24 branches) to yield the logic tree mean. The branch that yielded the best result for the logic tree mean for a given objective function is highlighted in bold. For reference, the initial rates for the optimisation are included, named “branch 0”. These are the actual production fractions as per calendar year 2017.

Events

Branch	Itr	Case	BIR	EKL	AMR_TJM_SDB	OWG_SCB	EKR_SZW_ZPD	FRB_KPD_SLO	SAP_TUS	ZVN_SPI
0	2017 actual		0.11	0.02	0.29	0.12	0.19	0.13	0.05	0.08
6	**8**	**7**	**-**	**-**	**-**	**-**	**0.32**	**0.27**	**0.12**	**0.29**
9	10	5	-	-	-	0.06	0.32	0.19	0.15	0.29
14	10	5	-	-	-	0.06	0.32	0.19	0.15	0.29
22	7	18	-	-	-	-	0.30	0.28	0.13	0.30
24	7	18	-	-	-	-	0.30	0.28	0.13	0.30

Table 4-2 Production fraction for each control across the various branches when optimised for Events

maxPGA

Branch	Itr	Case	BIR	EKL	AMR_TJM_SDB	OWG_SCB	EKR_SZW_ZPD	FRB_KPD_SLO	SAP_TUS	ZVN_SPI
0	2017 actual		0.11	0.02	0.29	0.12	0.19	0.13	0.05	0.08
6	**9**	**5**	**-**	**0.22**	**-**	**0.14**	**0.27**	**0.19**	**0.08**	**0.10**
9	8	5	-	0.01	-	0.23	0.34	0.29	0.13	-
14	5	11	-	0.01	-	0.24	0.36	0.27	0.12	-
22	7	12	0.01	0.17	-	0.14	0.35	0.23	-	0.11
24	7	1	0.04	0.22	-	0.15	0.22	0.16	0.09	0.12

Table 4-3 Production fraction for each control across the various branches when optimised for maxPGA.

maxPGV

Branch	Itr	Case	BIR	EKL	AMR_TJM_SDB	OWG_SCB	EKR_SZW_ZPD	FRB_KPD_SLO	SAP_TUS	ZVN_SPI
0	2017 actual		0.11	0.02	0.29	0.12	0.19	0.13	0.05	0.08
6	7	15	-	0.03	-	0.16	0.24	0.23	0.09	0.25
9	**5**	**16**	**-**	**-**	**-**	**0.16**	**0.25**	**0.21**	**0.11**	**0.27**
14	6	5	0.01	0.08	-	0.15	0.25	0.19	0.11	0.20
22	8	12	-	-	0.01	0.16	0.23	0.22	0.16	0.23
24	7	8	-	0.01	-	0.22	0.23	0.23	0.09	0.22

Table 4-4 Production fraction for each control across the various branches when optimised for maxPGV

pwPGV

Branch	Itr	Case	BIR	EKL	AMR_TJM_SDB	OWG_SCB	EKR_SZW_ZPD	FRB_KPD_SLO	SAP_TUS	ZVN_SPI
0	2017 actual		0.11	0.02	0.29	0.12	0.19	0.13	0.05	0.08
6	2	14	0.06	0.01	0.18	0.07	0.25	0.08	0.03	0.31
9	10	10	0.19	-	0.11	0.23	0.36	-	0.11	-
14	**8**	**17**	**0.19**	**-**	**-**	**0.29**	**0.39**	**-**	**0.13**	**-**
22	2	7	0.07	0.01	0.20	0.24	0.29	0.09	0.04	0.06
24	10	5	-	-	-	0.31	0.32	0.13	0.13	0.11

Table 4-5 Production fraction for each control across the various branches when optimised for pwPGV.

The results for the various optimisations are summarised in Figure 4-1, where the optimum production fraction is given for each control. It can be observed that only the **pwPGV** optimisation is utilizing the `[BIR]` control, and only the **maxPGA** optimisation is utilizing the `[EKL]` control. None of the optimisations are using the `[AMR_TJM_SDB]` control. The remaining controls are utilised to varying degrees by the different optimisations. **Event count** optimisation is exclusively producing from the South, while **pwPGV** optimisation avoids the more populated South-West by not utilizing the `[FRB_KPD_SLO]` and `[ZVN_SPI]` controls.

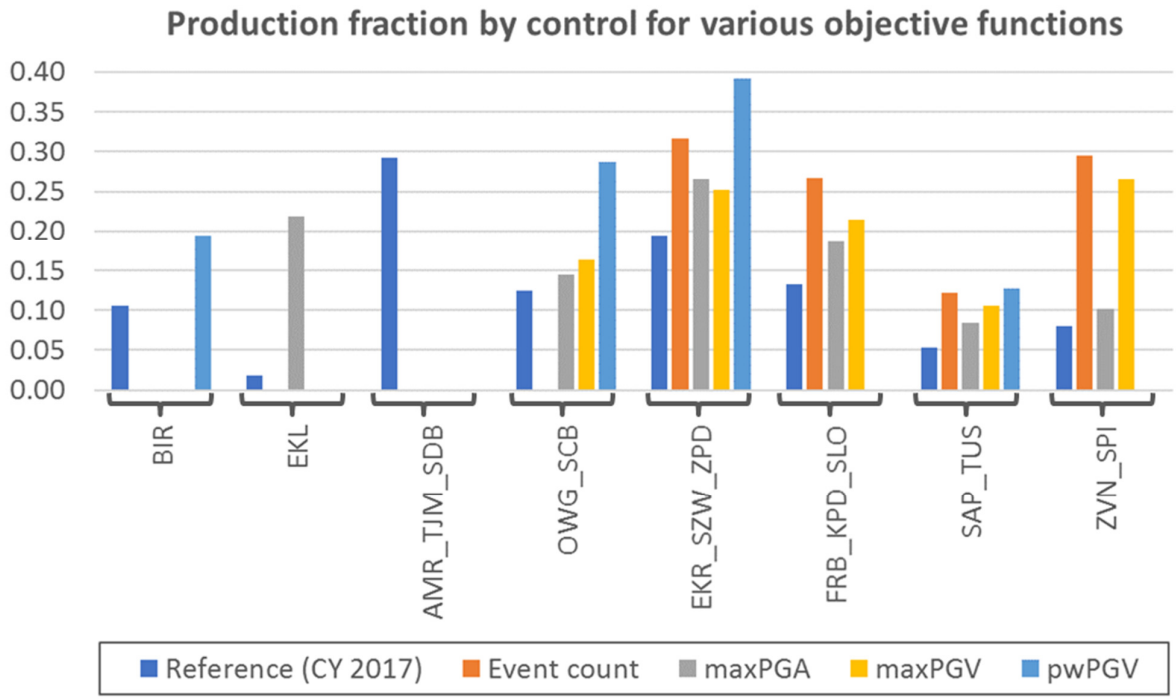

Figure 4-1 *Optimised production fractions of the controls for the different optimisation objective functions*

The improvements in hazard metrics are given in Table 4-6. These results are derived from the logic tree mean, i.e. the production fractions in Figure 4-1 evaluated across all 24 logic tree branches. The improvements vary between 5 and 14 percent depending on the metric that is optimised based on the objective function.

Objective	Hazard Metric				Improvement			
	Events	maxPGA	maxPGV	pwPGV	Events	maxPGA	maxPGV	pwPGV
Reference (2017 Act)	62.4	0.155	0.127	0.039				
Event Count	58.7	0.149	0.123	0.037	-6%	-4%	-3%	-5%
maxPGA	56.3	0.140	0.112	0.040	-10%	-10%	-12%	2%
maxPGV	56.0	0.150	0.112	0.039	-10%	-3%	-12%	-1%
pwPGV	53.6	0.152	0.110	0.039	-14%	-2%	-13%	0%

Table 4-6 *Optimisation results with respect to the HRA metrics*

In Figure 4-2 to Figure 4-6 the optimisation results are given in map view. In Figure 4-2 the reference case is displayed. The bottom left-hand panel gives the starting point for the optimisation: the reservoir pressure at 1/10/2018. The production fractions for the reference case (Table 4-1) are displayed as a bubble map in the top centre panel. The resulting depletion over the four-year optimisation period (1/10/2018 to 1/10/2022) is given in the top left-hand panel, and the associated event density in the bottom centre panel (in total, there are some 46 events expected in the North of the field, and 16 in the South). The subsequent PGA and PGV maps are in the left-hand side of the panel. For reference, the fault pattern is displayed in the background. Additionally, the main cities are highlighted as orange dots (labelled in the top right-hand panel), and the production clusters as black dots (labelled in the bottom right-hand panel).

This same visualisation panel was used to present the optimisation results for the various objective functions, but now the differences are displayed with respect to the reference case. Table 4-3 gives the results for the **Event count** optimisation. Note that the production fraction bubble map highlights the optimised production fractions in red, with the reference case production fractions kept in blue. All production is moved towards the southernmost 4 controls. As a result, the depletion has increased in the South (compared to the reference case) and reduced in the North. In terms of absolute depletion (top left-hand panel), there is even some pressure rebound just south of the city of Delfzijl. In total a reduction of some 9 events is expected as compared to the 62 events in the reference case. This constitutes of a reduction of 16 in the North, and an increase of 7 in the South. The net effect in terms of PGA and PGV gives an almost horizontal divide between increase and reduction with respect to the reference case (right-hand panel).

The optimisation with respect to **maxPGA** utilises the `[EKL]` and `[OWG_SCB]` controls, resulting in increased hazard in the South-West, rotating the net hazard effect by some 45^{o}.

The **maxPGV** optimisation does not utilise the `[EKL]` control, with more offtake from the central controls in the South, `[ZNV_SPI]` and `[SAP_TUS]`. This makes the net gain with respect to the reference case similar to the Event count optimisation, although less pronounced both in terms of event count and hazard. The net reduction in Event count has reduced to 6 (North: -12, South: + 6).

The **population weighted PGV** optimisation yields the biggest areal reduction in PGA and PGV including the entire South-West, although the magnitude of reduction is somewhat modest compared to the other optimisation strategies. This is achieved by opening up the Northern-most `[BIR]` control and not utilizing controls in the South-West that are close to the city of Groningen. The net reduction in events is also less than in the other strategies (+2 in the South and -6 in the North).

Optimization results – Reference Case (actual production fractions CY17)

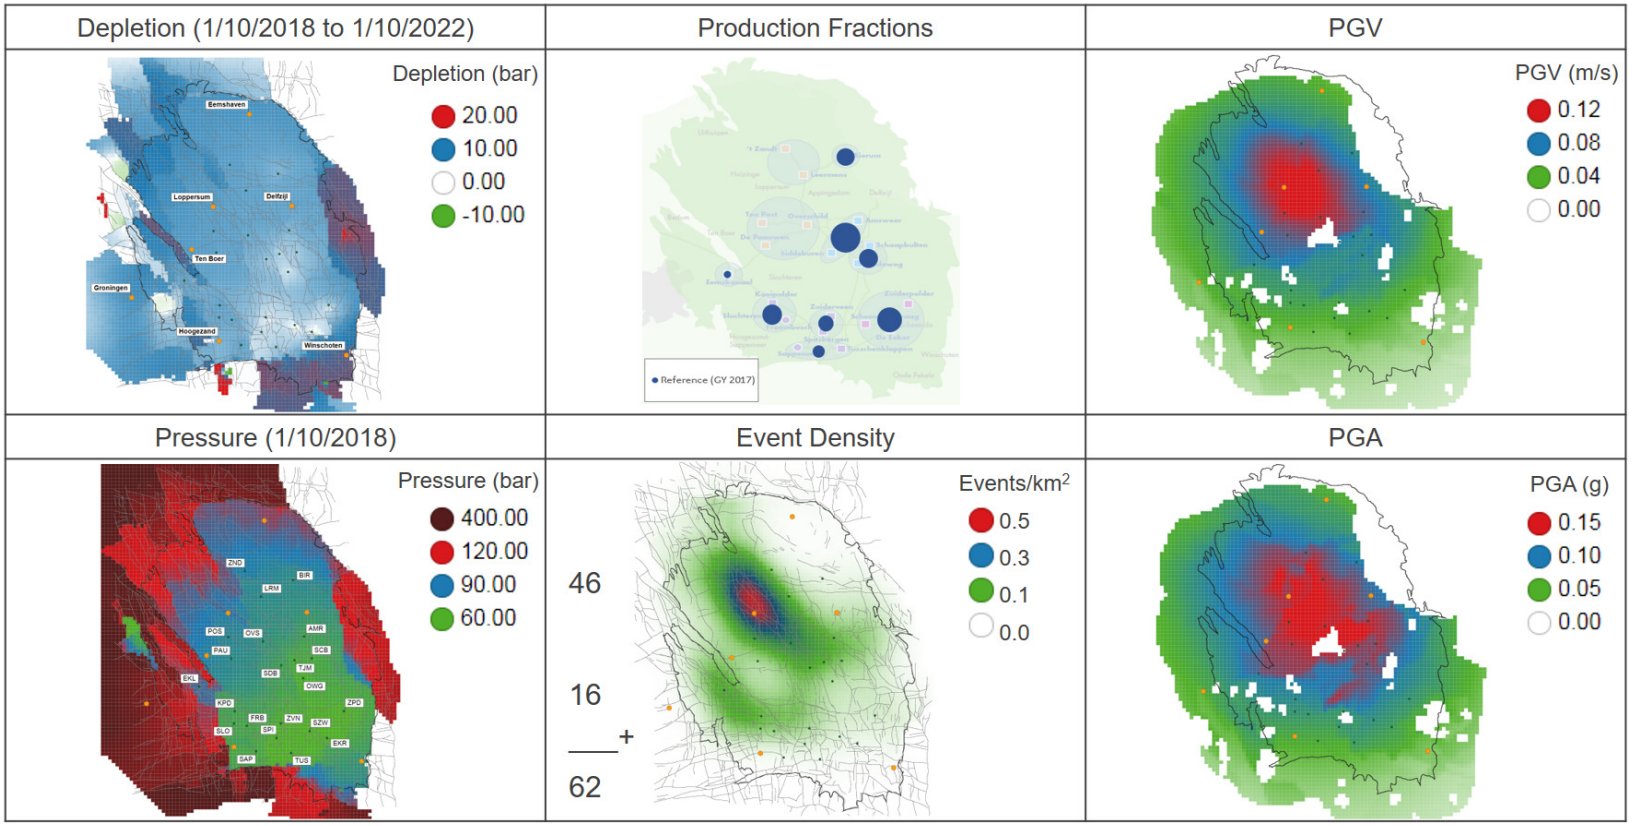

Figure 4-2 *Optimisation reference case – production distribution and subsequent impact on pressure, event density and hazard*

Optimization results – Event Count

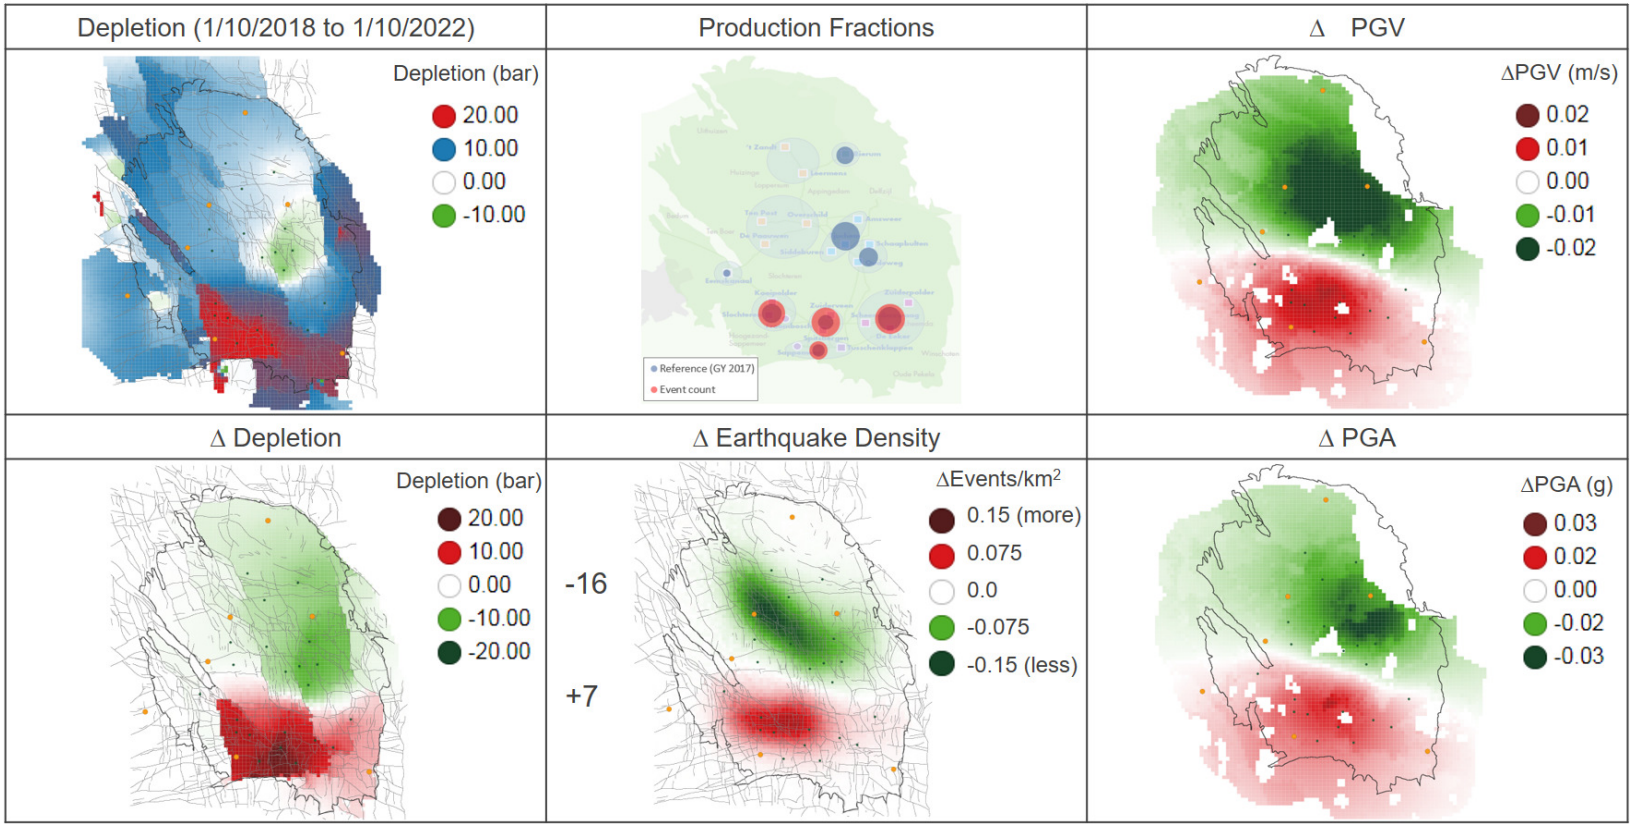

Figure 4-3 *Optimisation for Event count – Changes with respect to reference case (Figure 4-2) in terms of production distribution and subsequent impact on pressure, event density and hazard. The relative reduction in event count and maximum values of PGA and PGV are summarised in Table 4-6.*

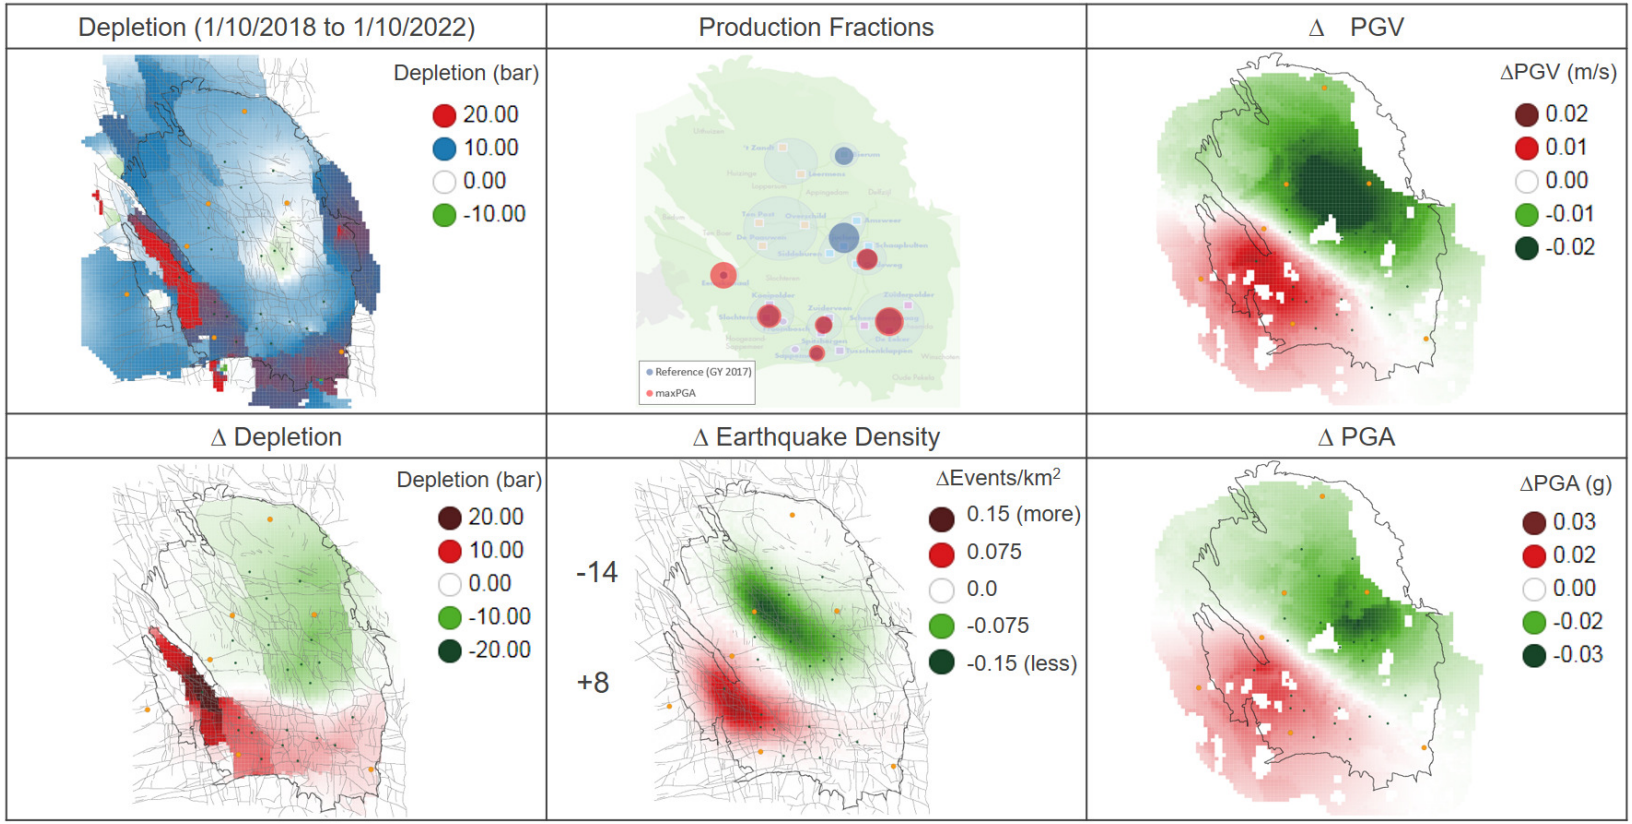

Figure 4-4 *Optimisation for maxPGA – Changes with respect to reference case (Figure 4-2) in terms of production distribution and subsequent impact on pressure, event density and hazard. The relative reduction in event count and maximum values of PGA and PGV are summarised in Table 4-6.*

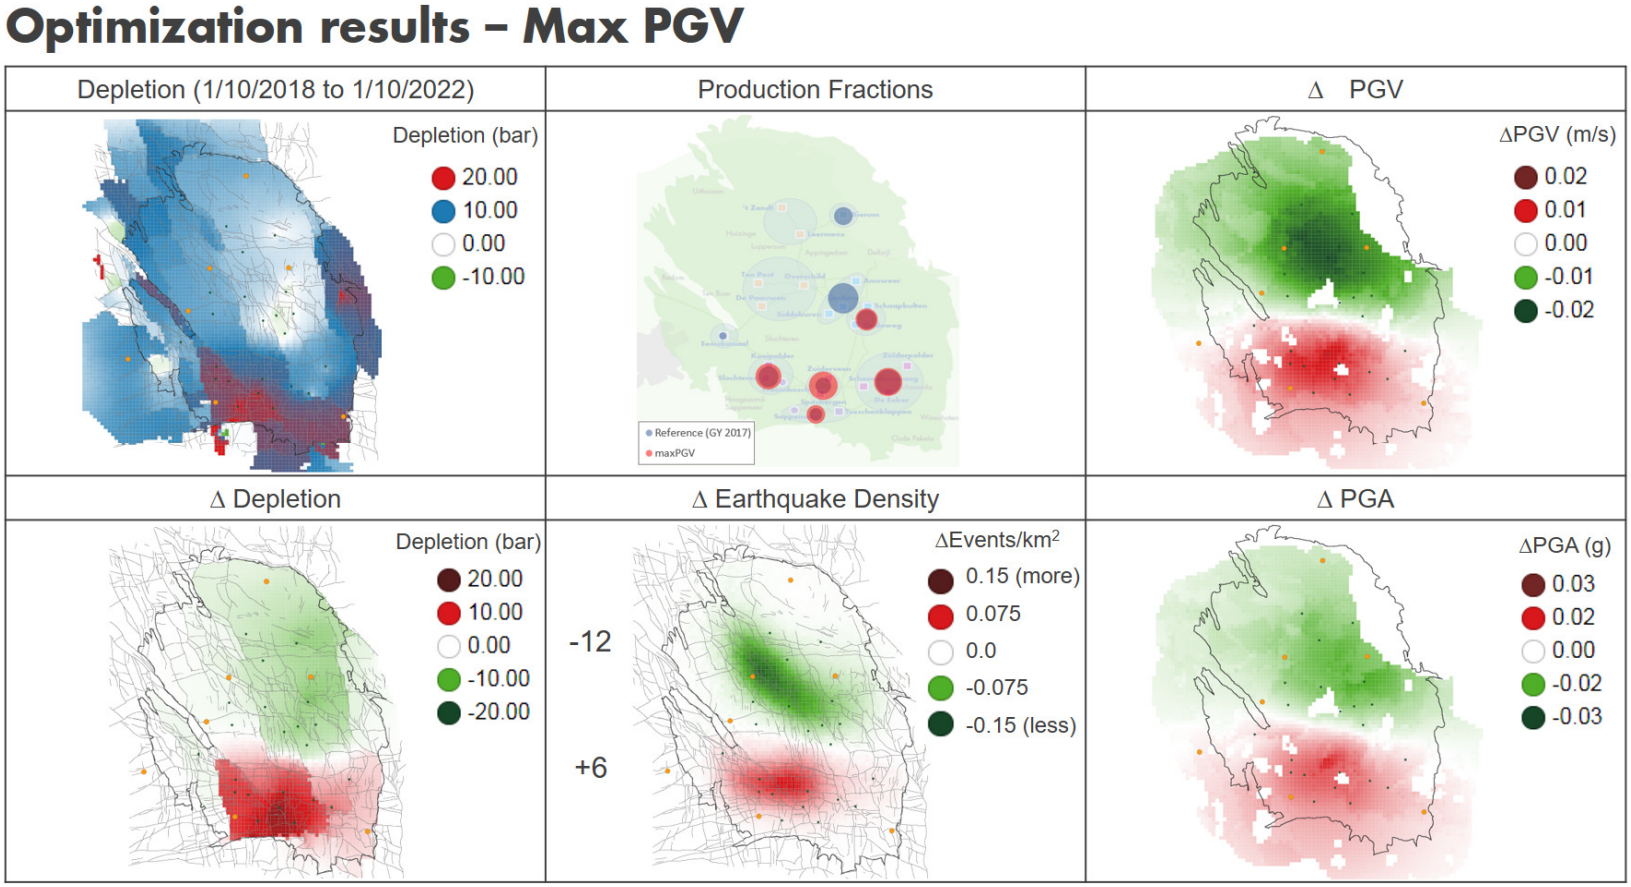

Figure 4-5 *Optimisation for maxPGV – Changes with respect to reference case (Figure 4-2) in terms of production distribution and subsequent impact on pressure, event density and hazard. The relative reduction in event count and maximum values of PGA and PGV are summarised in Table 4-6.*

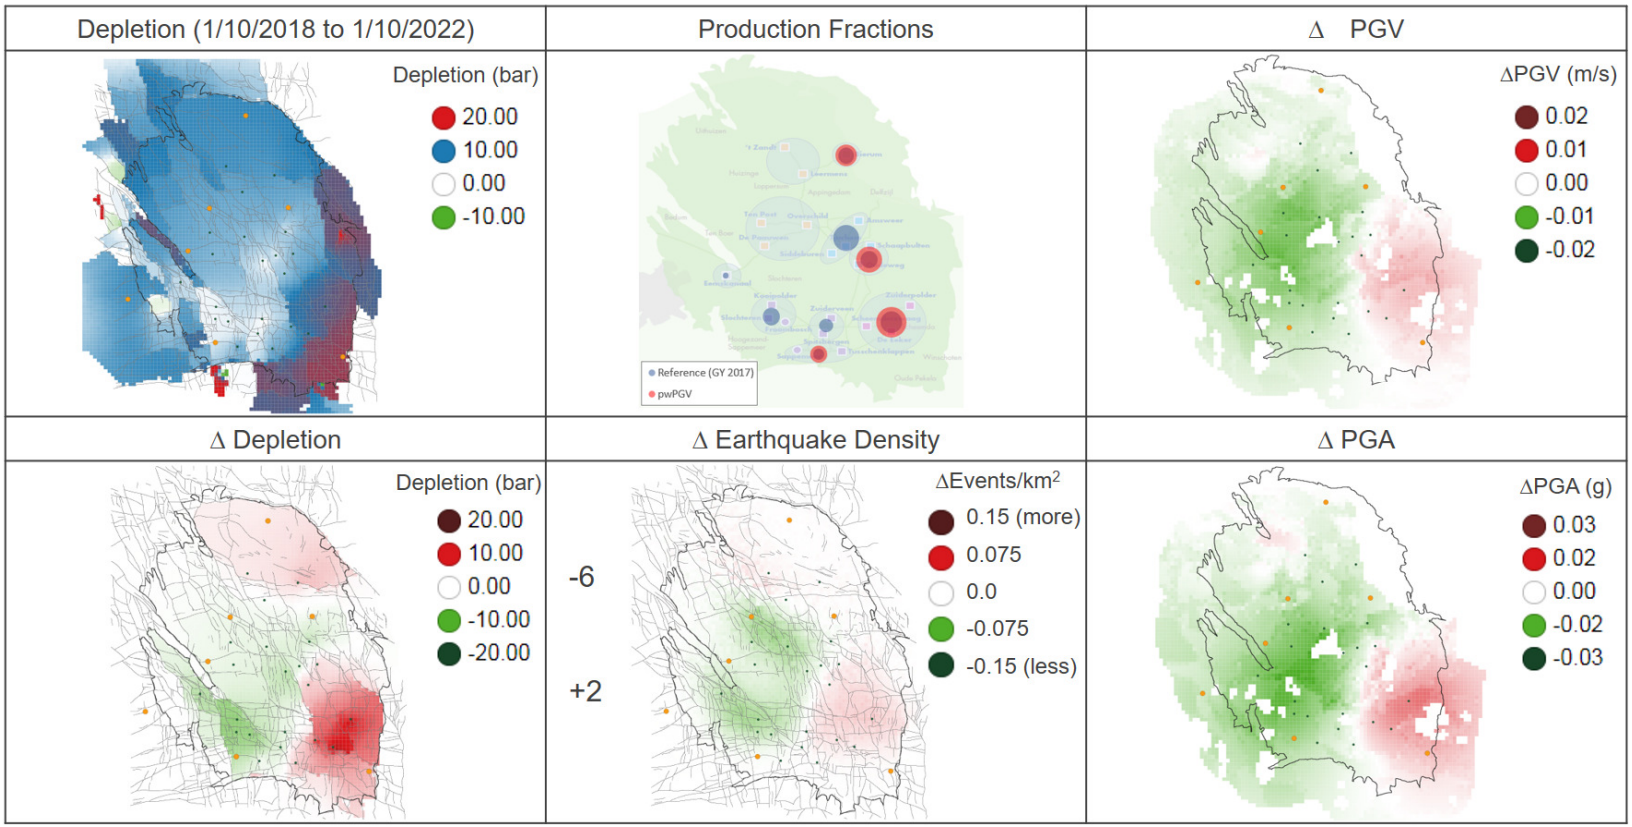

Figure 4-6 *Optimisation for pwPGV – Changes with respect to reference case (Figure 4-2) in terms of production distribution and subsequent impact on pressure, event density and hazard. The relative reduction in event count and maximum values of PGA and PGV are summarised in Table 4-6.*

5 Data driven analysis of model response

5.1 Random Forest proxy models

To further complement the understanding of the behaviour of the HRA model (rather than to only state the optimal production fractions), the complex HRA model was approximated with a simpler but still reasonably accurate proxy model that can be analysed more directly using experimental design type techniques, e.g. variable importance plots and partial dependence plots. The chosen class of proxy models are Random Forests. Although other models may have worked equally well, Random Forests have the following features that contribute to their particular utility in this case:

- Can capture non-linear effects and interactions between controls
- Good out-of-box performance without tuning of algorithm parameters
- Computational efficiency
- Diagnostic tools that allow to extract information like variable importance and partial dependence of model response on a subset of controls.

A detailed description of the underlying theory and further references are provided in Appendix A.

Dedicated Random Forest proxy models were built for each combination of branch and objective function, that can mimic the relation between model input (controls setting) and model output (optimisation metric), to be analysed further.

5.2 Sampling of the control space

To fit the Random Forest models, the control space[8] had to be sufficiently sampled. A detailed description of the sampling strategy is given in Appendix B. Table 5-1 summarises the applied boundary conditions.

control	production fraction	control	production fraction
BIR	$0.01 \le c_i \le 0.25$	FRB_KPD_SLO	$0.01 \le c_i \le 0.35$
EKL	$0.01 \le c_i \le 0.25$	SAP_TUS	$0.01 \le c_i \le 0.16$
AMR_TJM_SDB	$0.01 \le c_i \le 0.51$	ZVN_SPI	$0.01 \le c_i \le 0.30$
OWG_SCB	$0.01 \le c_i \le 0.31$	ZND_LRM	$0.01 \le c_i \le 0.36$
EKR_SZW_ZPD	$0.01 \le c_i \le 0.38$	PAU_POS_OVS	$0.01 \le c_i \le 0.60$

Table 5-1 Sampling space for control fraction of each control

Even though constraints were already applied on the c_i to ensure that (most of the) combinations would achieve the minimal production threshold, some of the more extreme combinations can still drop out. To fit the proxy models, only samples were considered that have at least an average production ≥15.675 Bcm/year (0.2 Bcm deviation). When applying ±0.2 Bcm deviation on total optimisation window (i.e. ±0.05 Bcm deviation annually) instead of annually 4,227 versus 4,366 samples passed the filtering rule

(on averaged over the forecasting period production volume). Limited impact when applying stricter production threshold was observed.

[8] The control space is span by all possible combinations of production fractions for the various controls (the (lumped) production clusters)

5.3 Overall model quality assessment

Random Forest models were trained on the filtered sample sets and the quality of each model is assessed using blind testing.

As shown in Table 5-2, it was established that the normalised RMSE (defined as the root mean squared error of the model normalised by the mean response of the MoReS-HRA coupling) is close to the actual variability in prediction responses from the HRA which is between 1% and 2% depending on the branch and metric, section 3.4.2.

The predictive quality of the model is evaluated out-of-sample. Further details on the methodology can be found in Appendix A. Once the model quality reaches a certain threshold value which comes close to the aleatoric uncertainty (between 1% and 2% depending on the HRA branch), the number of samples is considered sufficient. Through this procedure, it was established that between 3,500 and 4500 uniform samples per branch are sufficient to build a suitable proxy model. The exact number of control configurations that managed to make the minimal production constraints and the out-of-sample performance of the respective Random Forest models for the different metrics that were trained on the available samples is shown in Table 5-2. The estimated standard error for the reported out of sample R^2 was between 1% and 2%, hence there is little variability in this measure of model performance.

Model statistics for Events

Branch number	6	9	14	22	24
Number of samples	4366	4366	2855	2855	4366
Out of sample R^2	0.95	0.95	0.95	0.95	0.95
Normalized RMSE in %	1.2%	1.2%	1.2%	1.2%	1.2%

Model statistics for maxPGA

Branch number	6	9	14	22	24
Number of samples	4366	4366	2855	2855	4366
Out of sample R^2	0.88	0.89	0.90	0.87	0.84
Normalized RMSE in %	1.2%	1.1%	1.2%	1.3%	1.4%

Model statistics for maxPGV

Branch number	6	9	14	22	24
Number of samples	4366	4366	2855	2855	4366
Out of sample R^2	0.89	0.89	0.90	0.89	0.86
Normalized RMSE in %	1.6%	1.4%	1.5%	1.8%	1.9%

Model statistics for pwPGV

Branch number	6	9	14	22	24
Number of samples	4366	4366	2855	2855	4366
Out of sample R^2	0.86	0.89	0.91	0.85	0.78
Normalized RMSE in %	0.9%	0.7%	0.7%	1.4%	1.6%

Table 5-2 Random Forest model statistics for the various metrics

In order for the model diagnostics tools to be meaningful (e.g. variable importance and partial dependence plots), the controls need to be reasonably uncorrelated. The uniformly random sampling setup provides some asymptotic guarantees in that regard. To establish that no high correlations arise by chance, the correlations between the controls were additionally checked explicitly, see Table 5-3. It was indeed found that no strong correlations are present in the controls for the different branches. On average there is a slightly negative correlation between controls which is due to the fact that all the production fractions of all controls need to sum up to 1.

Correlations between controls for samples from branch 6

	BIR	EKL	AMR_TJM_SDB	OWG_SCB	EKR_SZW_ZPD	FRB_KPD_SLO	SAP_TUS	ZVN_SPI	ZND_LRM	PAU_POS_OVS
BIR	1	-0.06	-0.11	-0.08	-0.11	-0.08	-0.04	-0.08	-0.10	-0.13
EKL		1	-0.12	-0.07	-0.11	-0.1	-0.04	-0.08	-0.06	-0.11
AMR_TJM_SDB			1	-0.15	-0.15	-0.15	-0.06	-0.11	-0.17	-0.2
OWG_SCB				1	-0.13	-0.11	-0.06	-0.10	-0.10	-0.15
EKR_SZW_ZPD					1	-0.15	-0.06	-0.14	-0.14	-0.12
FRB_KPD_SLO						1	-0.07	-0.10	-0.12	-0.16
SAP_TUS							1	-0.06	-0.06	-0.08
ZVN_SPI								1	-0.13	-0.14
ZND_LRM									1	-0.15
PAU_POS_OVS										1

Table 5-3 Correlation between controls for samples from branch 6

5.4 Random Forest model analysis

Using the Random Forest models, variable importance plots and one dimensional partial dependence plots were generated.

Variable importance plots

Variable importance plots are used to estimate the overall contribution (including non-linear and interaction effects) of a control on the (Random Forest) model response. An example plot is contained in Figure 5-1, which shows a variable importance plot for a RF model of branch 6 that predicts pwPGV. The variables are ordered decreasingly with respect to their cumulative impact on pwPGV. The associated standard error in the estimates are indicated by the whiskers on the bars. This type of plot does not allow to draw any conclusions about the nature of the relationship between a variable and the objective function, as in "an increase/decrease in production from region X leads to an increase/decrease in metric Y". However, this plot gives an indication about the order in which controls should be investigated based on their impact on the metric. How the average effect of changes in a control can be assessed will be discussed in section about partial dependence plots. Figure 5-1 implies that changes to controls `[POS_PAU_OVS]`, `[EKR_SZW_ZPD]` and `[EKL]` have the largest impact on the modelled pwPGV response for branch 6. Additionally, variables that are deemed insignificant in this representation, i.e. variables whose estimated increase in MSE is close to the standard error of the measurement are variables that can be changed without significant changes in the general model response. Based on Figure 5-1 this would imply that changes to controls `[ZND_LRM]`, `[ZVN_SPI]` and `[SAP_TUS]` have virtually no effect on the model response for branch 6. Since the

model is known to be a good proxy for the coupled Mores HRA setup, it can be concluded that those parameters are essentially in the model null-space.

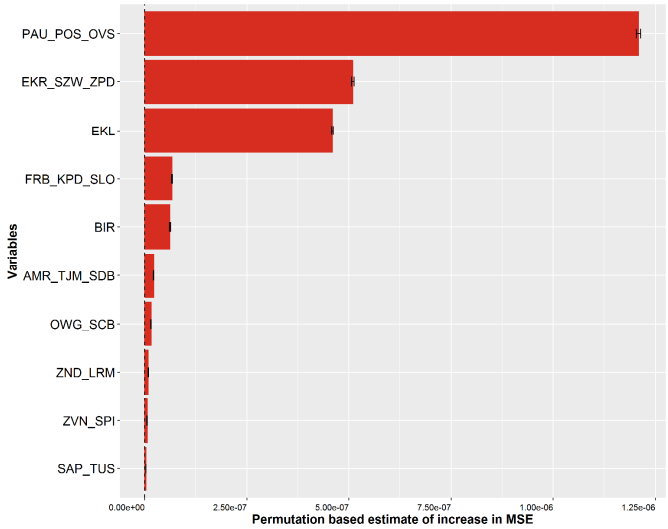

Figure 5-1 *Variable importance plot for population weighted PGV (pwPGV) on Branch 6*

Partial dependence plots

The partial dependence analysis provides a qualitative indication of whether, for a fixed total field production, increasing or decreasing the relative offtake percentage of a model control (or production cluster) would be beneficial for reducing a prediction metric (such as maxPGA, pwPGV, etc.). A set of example plots for the average modelled pwPGV response on changes in individual controls is shown in Figure 5-2.

For any (deterministic) model realisation, the total of the production fractions over all controls adds up to 1. Changing the production fraction for a single control would involve having to make a choice on how to offset that change by an opposite change in the total production fraction from all other controls. However, any specific choice would impact on the outcome. The power of the partial dependence plots is that it allows to evaluate the impact of a single control with respect to an “averaged” response of the rest of the system.

A positive slope means that an increase in the production fraction of that specific control on average increases the value of the objective function (i.e. increases the Hazard/Risk). A negative slope means that increasing the production fraction on average leads to a reduction in the value of the objective function (i.e. getting closer to the optimum). A zero slope means that the value of the objective is essentially not affected by the production fraction of that control.

It was found that the overall relationships are mostly linear and that there are no strong interaction effects (if that were not the case, higher dimensional partial dependence plots would be required).

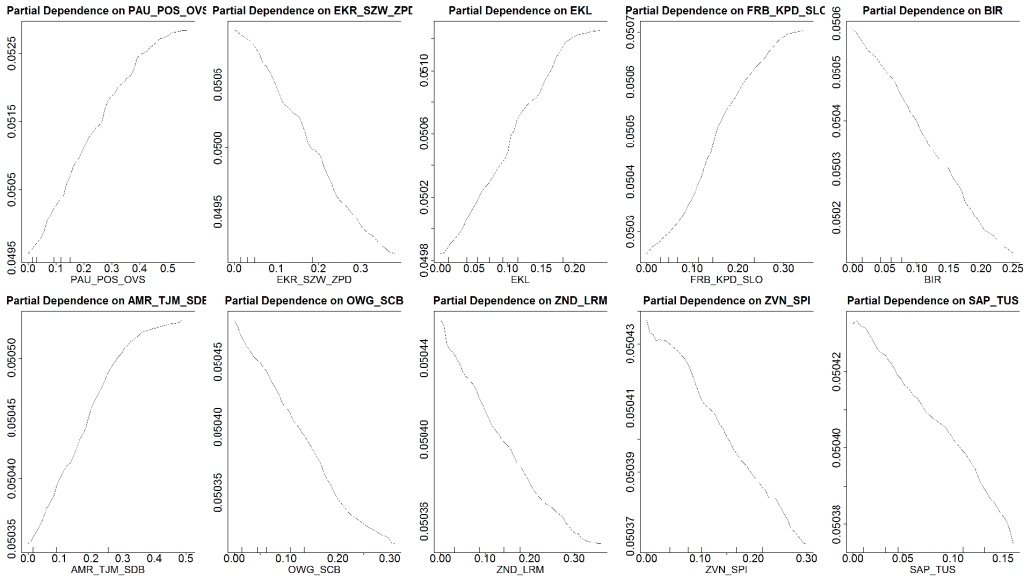

Figure 5-2 Partial dependence plot for population weighted PGV (pwPGV) on Branch 6

In Figure 5-3 to Figure 5-6 partial dependence analyses are provided for the various objective functions. Each figure contains four elements:

- On the top left-hand side, a trellis of the partial dependence plots for all 10 controls are given with respect to the seismicity metric. For each control, the 5 branches[9] are plotted as a dedicated colour.
- On the bottom left-hand side, the corresponding best fit linear slopes are given for all controls, again colour coded by branch. The arithmetic average from the 5 branches is given as a yellow dot.
- On the bottom-centre graph, the quality of fit for the linear slopes is given (R^2), again with a yellow dot for the arithmetic average over the 5 branches.
- The map on the right-hand side super-imposes the average slope for each control on top of the map that indicates the areal location of the 10 controls over the field. The averaged slopes are plotted in a traffic light style.
    - The red part of the colour spectrum reflects the positive slopes. For those controls, an increase in the production fraction will increase the objective function (i.e. increase the hazard or risk).
    - The green part of the colour spectrum reflects the negative slopes, hence increased production from those controls will reduce the hazard or risk.
    - The white controls are indifferent, more or less production will not impact the hazard or risk.

It can be observed from Figure 5-3 to Figure 5-6 that the slopes of these partial dependence curves are highly consistent across all branches: they are either all positive, all negative, or all zero. All slopes are also roughly linear; the relationships can be approximated by fitting a linear trend line through each slope (R^2 values close to 1). For each control, the partial dependence curves for the various

[9] As described in section 2.6, only 5 branches were evaluated from the total 24 branch uncertainty tree for Hazard

branches generally show a fairly linear behaviour as a function of its production fraction. Given this behaviour, the controls were regarded fairly independently. For each control, best fit linear slopes were derived from the partial dependence curves for all branches. Next the average from these slopes was taken. These “averaged slopes” were visualised on a map.

Some clear patterns stand out from Figure 5-3 to Figure 5-6. The `[PAU_POS_OVS]` control in the middle of the seismic-prone Loppersum area has consistently the most negative impact on the seismic hazard and risk. The `[EKR_SZW_ZPD]` control in the South-East of the field is consistently positive.

For each objective function, the consistency of the partial dependencies across the various branches clearly comes out: either all positive, negative, or (near) zero. This provides a clear guidance for the optimisation. There is an epistemic uncertainty range, which is span by the branches of the hazard tree. Despite the fact that we don’t know which branches best approximate the truth, we can still optimise because all branches reflect the same directionally.

Depending on the objective function: the controls favour production from the more densely populated South of the field for hazard, and the sparsely populated North of the field for risk.

Figure 5-3 Analysis of the partial dependence curves for Events

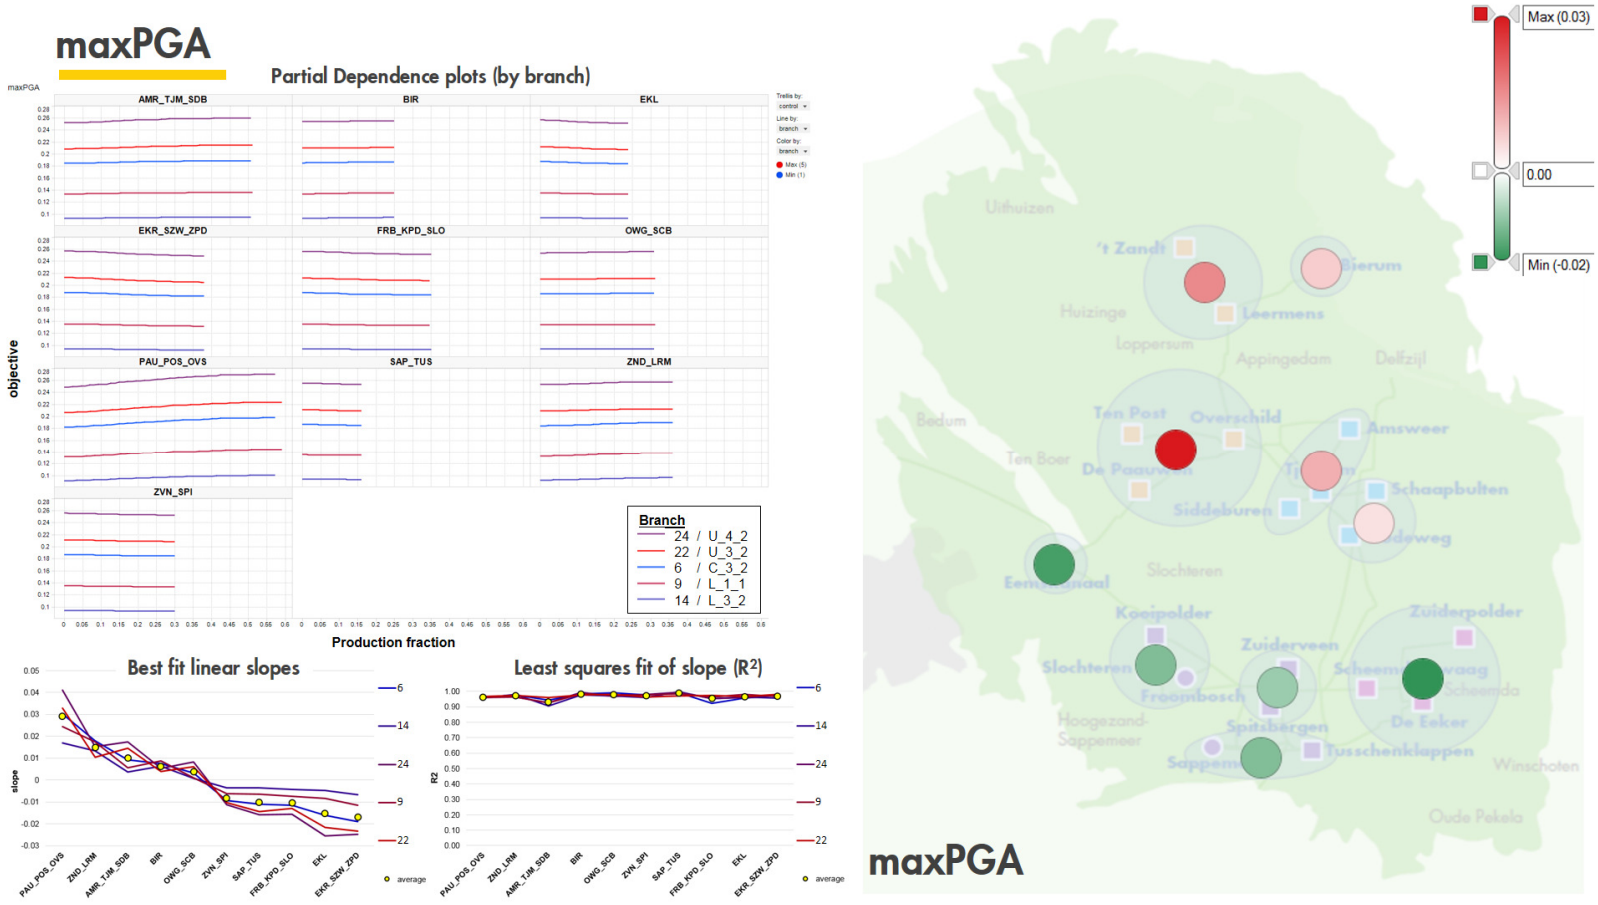

Figure 5-4 Analysis of the partial dependence curves for Maximum PGA

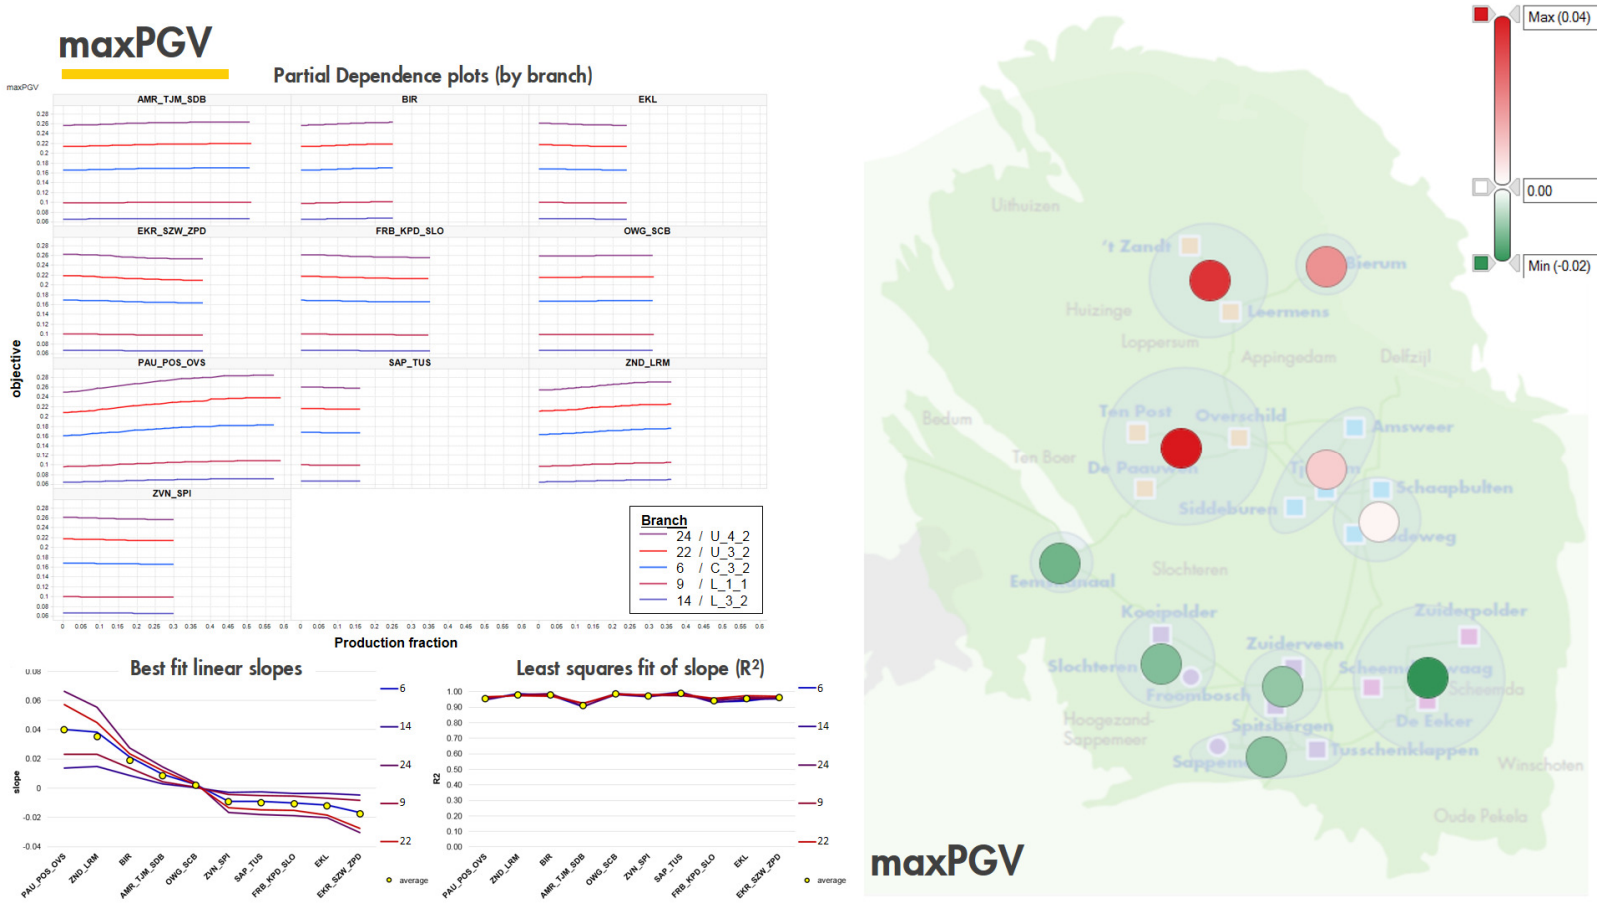

Figure 5-5 Analysis of the partial dependence curves for Maximum PGV

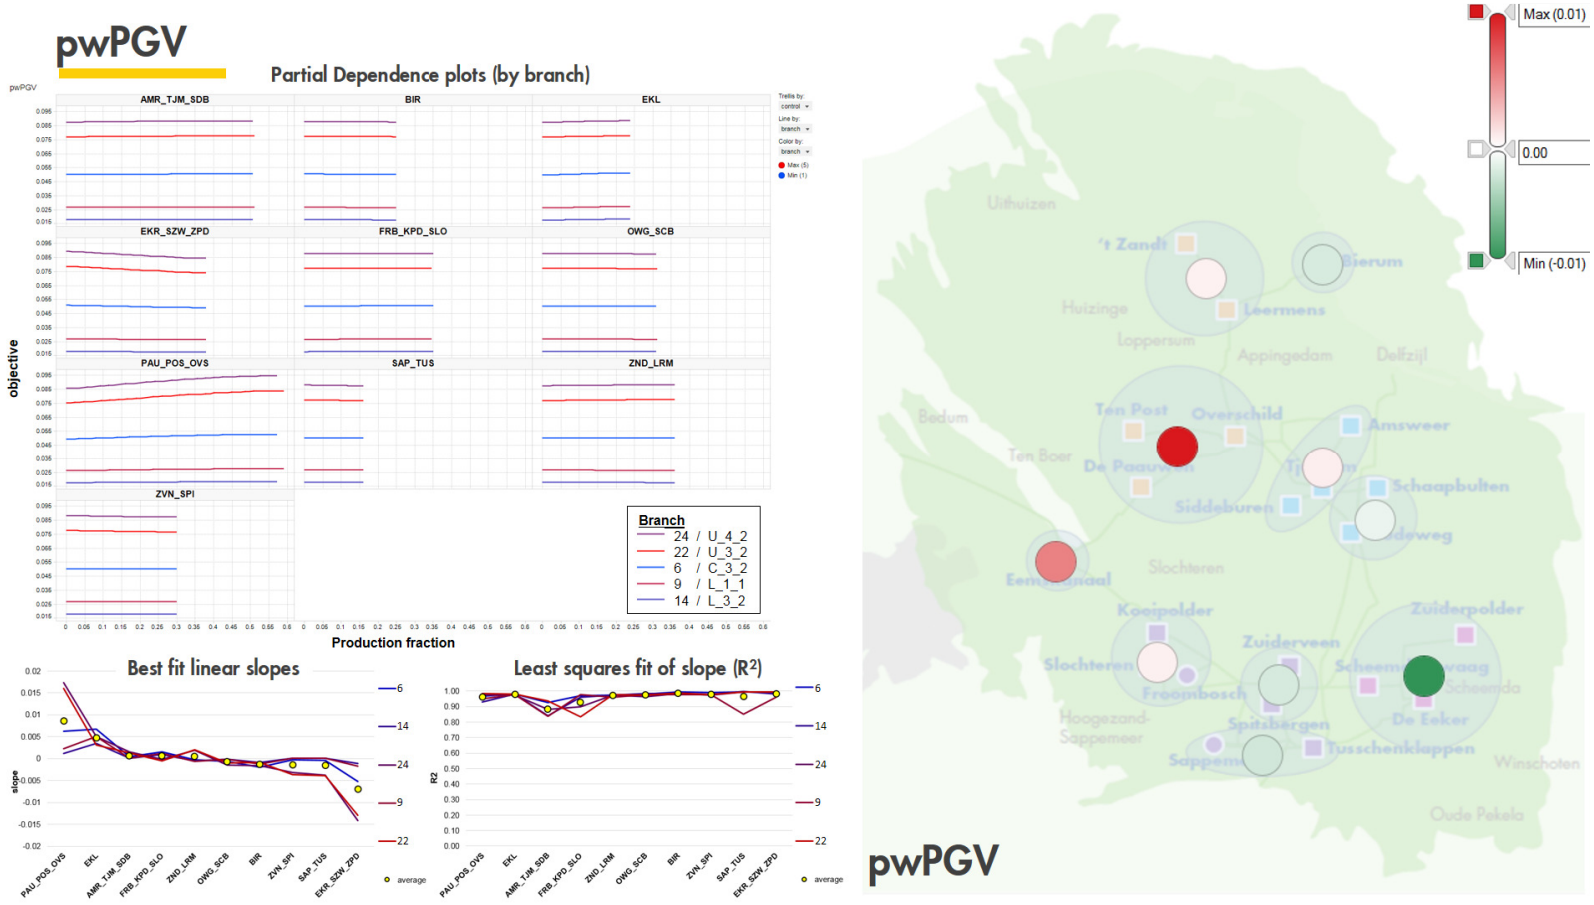

Figure 5-6 *Analysis of the partial dependence curves for Population Weighted PGV*

5.5 Observations from Random Forest models

The results are generally consistent across the examined branches, and it was concluded that there is no need to run more branches on lower and higher sides. There is a good consistency between two lower branches, and between two upper branches for all metrics

Production from the [EKR_SZW_ZPD] control consistently has a favourable impact on seismic hazard and risk for all objective functions. Likewise, production from the [PAU_POS_OVS] control consistently has a negative effect.

[EKL] switches from negative to positive between metrics.

It was a valid assumption to exclude the [PAU_POS_OVS] and [ZND_LRM] controls from the optimisation (chapter 4). In the Random Forest analysis these controls were found to consistently have a negative impact on any of the optimisation metrics.

5.6 Consistency between Random Forest models and Optimisation runs

To check for consistency between the Random Forest models and the optimisation results (chapter 4), the optimum production fractions for the various controls (Table 4-2 to Table 4-5) were superimposed on the Random Forest maps for the various objective functions, Figure 5-7.

In Figure 5-7 the optimum production fractions are displayed as bar charts, with dedicated bars for all evaluated branches of the logic tree. The first bar (indexed by 0) reflects the initial rates, the starting point of the optimisation runs. The optimised branch that yielded the best overall improvement is highlighted (diagonal stripes) in each branch.

Each bar chart also includes a horizontal line, to indicate how far each control is opened up with respect to its total production capacity. These capacities were established from the total production capacity in 2019, normalised to the Basispad Kabinet demand in 2019 (20.4 N.Bcm), Table 5-4.

Production capacity in 2019

Control	BIR	EKL	SDB TJM AMR	SCB OWG	ZPD SZW EKR	SLO KPD FRB	TUS SAP	SPI ZVN	ZND LRM	OVS POS PAU	total
production capacity (N.Bcm/y)	4.92	4.94	10.40	6.26	7.64	7.15	3.25	6.05	7.19	12.17	69.96
normalized production capacity	0.24	0.24	0.51	0.31	0.37	0.35	0.16	0.30	0.35	0.60	3.43

Table 5-4 Production capacities in 2019 (at 90% load factor), normalised to 20.4 N.Bcm (Basispad Kabinet in 2019)

It can be observed from Figure 5-7 that there is generally a good agreement between the Random Forest maps and the optimised production fractions. The “red controls” generally get zero or very low production fractions assigned, and the “green controls” high fractions (in cases all the way up to its maximum capacity).

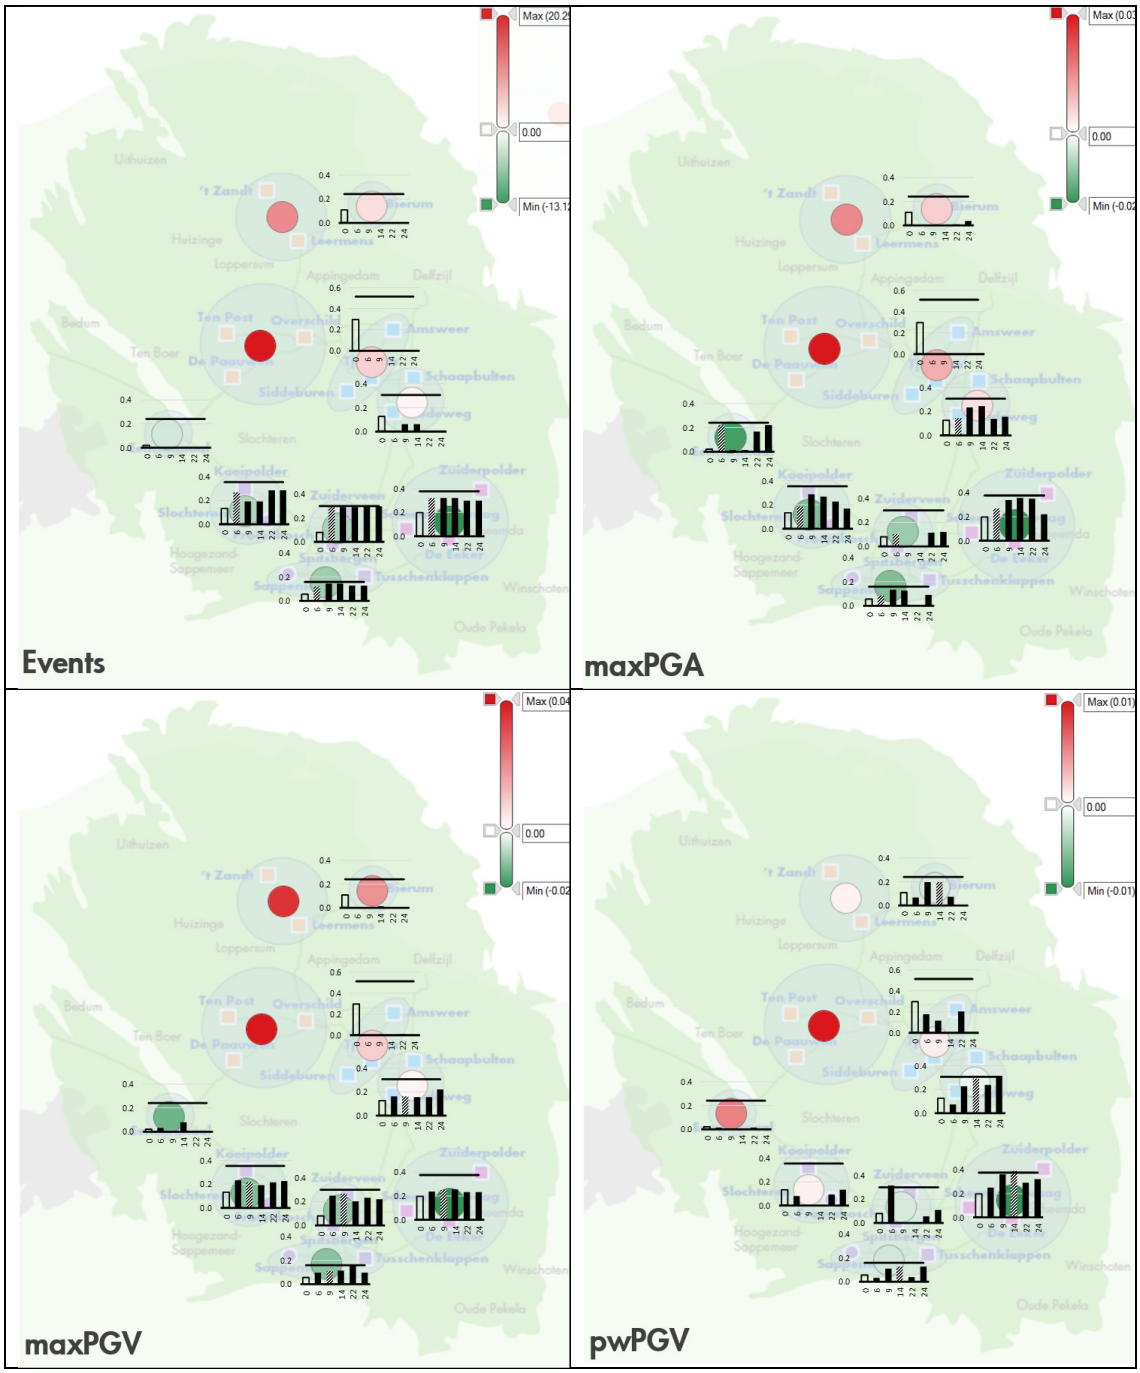

Figure 5-7 Comparison of the Random Forest models (maps) to the optimisation results (superimposed bar charts).

6 Operationalisation of results

6.1 More realistic reflection of operational envelope

The production optimisation described in chapter 4 was consciously implemented to evaluate the theoretical maximum scope for optimisation across the full control space. Each control was allowed to produce anywhere between zero and maximum capacity, with maximum capacity imposed by means of as a simple THP constraint. However, in reality the operational space is smaller. For instance, production clusters are constrained by minimum gas throughput at the production facilities and cannot physically produce at e.g. 1% of maximum capacity. The maximum capacity, in turn, is not governed by a single tubing head pressure. The production rates that can maximally be achieved are governed by several factors, including the compressor envelope, possible compressor speed restriction due to noise constraints, and the backpressure that is exerted on the compressor outlet. The outlet pressure at a production cluster is a function of its location in the pipeline-ring, the number of transfer stations that are in use and their respective rates, the operational settings within the ring (e.g. whether or not the double lines can both be used, possible high/low pressure ring split, etc.), and the backpressure within the GTS network.

6.2 Model description

The importance of accurately modelling the surface facility constraints has long been realised by NAM. Over the last 30 years this has evolved into the Integrated Gas Production System Model (IPSM) software toolsuite called GenREM, a robust, strongly integrated business tool supporting operational gas field development and forecasting. GenREM is in operation for all NAM operated gas fields in the Netherlands and has gained countless capabilities over the years, continuously evolving to address the changing technical and regulatory demands of both the gas fields and the stakeholders of NAM.

Development of GenREM has been a combined NAM/ORTEC effort between mathematicians, IT specialists and Petroleum engineers, resulting in a time-step driven gas capacity and production forecasting tool, using a detailed gas facility treatment network definition. This network solver can be coupled to the Mores subsurface reservoir simulator, resulting in an integrated model capturing the entire Groningen system up to the export points to GTS.

Modelling the Groningen Ring, as previously described in [4], is a challenge for a surface facilities model. For this purpose, a dedicated surface network ring solver has been developed. Being a pressure balanced solver, GenREM is completely thermal driven. Every cluster is modelled in detail, defining every physical Groningen surface device in place. All the cooling devices are based on Heat Transfer Research Institute (HTRI) modelling. In *Figure 6-1* a typical layout of the devices associated with a cluster location is shown.

The Groningen specific design of the centrifugal compressors is completely implemented, up to the level of its efficiency, which is based on a neural network implementation based on actual conditions. The compressor characteristics are based on the envelopes and checked against anti-surge control and speed line limits. Recycling is fully implemented. Multi stage power control usage has been specifically designed for GenREM in close cooperation with the NAM and Siemens engineers (the supplier of the compressors).

GenREM handles constraints for numerous physical parameters like flow, pressure and gas quality throughout the whole surface definition to ensure realistic results of the simulation. Calibration of the input parameters is established by means of big data analysis of real-time production data.

In addition to a realistic representation of the operational envelope, the operational model also includes planned maintenance downtime. Unplanned downtime is captured through an uptime fraction based on historical availability data.

GenREM is being used for short, medium and long-term forecasting of gas production and capacity. Additionally, it is also being used for electricity forecasting due to the high-power demand of the compressors.

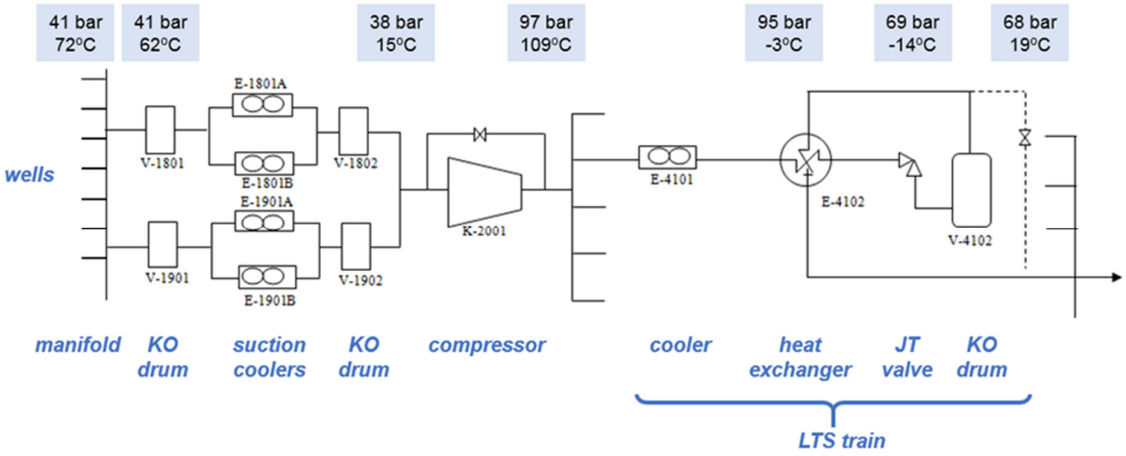

Figure 6-1 Typical layout of surface facilities at a cluster location

It should be realised that GenREM assumes that the installed equipment is always available with no functionality deterioration takinges place over time, other than an assigned uptime factor. Additionally, the restrictions and distribution of the gas across the custody transfer stations into the GasUnie operated pipeline network is not addressed in this optimisation.

6.3 Simulation setup

6.3.1 Production profile

The runs were done for the “Basispad Kabinet” scenario for an average temperature (section 2.5.1) 4-year period starting 1/10/2018. While the optimisation was based on a simplified annual demand profile, these runs are based on the monthly demand profile, as defined in the “Basispad Kabinet”, to ensure that monthly demand variation can be met.

Figure 6-2 *Monthly demand profile as per "Basispad Kabinet". Analysis is done for the first 4 gas years.*

6.3.2 Reference case – WP2016 Addendum June 2018 (EZK)

The reference case for benchmarking the impact of the optimisation was chosen to be the seismic risk assessment for the "Basispad Kabinet" [20], as requested by the Ministry of Economic Affairs, hereafter referred as the "EZK" profile. This HRA was to evaluate the impact of the Ministerial decision to close-in the Groningen field by 2030. The Ministry had prescribed the scenario specifications in its 2/5/2018 Expectation Letter, Reference [14]. The production distribution for this scenario was requested to be based on the 2017 Optimisation Study, Reference [4]. This was implemented by means of a start-up list. Groups of production clusters (aligned with the specified Ministerial regions) are opened up until the prescribed production profile can be delivered. The resulting "EZK" production profile is given in Figure 6-3, with a breakdown by Ministerial region.

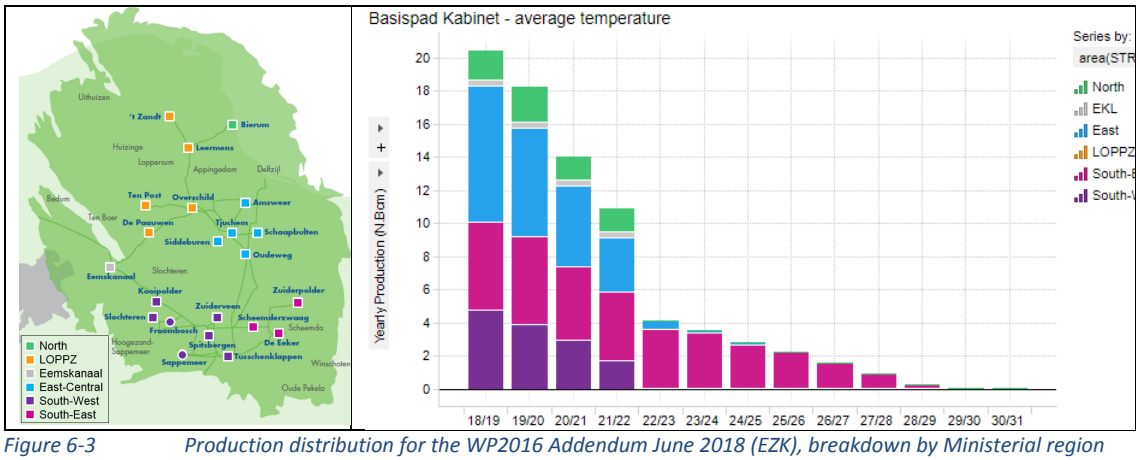

Figure 6-3 *Production distribution for the WP2016 Addendum June 2018 (EZK), breakdown by Ministerial region*

6.3.3 Production scenarios – Start-up list

The same approach as used for the EZK profile was applied here to test the operational viability of the optimised scenarios described in chapter 4. The same 8 controls (section 4.1.1) were implemented as a start-up list in the surface network model. In order to meet monthly gas demand (section 6.3.1), the surface network model starts by opening up the first control on the start-up list, up to a maximum

utilisation factor. If the associated production falls short of the demand, the next control from the start-up list is opened. This process continues until demand is met.

A dedicated start-up list for each objective is given in Table 6-1, based on the Random Forest partial dependence analysis, Figure 6-4. The "EZK" reference case is also included, which was based on the same logic but using the results of the 2017 Optimisation Study [4].

Region	Control	EZK	Events	maxPGA	maxPGV	pwPGV
North	[BIR]	6 mln m^3/d	7	7	8	4
Eemskanaal	[EKL]	2 mln m^3/d	5	2	2	8
South-East	[EKR_SZW_ZPD]	1	1	1	1	1
East-Central (1)	[OWG_SCB]	2	6	6	6	5
East-Central (2)	[AMR_TJM_SDB]	4	8	8	7	7
South-West (1)	[ZVN_SPI]	3	2	5	5	2
South-West (2)	[SAP_TUS]	5	3	4	4	3
South-West (3)	[FRB_KPD_SLO]	6	4	3	3	6

Table 6-1 *Production start-up list for achieving total required field production. Starting from number 1, groups of clusters are sequentially opened-up by the surface network model until the total required production can be achieved. The sequence was derived based on insights from the production optimisation study*

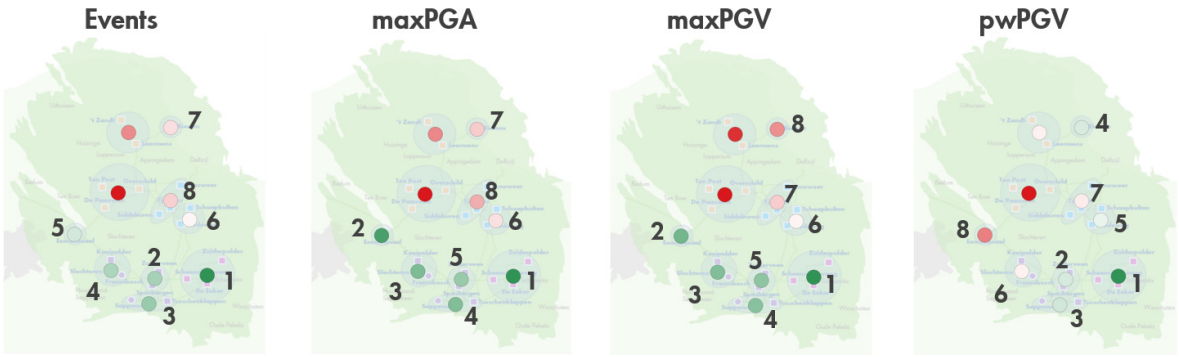

Figure 6-4 *Start-up list as derived from the Random Forrest analysis (section 5.4)*

6.3.4 Model resolution

Results from the full fidelity (GENREM) model are evaluated for hazard using 250m surface resolution (compared to 500m resolution used in chapter 4).

6.4 Scope for optimisation

The gas production capacity of the Groningen field is significantly higher than what is required by the production scenario Basispad Kabinet, even when excluding the Loppersum clusters. Figure 6-5 shows the production capacity prediction associated with the EZK reference case. This excess capacity provides room to redistribute production. In Figure 6-6 the monthly production and capacity for the EZK reference case is shown for the various Ministerial regions. It is clear that the South-East is essentially producing at capacity, as expected given its high priority on the start-up list. Most excess capacity is available in the South-West.

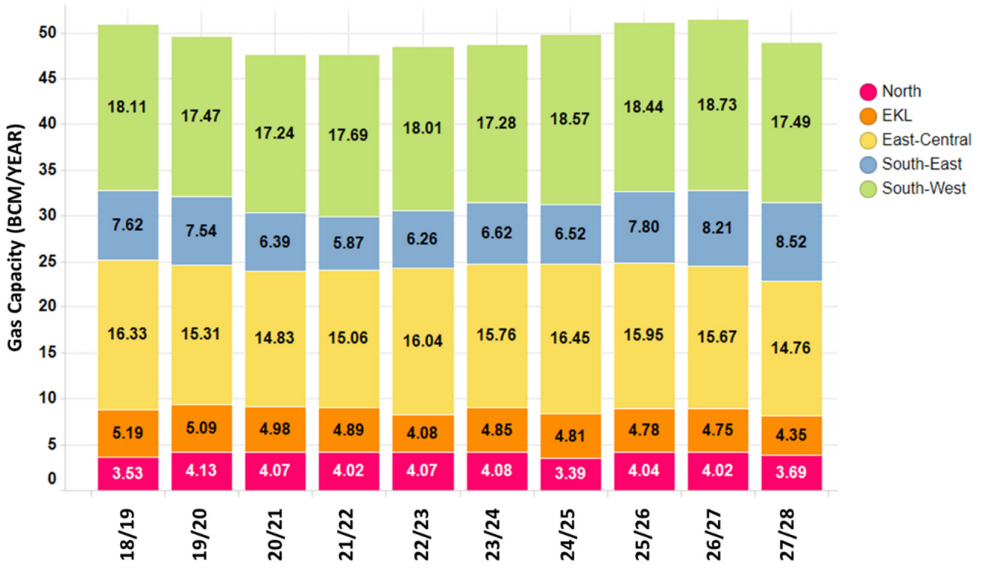

Figure 6-5 *Field gas capacity associated with the EZK reference case (excluding Loppersum clusters).*

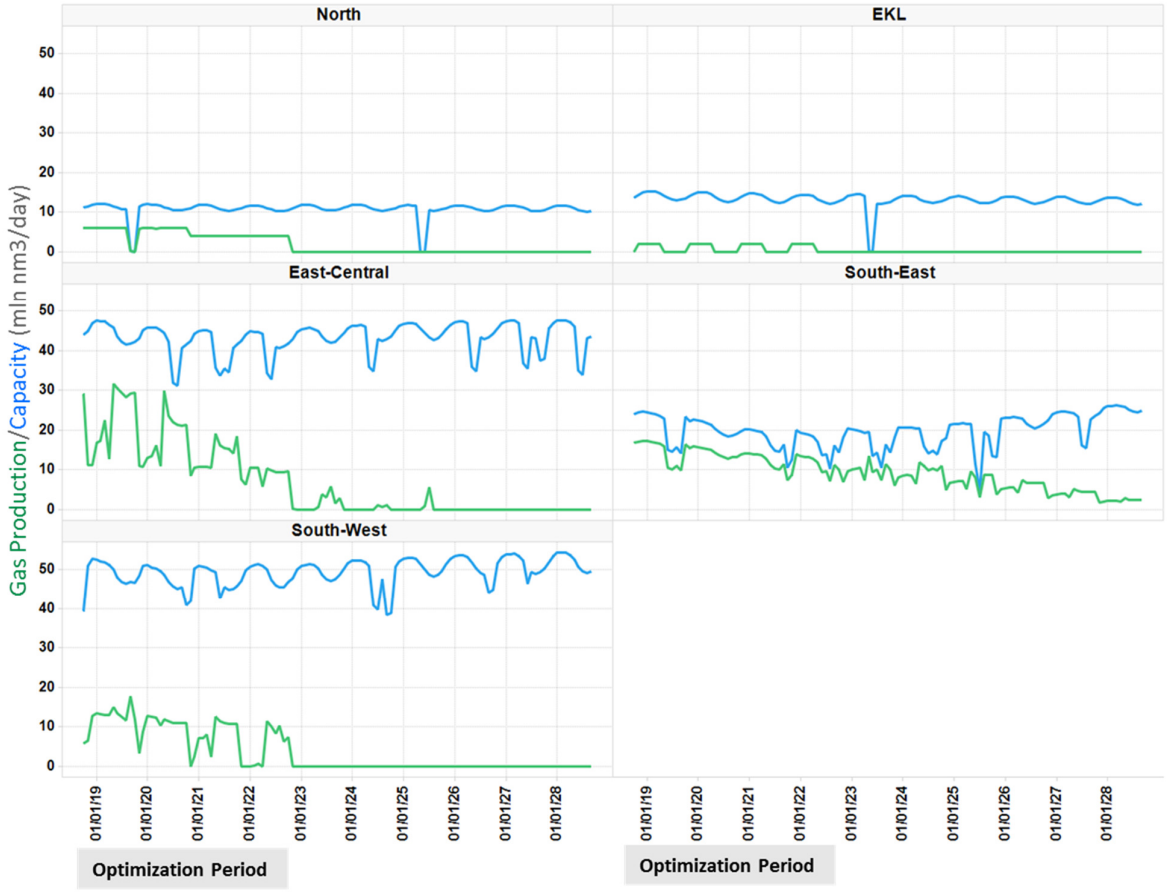

Figure 6-6 *Gas production and capacity as per EZK reference case.*

6.5 Results – Hazard

The model results for the different optimisation scenarios are given in Table 6-2. The biggest reduction can be achieved in **Event count** and **maxPGV** (both about 10% reduction). These results are with respect to the EZK reference case. It is worth noting that the EZK case in itself yields an improvement over the reference case used for the optimisation in chapter 4, which used fixed regional fractions based on 2017 actual production data. The fixed fractions case has also been evaluated in this full fidelity model, referred to as **17Actual** in Table 6-2. With respect to this original reference case, the improvements in the optimisation scenarios are similar to those obtained in the less constrained analysis in chapter 4 (Table 4-6).

The amount of gas being produced from a given region, e.g. East-Central (2) in Table 6-1, is not only a function of its order in the start-up list but also how much of its capacity is being utilised. A high utilisation will maximise volume withdrawal from clusters high in the start-up sequence, but at the same time also increase production fluctuations from clusters low in the start-up sequence (due to fluctuating demand). It should be noted that, based on historic performance, the achieved utilisation range between 70% and 80%. A 70% utilisation factor was found to be a good compromise between optimising production distribution while minimising regional fluctuations. A maximum utilisation of 90% was also tested and found to only give marginal improvements in seismic hazard, as tabulated in Table 6-2. The remaining analysis will focus on the 70% utilisation cases.

Scenario	Utilisation Factor	Hazard Metric				Improvement			
		Events	maxPGA	maxPGV	pwPGV	Events	maxPGA	maxPGV	pwPGV
EZK		61	0.160	0.127	0.041				
17Actual		63	0.162	0.131	0.042	4%	1%	4%	3%
Event Count	70%	55	0.157	0.116	0.041	-10%	-2%	-9%	2%
	90%	55	0.160	0.115	0.041	-10%	0%	-9%	2%
maxPGA	70%	56	0.152	0.116	0.042	-8%	-5%	-9%	4%
	90%	55	0.151	0.115	0.042	-10%	-6%	-9%	4%
maxPGV	70%	56	0.155	0.116	0.042	-8%	-3%	-9%	4%
	90%	55	0.152	0.115	0.042	-10%	-5%	-9%	5%
pwPGV	70%	58	0.153	0.124	0.041	-5%	-4%	-2%	0%
	90%	57	0.152	0.122	0.040	-7%	-5%	-4%	-2%

Table 6-2 Optimisation results for Hazard using the surface network model

In Figure 6-8 to Figure 6-11 the areal impact of the different scenarios is investigated. The observations are largely in line with those from section 4.2.

- All scenarios shift production to the South-West region compared to the **EZK** reference case, Figure 6-8. The result is more depletion (coloured red) in the South-West with respect to the reference case. Consequently, less production is required from the East-Central region, with corresponding less depletion (coloured green). Depletion in the North-West, that includes most of Loppersum, remain largely similar between all scenarios. The **maxPGV** and **maxPGA** scenarios shifts more production to the far South-West compared to the **Event count** and **pwPGV** scenarios.

- The impact on earthquake density is shown in Figure 6-9. The **EZK** reference case estimates 61 tremors over the four-year period, with 47 in the northern main area of seismicity and 14 in the southern area. As expected, earthquake density is higher in the South-West for the **maxPGA** and **maxPGV** cases due to the increased depletion. The **Event count** and **pwPGV** cases show the best improvement ratio of 2:1, with 2 earthquakes less in the North for every 1 added in the South.
- The improvement in hazard, Figure 6-10 and Figure 6-11, is consistent with the observations for depletion and earthquake density. Significant improvement in the East-Central area does come with some increase in hazard in the South-West.

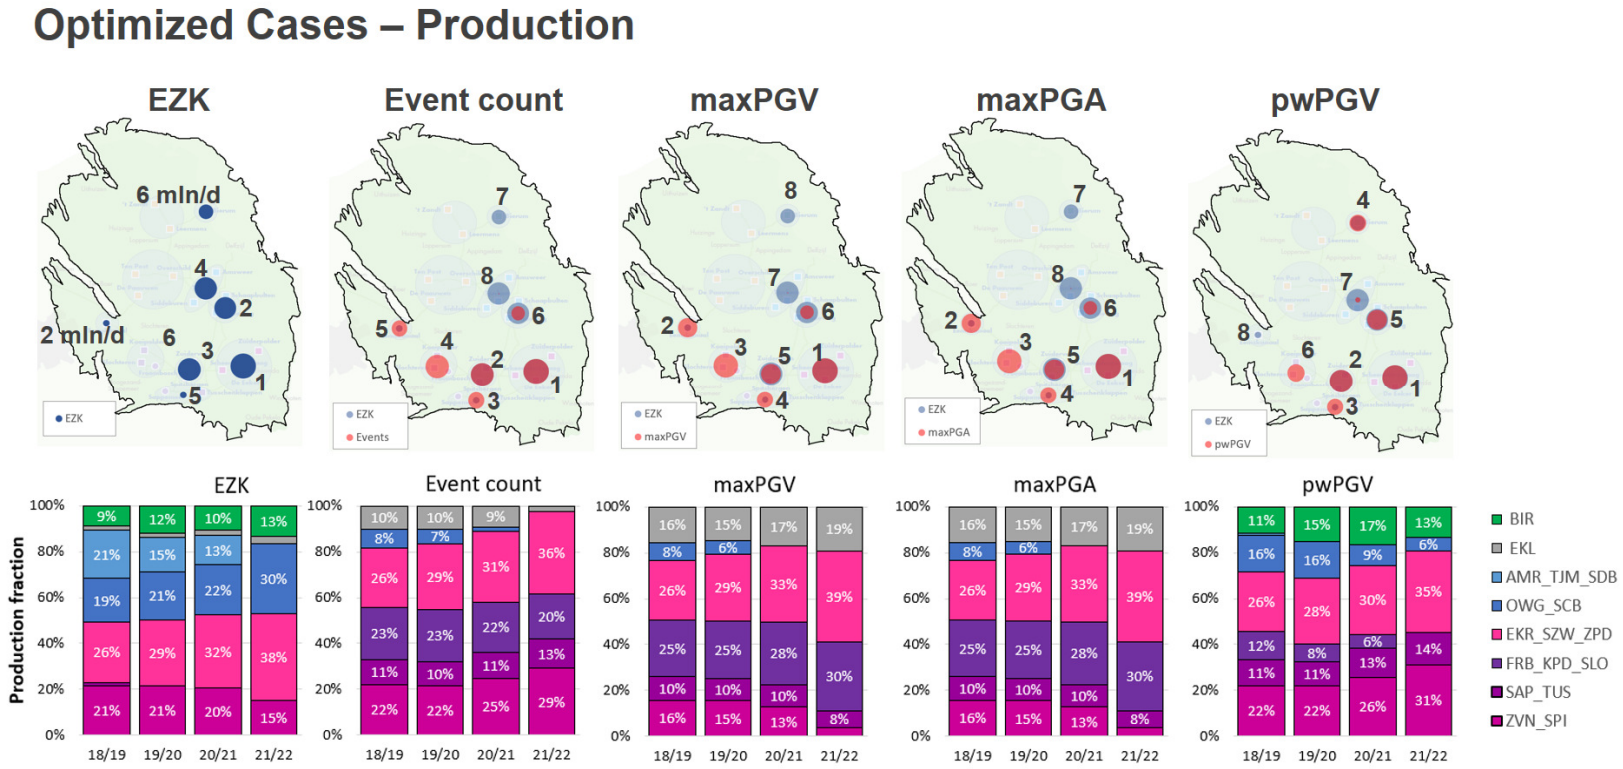

Figure 6-7 *Production distribution for various objective functions, and difference with respect to Basispad Kabinet*

Optimized Cases – Depletion

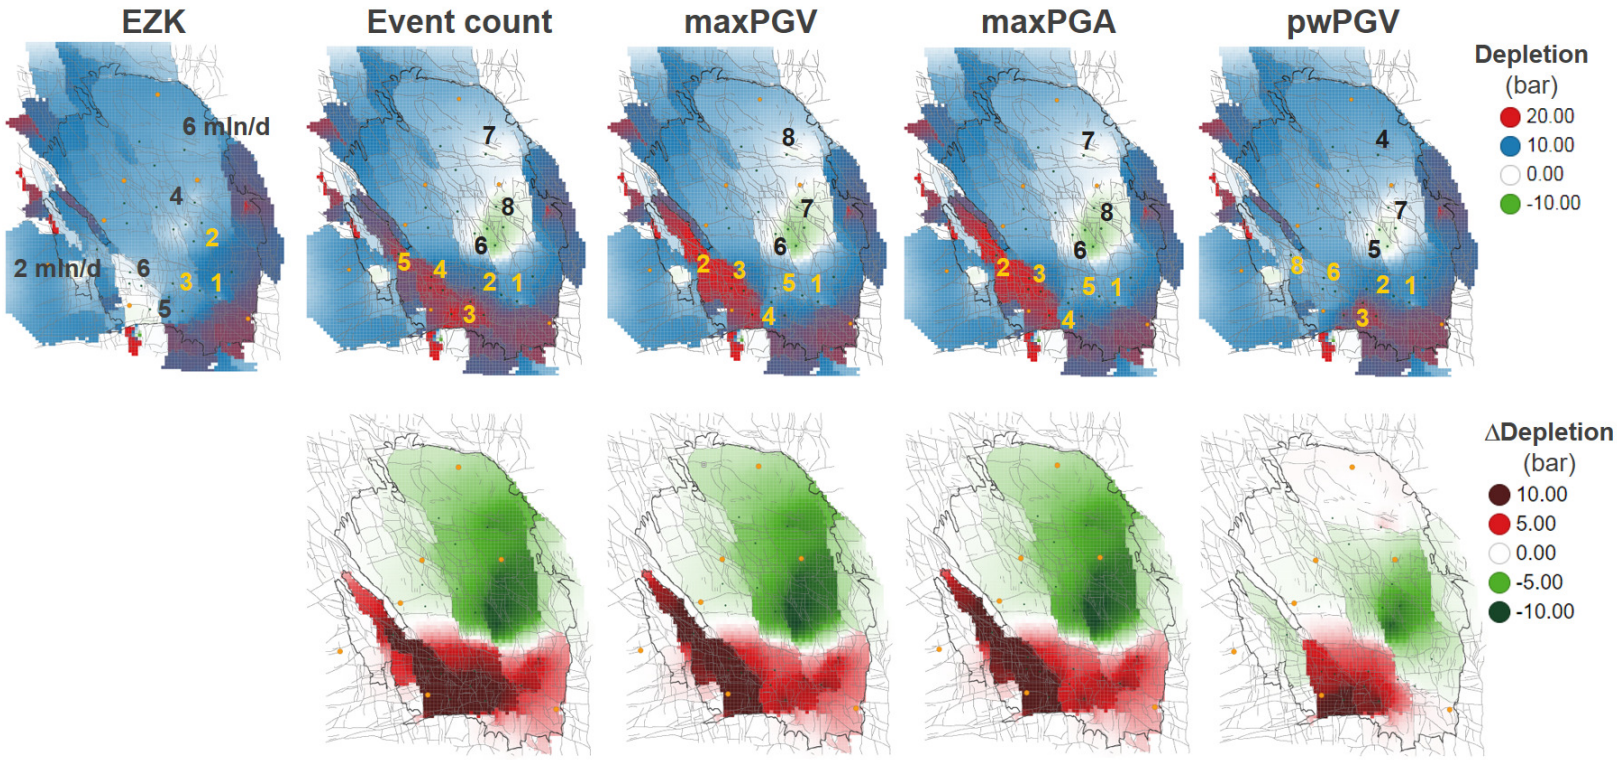

Figure 6-8 *Pressure depletion maps for various objective functions, and difference with respect to Basispad Kabinet*

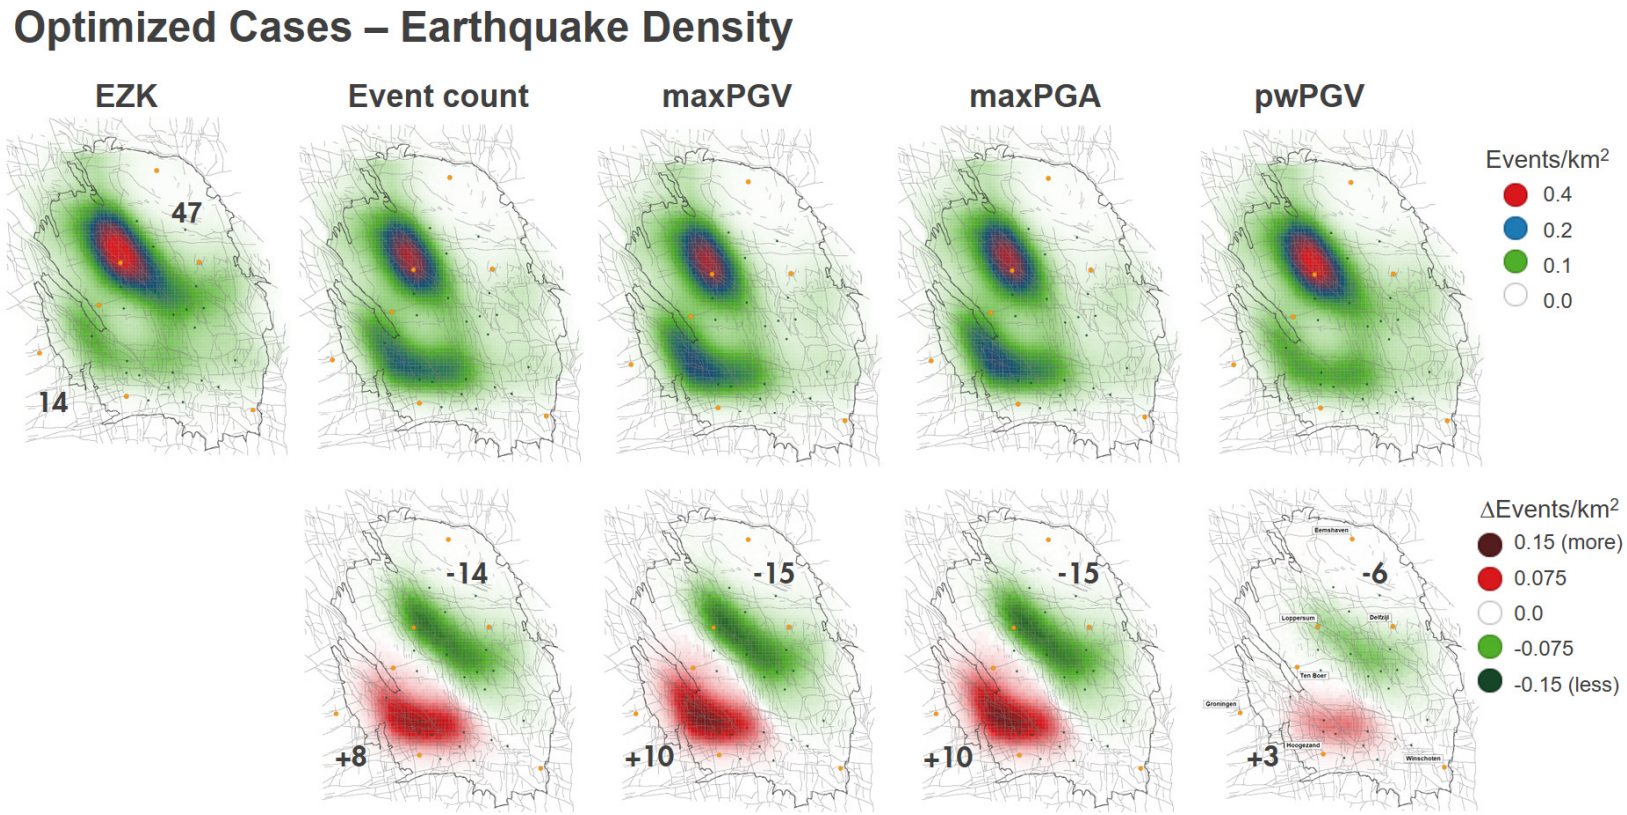

Figure 6-9 *Earthquake density maps for various objective functions, and difference with respect to Basispad Kabinet. The relative reductions in event count are summarised in Table 6-2.*

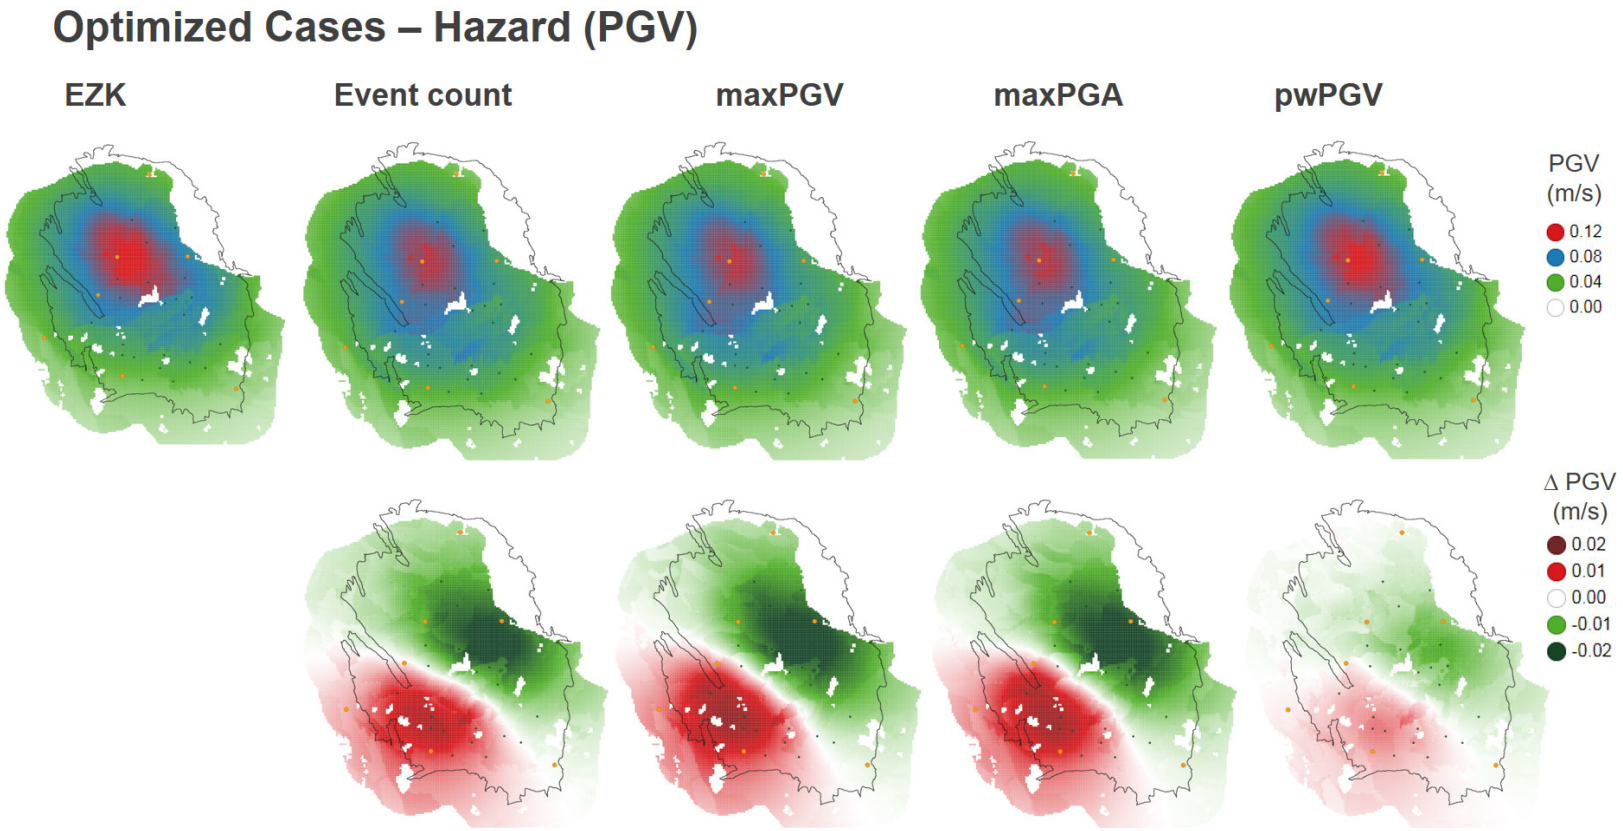

Figure 6-10 *PGA maps for various objective functions, and difference with respect to Basispad Kabinet (EZK). The relative reductions in maximum values of PGV are summarised in Table 6-2.*

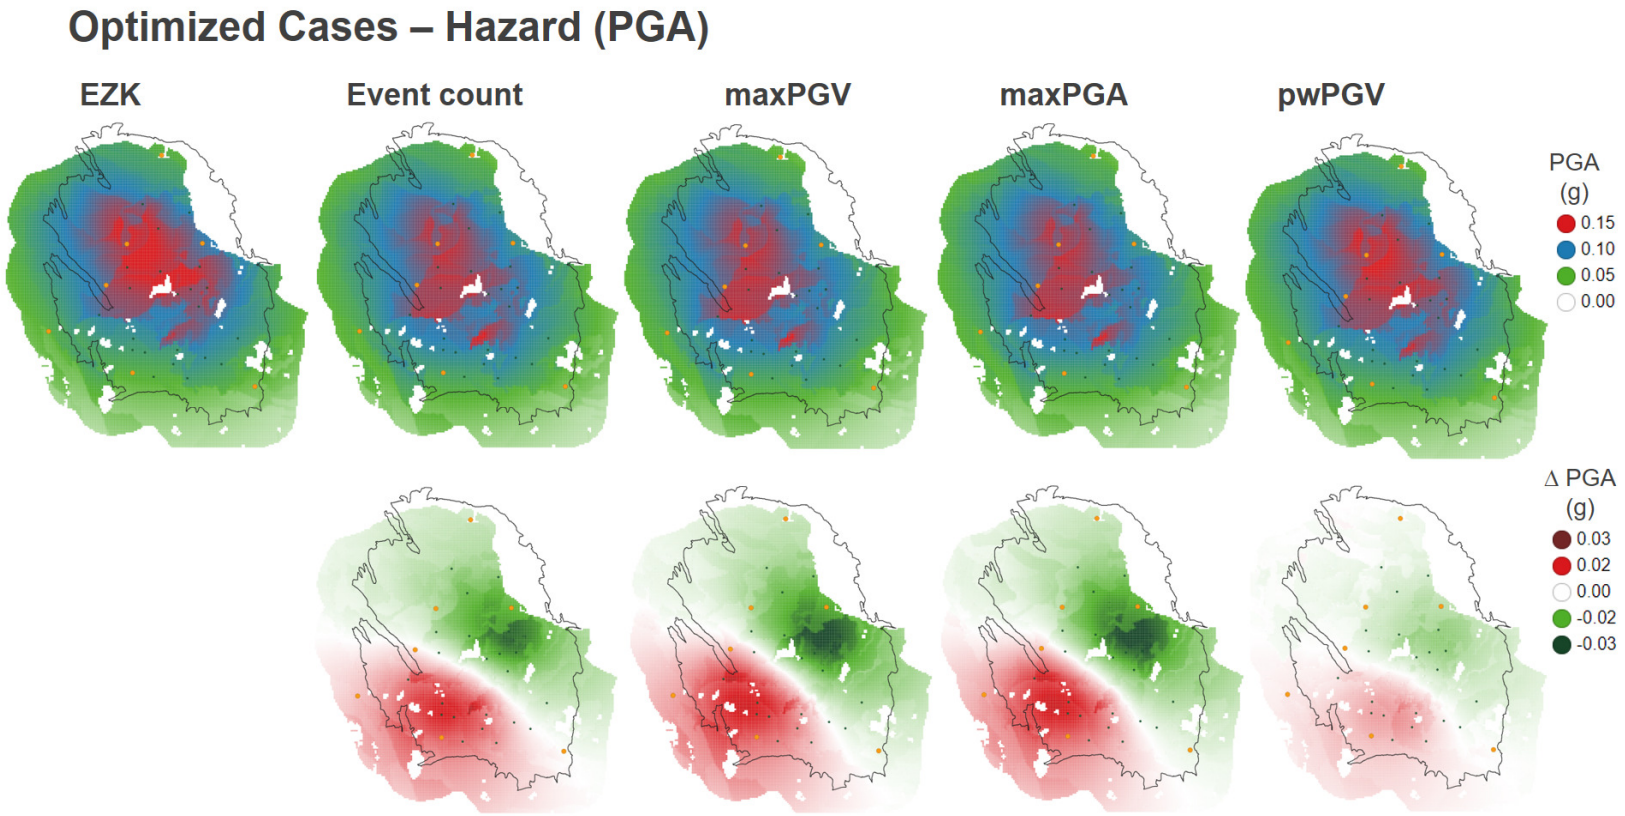

Figure 6-11 *PGA maps for various objective functions, and difference with respect to Basispad Kabinet (EZK). The relative reductions in maximum values of PGA are summarised in Table 6-2.*

6.6 Results – Risk

6.6.1 From hazard to risk – brief overview

As explained in section 2.4.1, given the finite calculation capacity the optimisation work was restricted to the Hazard domain. Optimised production distributions were established using seismicity- (event-count) and hazard-related (maximum PGA, maximum PGV, population-weighted PGV) metrics only. However, once these optimised production distributions are established, they can be evaluated using the full risk tree.

As indicated in section 2.4.4 which introduces pwPGV as an optimisation metric, the move from hazard to risk introduces additional heterogeneities, which make the relationship from hazard to risk highly non-linear. The following elements prevent the derivation of a simple one-to-one relationship between changes to the hazard distribution and changes to the overall risk:

- **Population distribution**
  As indicated in Figure 2-3, most of the population resides in the greater Groningen city area, which mainly lies outside the Groningen field outline. Therefore, all things being equal, changes to the hazard in that region would results in far greater impact on the overall risk (i.e. buildings above the Meijdam norm) than a similar change elsewhere in the field.

- **Distribution of building typologies**
  Beyond the spatial distribution of the population throughout the field, the nature of the buildings themselves plays a crucial role in the translation of hazard to risk. As in the case of the population distribution, had all the buildings been of identical type, changes to the hazard would be readily relatable to changes in risk (accounting for the population distribution). As observed in [21] (and references therein), the majority of buildings not meeting the Meijdam norm have one of four typologies as their primary structural systems: URM1L, URM4L, URM3M, and URM8L. Plotting their spatial distribution throughout the area of interest (see Figure 6-12), it is clear that their distribution, for the most part, is far from uniform. That, in turn, means that a reduction in hazard could result in a lesser reduction of overall risk.

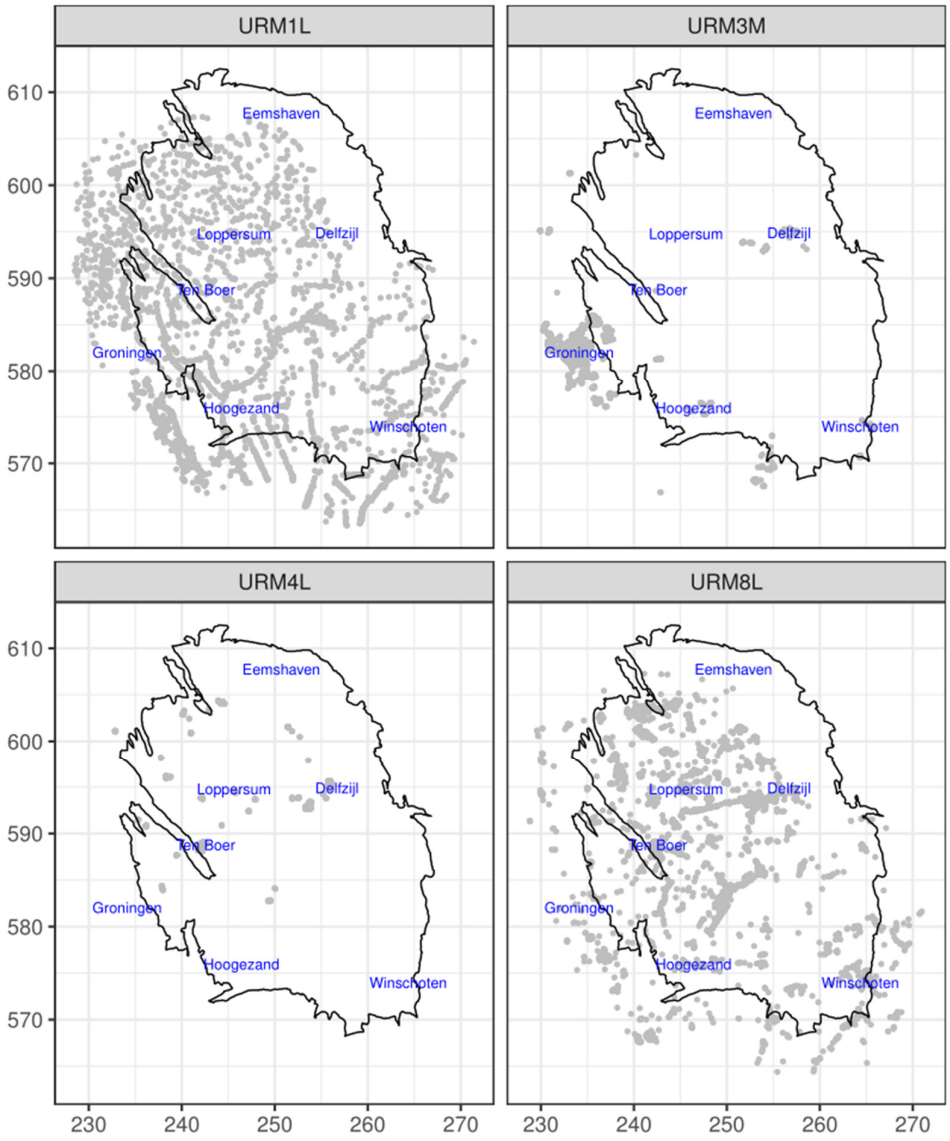

Figure 6-12 Spatial distribution of the vulnerable structural systems throughout the area of interest

6.6.2 Risk outcomes and discussion

When reflecting the facilities and operational constraints in the production distributions as optimised for hazard, it appears that there is significant scope to reduce the number of buildings that exceed the Meijdam risk norm. This can be seen in in Table 6-3 which lists the number of buildings exceeding the Meijdam norm for each of the optimised distributions, including a comparison to the reference case. Note that the exceedance numbers in Table 6-3 are for the 5-year period 1/1/2018 to 31/12/2022. A breakdown including the annual numbers is given in Figure 6-13.

Scenario	Total	Improvement	Unchanged	Removed	Added
EZK	1136				
Event	760	-33%	715	421	45
maxPGA	752	-34%	674	462	78
maxPGV	780	-31%	677	459	103
pwPGV	1033	-9%	1025	111	8

Table 6-3 *The number of buildings with Mean LPR greater than 1E-5 for the different optimised scenarios. **Unchanged** buildings are shared between scenarios, **Removed** are those exclusive to EZK reference case and **Added** are those exclusive to the scenario in question.*

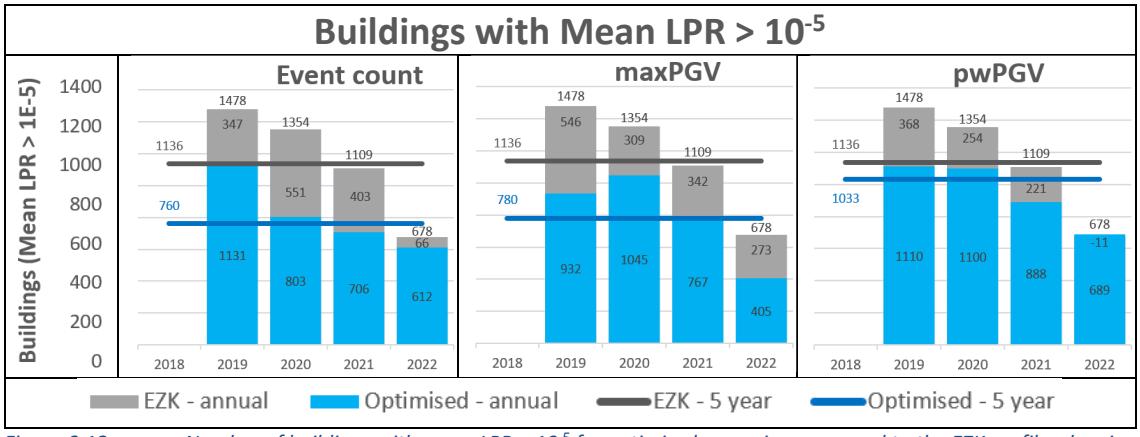

Figure 6-13 *Number of buildings with mean LPR > 10^{-5} for optimised scenarios compared to the EZK profile, showing both the yearly and the five yearly numbers.*

For a more detailed view of the impact of the optimisation on risk, consider the maps shown in Figure 6-14. Each row, associated with one optimisation objective function, contains three maps showing buildings that are – from left to right – in excess of the norm in both reference and optimised cases; no longer above the norm in the optimised distribution; and, finally, buildings that were below the norm in the reference case, but are above the norm in the optimised case.

As can be seen from the maps, buildings with higher LPR than in the reference case are mainly located in the Ten Boer/ Eemskanaal area.

Optimisation impact on buildings ≥ 10^{-5} (Mean LPR)

70% utilisation of production capacity

	Unchanged	Removed	Added
Events (760 buildings)	715 buildings	421 buildings	27 buildings
maxPGA (752 buildings)	674 buildings	462 buildings	78 buildings

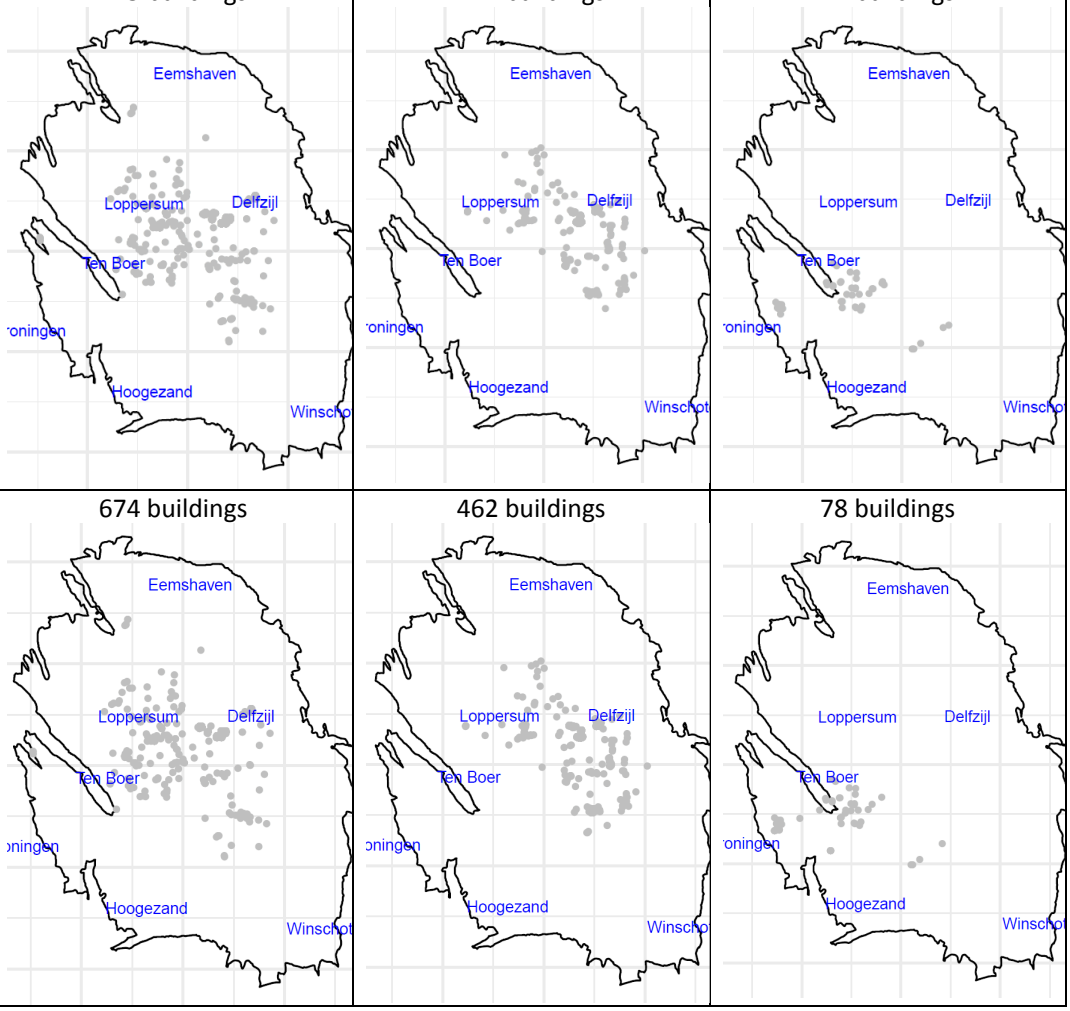

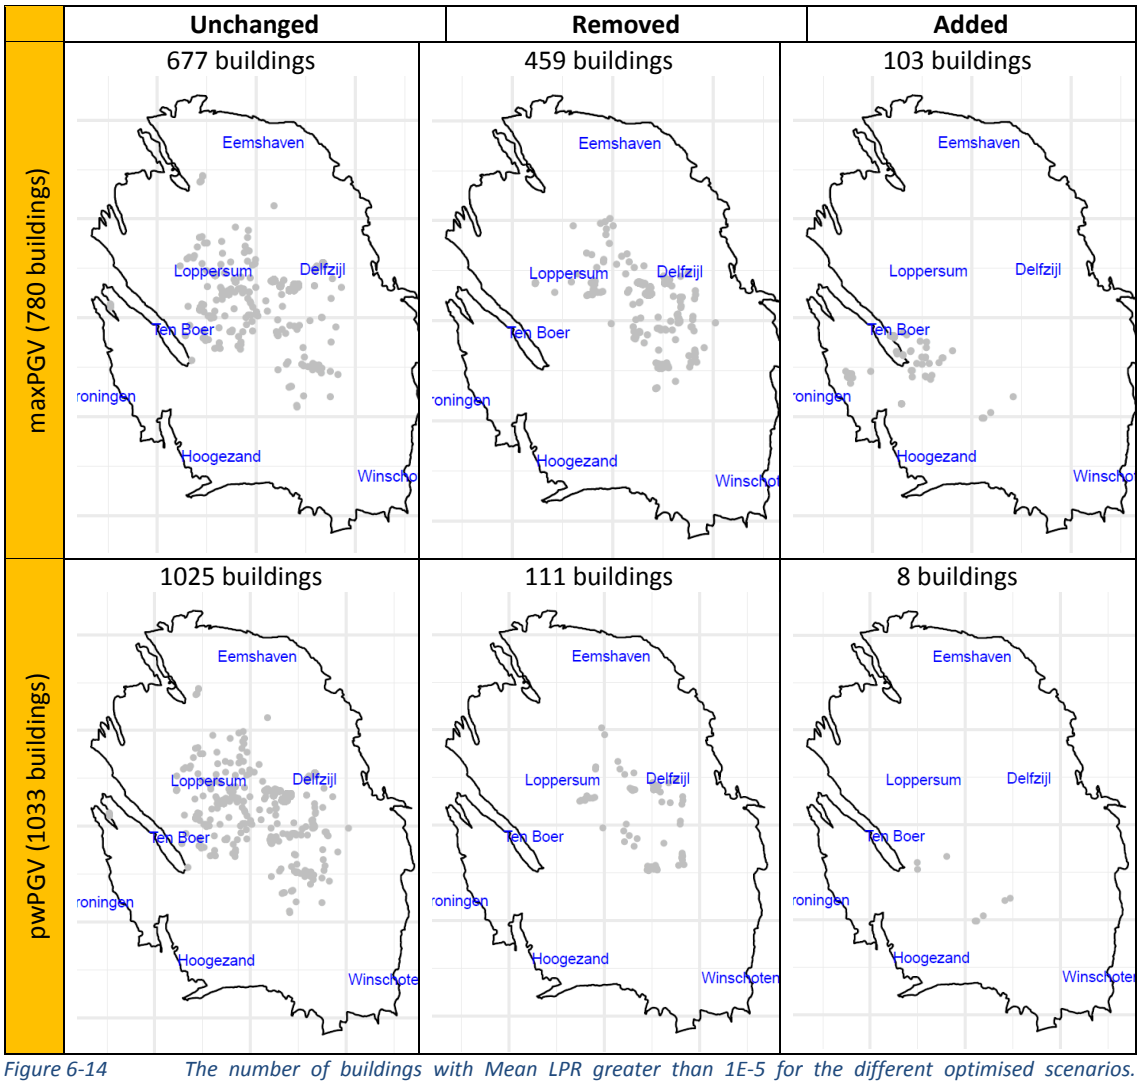

Figure 6-14 *The number of buildings with Mean LPR greater than 1E-5 for the different optimised scenarios. Unchanged buildings are shared between scenarios, Removed are those exclusive to EZK reference case and Added are those exclusive to the scenario in question*

6.7 Non-linear relation between hazard and risk

Optimising directly on risk is not computationally feasible, as outlined in section 2.4.1, but as evident from this chapter, using hazard as a proxy for risk does seem to generally be a valid approach. However, while the dependence of risk on hazard is well understood and modelled, it cannot be distilled into a simple correlation without sacrificing accuracy and applicability. It is thus worth re-visiting some aspects of the chosen hazard metrics.

- The metrics **maxPGV** and **maxPGA** are both single points within the modelled area. Due to the bimodal behaviour of the earthquake density, with peaks both in the Loppersum and South-West area, a relative small reduction in e.g. **maxPGV** can result in quite a large increase in hazard in the South-West as production is shifted in that direction, as depicted in Figure 6-5.
- **Event count** and **pwPGV** metrics are both integrating hazard across the entire modelled area and thus providing a more holistic hazard metric.

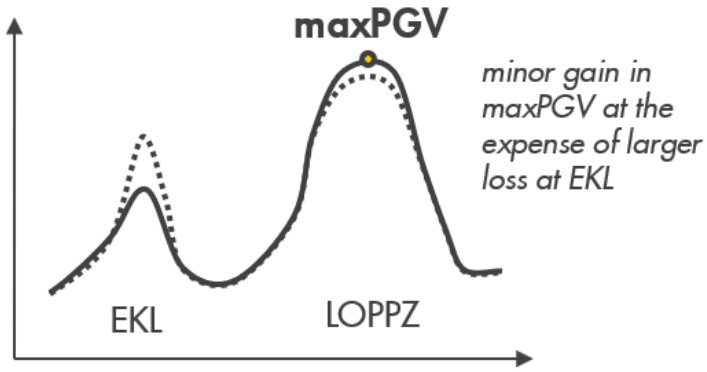

Figure 6-15 Schematic of point-based hazard behaviour.

In addition to the cautionary notes on the hazard metrics, there are also subsurface aspects of the field that has a direct impact on the resulting risk. While most of the field is well connected in terms of pressure communication, the Eemskanaal cluster is in a fault block that is significantly baffled from the remaining South-West area of the field. The associated rock volume is such that 1 bcm of production from here yields more depletion and compaction than elsewhere, and this is largely contained locally. Combined with its location being relatively far from the points of maximum PGV and PGA, is resulting in its positive ranking in the random forest analysis for maximum PGV and PGA.

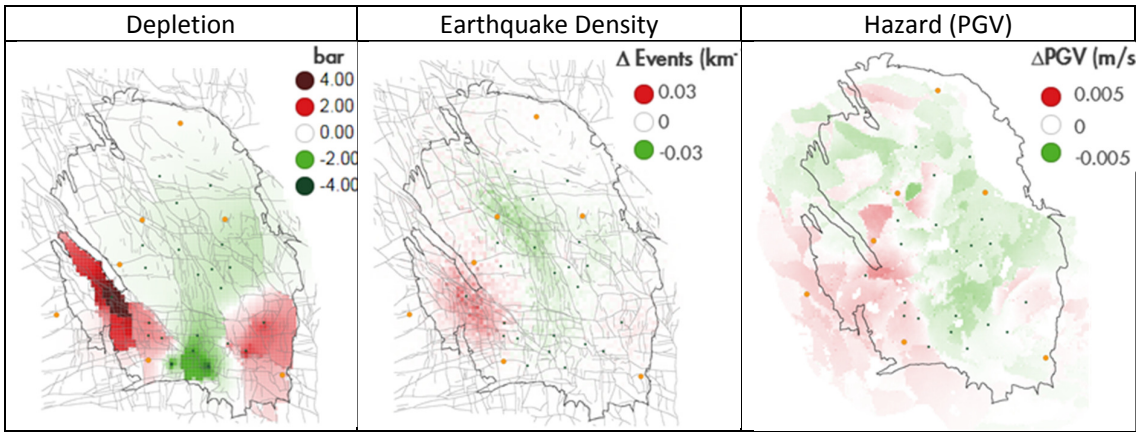

Figure 6-16 *Change in areal properties for maxPGV optimisation going from 70 to 90% utilisation factor.*

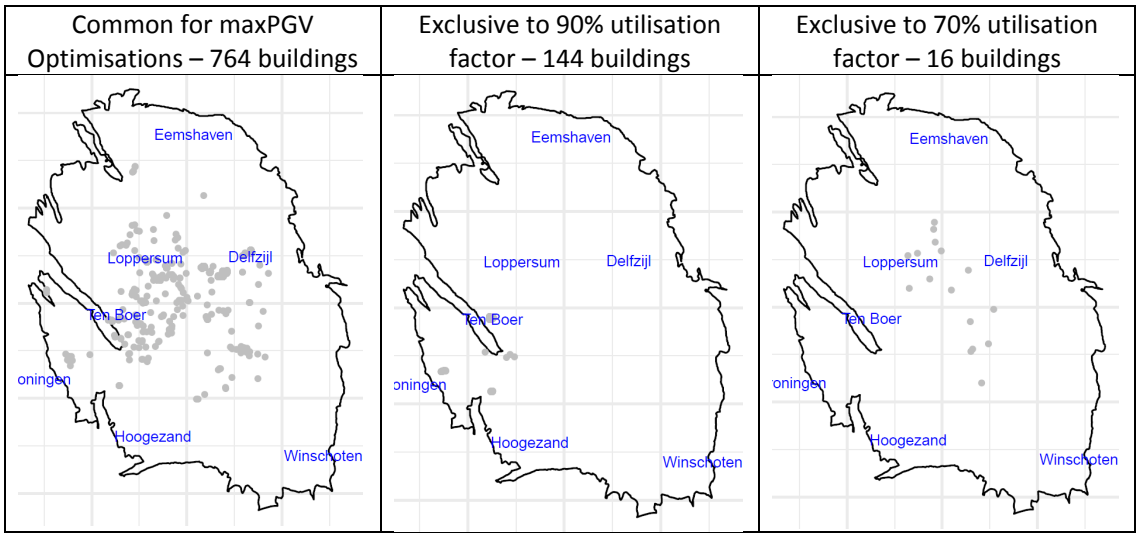

Figure 6-17 *Building above the norm 1E-5 (Mean LPR) for the maxPGV optimised case for different utilisation factors. 70% utilisation factor has 780 buildings while 90% has 908 buildings exceeding the norm.*

The impact on risk has been further analysed using increased utilisation factor. Going from 70 to 90 percent utilisation has relatively modest impact on the hazard metrics as shown in Table 6-2. In Figure 6-16 the areal impact of the increased utilisation is shown. It is worth noting that the largest impact is in the Eemskanaal area. In terms of gas production from the Eemskanaal cluster this represents an increase from 3 to 3.5 bcm for gas-year 2018/19.

While the hazard metrics show a slight decrease, risk does not behave similarly. In fact, the increased hazard in the South-West results in more houses exceeding the norm, see Figure 6-17.

7 Production Distributions as per "Bouwstenen" document

In the document "Bouwstenen voor Operationale Strategie Groningenveld 2018/2019" (here referred to as the "Bouwstenen document", Reference [22], NAM outlined two possible production strategies ("Inzetstrategiën") resulting in different volume distributions across the regions of the field. In summary, these strategies are:

- **"Inzetstrategie 1"**
  This strategy uses an ordered start-up list with higher priority given to those clusters in the less seismically active areas. This strategy results for an average year in an approximate equal volume offtake from the three main regions; South-West, South-East and East-Central. Bierum contributes by continuous production (approximately 15% of field total) and Eemskanaal only as capacity provider.

- **"Inzetstrategie 2"**
  This strategy is close to how Groningen has been produced in previous years, with a fixed volume distribution across regions. This is irrespective of the total annual production volume (or temperature scenario), resulting in a distribution with the lowest achievable regional fluctuation exceedance. The Bierum cluster is to be produced as per strategy 1 and the Eemskanaal cluster is to be produced at a maximum share of 5%.

Although the operational production strategies defined in the "Bouwstenen document" only apply to gas year 2018/2019, it is possible to analyse the associated behaviour in terms of hazard and risk when assuming that the same strategy is maintained for subsequent years. An exact like-for-like comparison is however not feasible. In the "Bouwstenen document" a detailed analysis is done to establish an expectation for daily demand fluctuations in gas year 2018/2019, based on the actual temperature profiles from the past 31 years. Such resolution is not feasible in the HRA analysis, which only evaluates hazard and risk on an annual basis based on a monthly demand profile based on a single temperature profile (average temperature used for this study). Also, the minimisation of production fluctuations was not part of this study. With these limitations, the following comparison was made:

- "Inzetstrategie 1" is based on a start-up list[10] derived from insights from the 2017 optimisation study, Reference [4]. To evaluate for hazard and risk the associated volume distribution mentioned above (for gas year 2018/2019) was reflected in new optimisation case "**Scenario 1**". This new scenario is then evaluated in terms of hazard and risk.

- The annual volume distribution of "Inzetstrategie 2" is close to the scenario evaluated in chapter 6 called "**17Actual**", that has fixed production fractions over time, as shown in Figure 7-1.

[10] The start-up list approach is described and applied in chapter 6.

Region	Control	Scenario 1	EZK Reference Case
North	[BIR]	6 mln m3/d	6 mln m3/d
Eemskanaal	[EKL]	2 mln m3/d	2 mln m3/d
South-East	[EKR_SZW_ZPD]	1	1
East-Central (1)	[OWG_SCB]	4	2
East-Central (2)	[AMR_TJM_SDB]	5	4
South-West (1)	[ZVN_SPI]	2	3
South-West (2)	[SAP_TUS]	3	5
South-West (3)	[FRB_KPD_SLO]	6	6

Table 7-1 *Modified production start-up list for scenario 1, reflecting strategy 1 from the Bouwstenen document.*

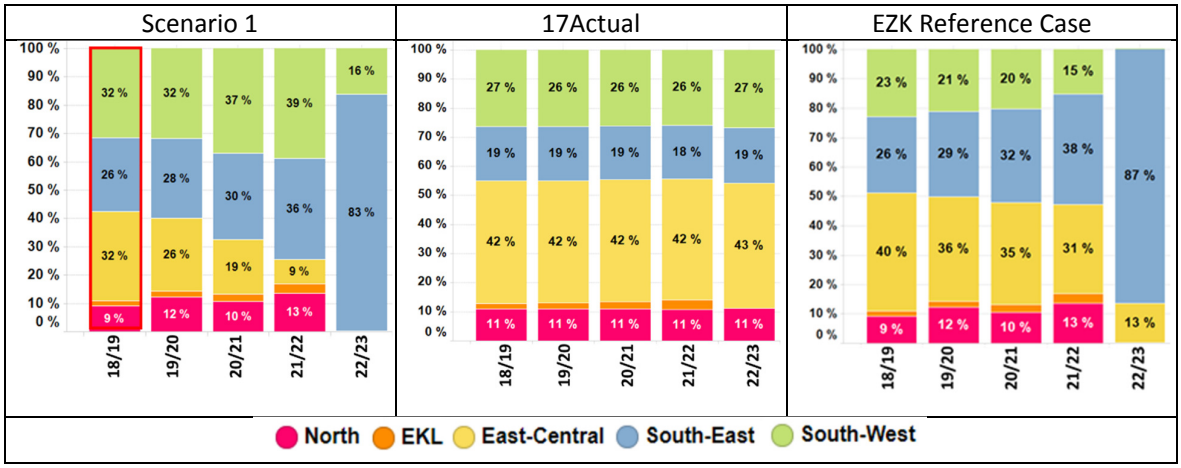

Figure 7-1 *Relative regional production volumes for Scenario 1, reflecting production strategy 1 from "Bouwstenen" document, 17Actual reflecting strategy 2 and the EZK reference case. The resulting regional volumes for gas year 2018/19, marked in red for scenario 1, correspond to volume distributions reported in the "Bouwstenen" document for strategy 1.*

The resulting relative region production volume over time for an average year is shown in Figure 7-1. As mentioned in section 6.4, South-East production region is largely capacity constrained for both scenarios, at a share of about 26% of field total (in gas year 2018/2019). In the EZK scenario the relative contribution of East-Central is higher when compared to Scenario 1, and for South-West vice versa. This reflects insights obtained from this study. Note that the resulting volume distribution for the gas year 2018/2019 for Scenario 1 (red box Figure 7-1-left) is nearly identical to the volume distribution resulting from an average year under strategy 1 in the "Bouwstenen" document. Over time, as production volume decreases, the contribution from the lower priority clusters decrease, evident in the production share from East-Central.

The resulting hazard metrics are shown in Table 7-2 with corresponding risk metrics in Table 7-3 and Figure 7-2. The areal distribution of properties, going from depletion to risk, are shown in Figure 7-3 to Figure 7-7. It is worth noting that Scenario 1 is very similar to the population weighted PGV case from chapter 6. The total number of buildings with mean LPR greater than 1E-5 is similar with 1037 building versus 1033 for the pwPGV case. The areal difference between the scenarios is shown in Figure 7-8.

While risk has not been evaluated for the case **17Actual**, corresponding to "Inzetstrategie 2", the hazard metrics are the worst of those evaluated, as evident from Table 7-2.

In general, it can be concluded that Strategy 1 defined in the "Bouwstenen" document is reducing seismic risk similarly to the pwPGV optimised scenarios described in this study. Based on the resulting volume distribution, the impact on hazard and risk should be an improvement over the EZK reference case.

	Events	maxPGA	maxPGV	pwPGV	Events	maxPGA	maxPGV	pwPGV
EZK Reference	**61**	**0.160**	**0.127**	**0.041**				
Scenario 1	**59**	**0.155**	**0.126**	**0.040**	**-3%**	**-3%**	**-1%**	**0%**
17Actual	**63**	**0.162**	**0.131**	**0.042**	**4%**	**1%**	**4%**	**3%**

Table 7-2 *Hazard metrics for Scenario 1, that is a reflection of Stategy 1 of the "Bouwstenen" document, and 17Actual, that is a reflection of Strategy 2. Percentage changes are w.r.t. the EZK reference case.*

	Total	Improvement	Common	Added	Removed
EZK Reference	**1136**				
Scenario 1	**1037**	**-9%**	**1030**	**7**	**106**

Table 7-3 *Buildings with Mean LPR ≥ 1E-5. **Common** buildings are shared between scenarios, **Added** are those exclusive to Scenario 1 and **Removed** are those exclusive to the EZK reference case.*

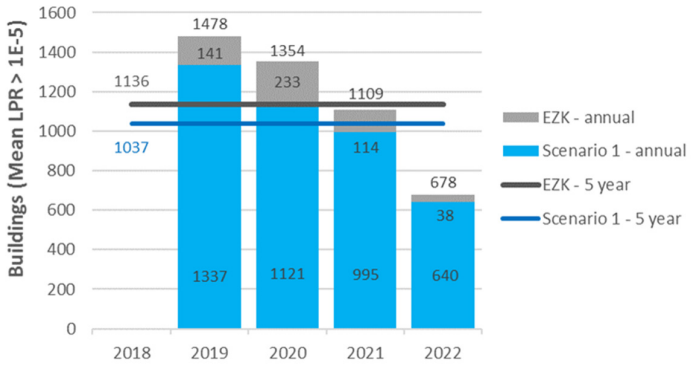

Figure 7-2 *Buildings with Mean LPR ≥ 1E-5, comparing Scenario 1 to the EZK reference case.*

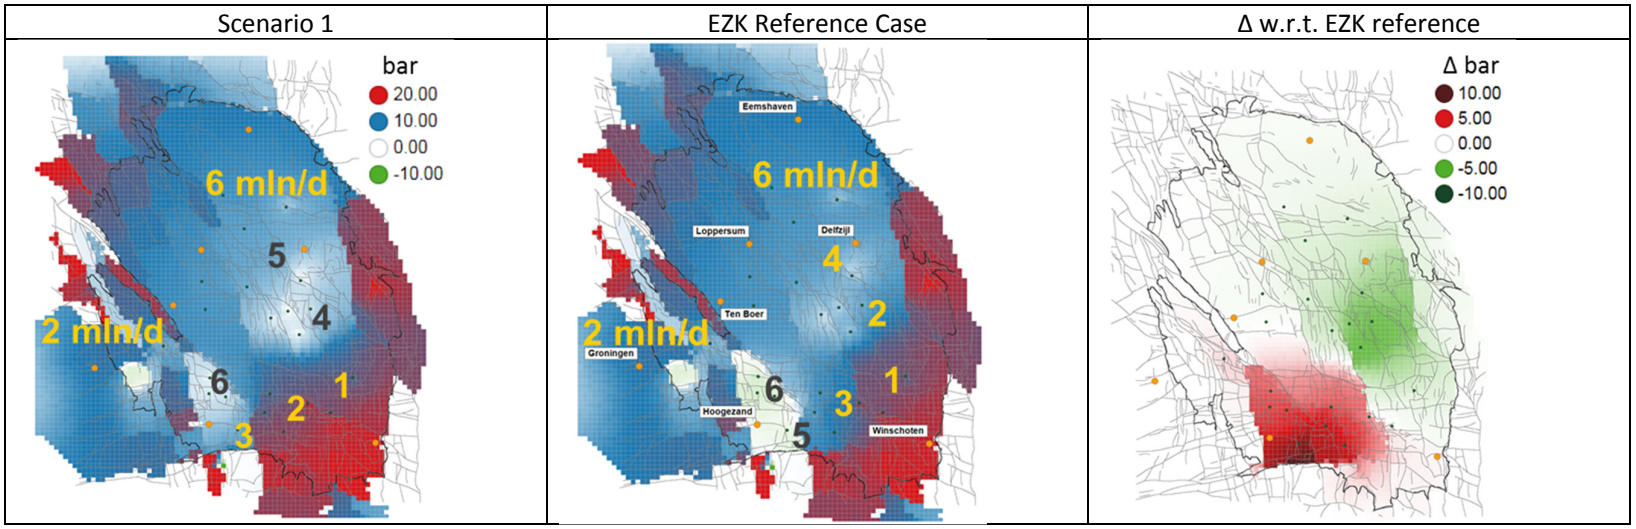

Figure 7-3 Pressure depletion over 4 years from gas year 2018/19 until 2021/22 for scenario 1, with comparison to EZK reference case.

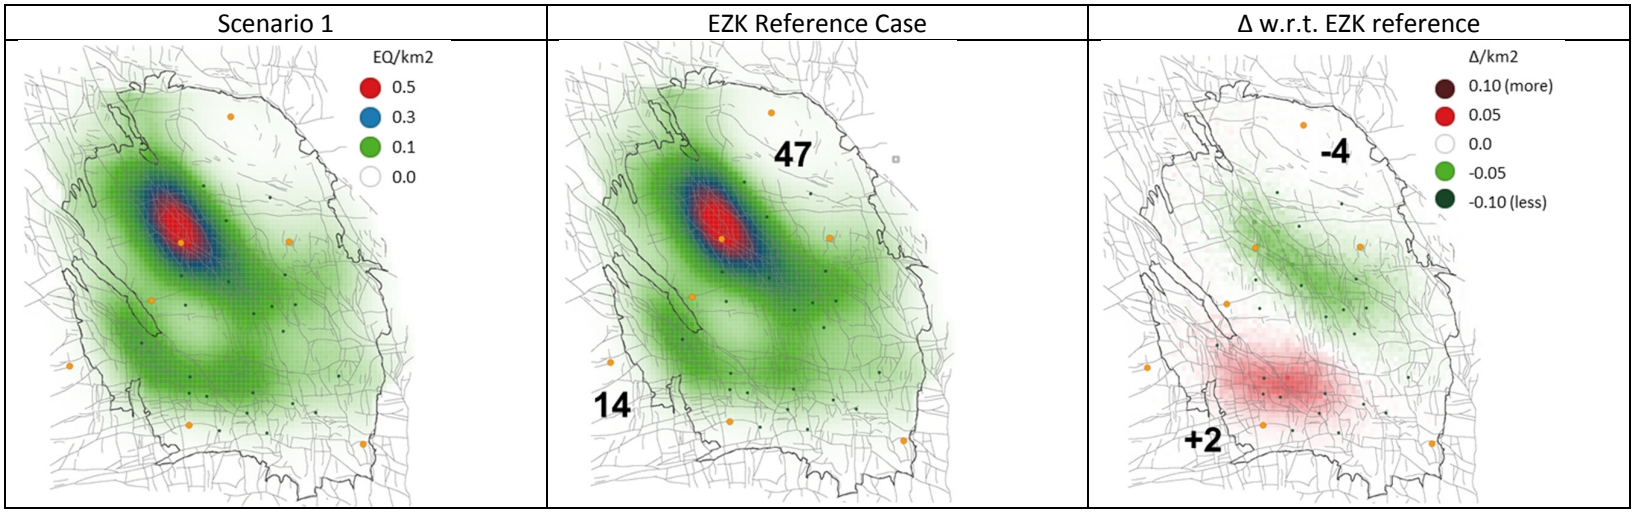

Figure 7-4 Earthquake density over 4 years from gas year 2018/19 until 2021/22 for scenario 1, with comparison to EZK reference case.

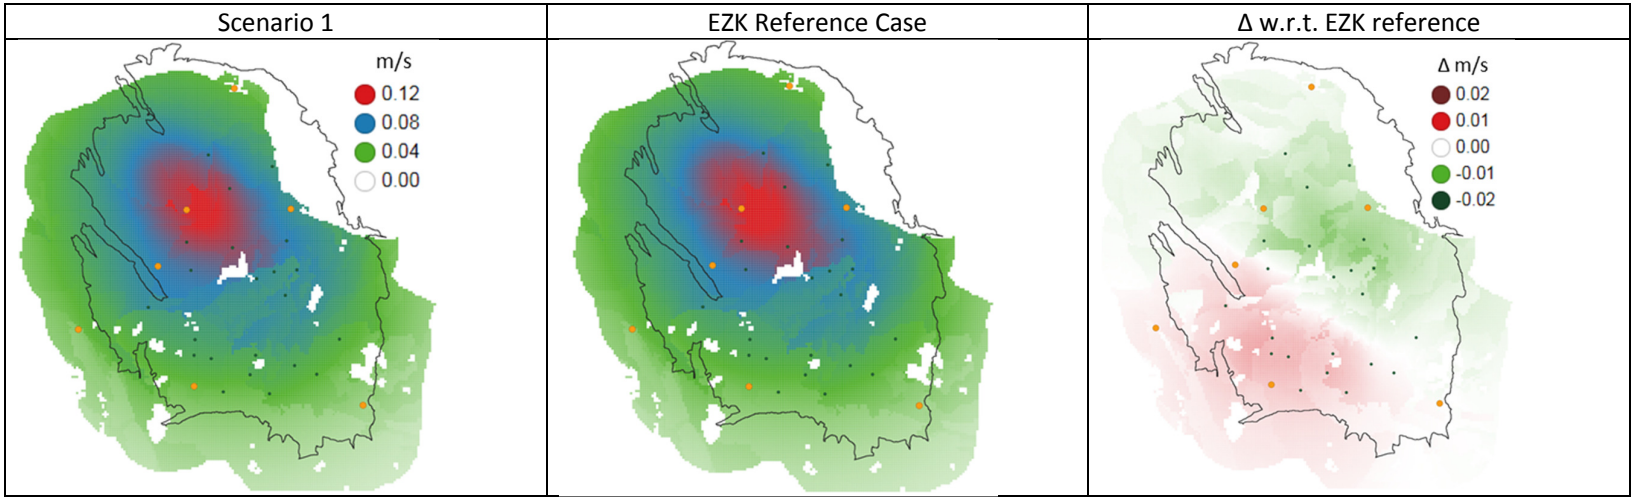

Figure 7-5 PGV over 4 years from gas year 2018/19 until 2021/22 for scenario 1, with comparison to EZK reference case.

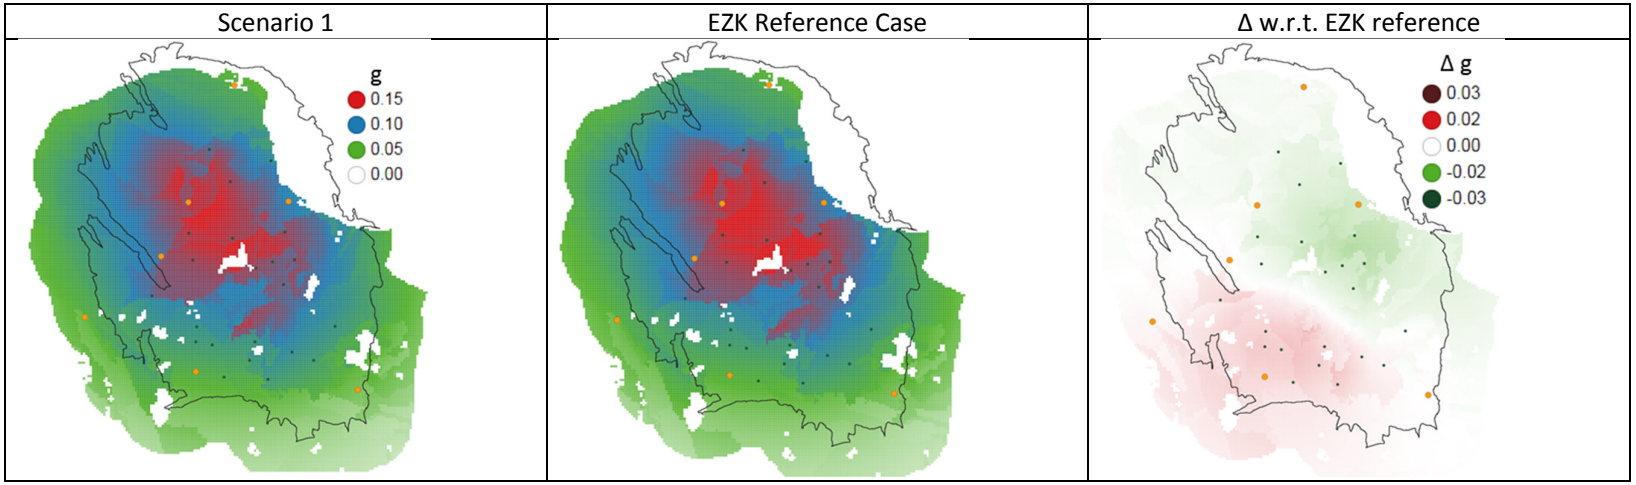

Figure 7-6 PGA over 4 years from gas year 2018/19 until 2021/22 for scenario 1, with comparison to EZK reference case.

Above norm in both scenarios (1030)	Above the norm in Scenario 1 (7)	Above the norm in EZK Reference (106)
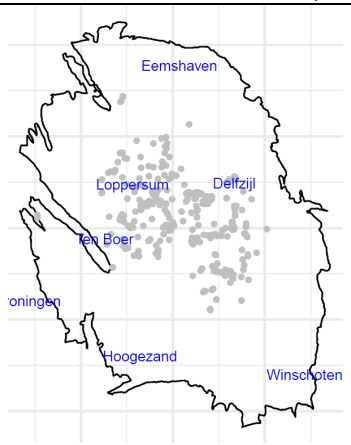	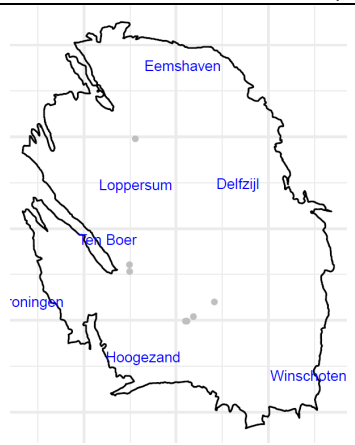	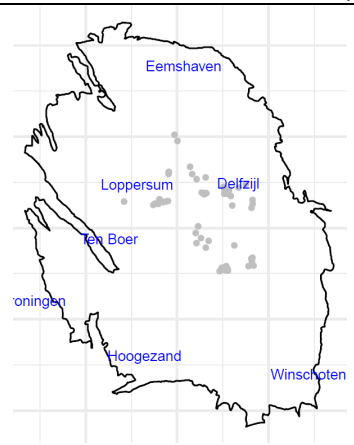

Figure 7-7 Building above the norm 1E-5 (Mean LPR) for scenario 1, with comparison to EZK reference case.

Above norm in both scenarios (1028)	Above the norm in Scenario 1 (9)	Above the norm in pwPGV case (5)
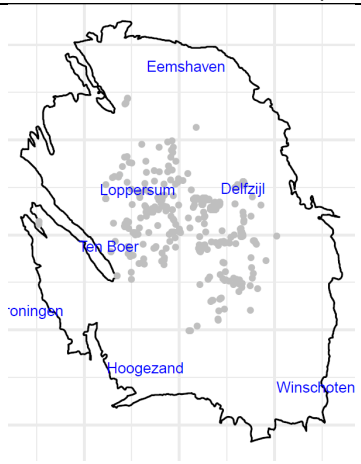	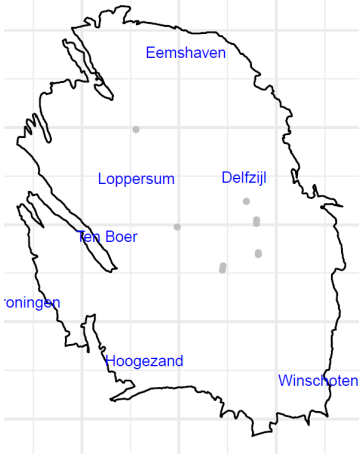	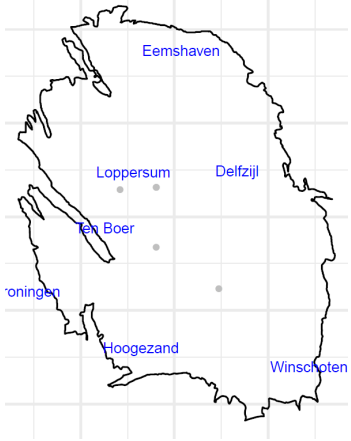

Figure 7-8 Building above the norm 1E-5 (Mean LPR) for scenario 1, with comparison to the optimised pwPGV case.

8 Concluding remarks

On the instruction of SodM, NAM has investigated the optimum production distribution for various hazard and risk metrics. Following the updated outlook on production (Basispad Kabinet), the decision on the distribution of production will be taken by the Minister of EZK. Therefore, the aim of this report is to establish optimal production distributions for several optimisation objectives, and the resulting hazard and risk associated with these.

This report presents the options for the optimisation of the distribution of the gas production over the Groningen field, based on scientific research. It includes the impact each change has for the population in Groningen as a whole and the regional differences. We have shown that seismicity in the 2018-2022 timespan can be reduced by redistributing production over de Groningen field. Dependent on the metric to be optimised (nuisance, hazard, risk), different optimised distributions result. Taking account of operational limitations, risk reduction or nuisance reduction up to 10% is possible relative to the reference case (WP2016 Addendum for Basispad Kabinet, June 2018). However, this represents a choice that needs to be made by the Minister of EZK.

The V5 version of the HRA model has improved the spatial resolution of modelled seismicity. This update with respect to the previous V4 model version yields a systematic shift of the hazard from West to East. Consequently, the 2018 optimisation yields different areal distributions as compared to the 2017 study for the various optimisation objectives investigated. The Eemskanaal and South-West areas are now less hazard prone, whereas the Eastern clusters are more prone to hazard. The combined effect is that the optimisation favours more offtake from the South as compared to the 2017 results.

Despite the update in the underlying models, the analysis presented in the report “Bouwstenen voor Operationale Strategie Groningenveld 2018/2019” was found to be robust. In terms of hazard, “Inzetstrategie 1” aligns roughly with the optimisation for population-weighted PGV as established in this study. Still greater reductions in hazard metrics can be achieved (albeit with spatial differences) when optimising production distribution for other objective functions. Detailed analysis (production fluctuations and temperature sensitivity) as done in Bouwstenen document for gas year 2018/2019, has not been performed for the various optimised scenarios.

9 References

[1] “Technical Addendum to the Winningsplan Groningen 2016 - Production, Subsidence, Induced Earthquakes and Seismic Hazard and Risk Assessment in the Groningen Field, PART I – Summary and Production,” NAM, Assen, 1/4/2016.

[2] Minister of Economic Affairs, “Instemmingsbesluit winningsplan Groningenveld,” 30/9/2016.

[3] Minister of Economic Affairs, “Wijziging instemmingsbesluit winningsplan Groningenveld,” 24/5/2017.

[4] L. Geurtsen and P. Valvatne, “Optimisation of the distribution of production over the Groningen field to reduce Seismicity,” December 2017.

[5] SodM, “Beoordeling "Optimisation of the Production Distribution over the Groningen field to reduce Seismicity",” 30/11/2017.

[6] SodM, “Reactie SodM op plan van aanpak "Optimisation of the Production Distribution over the Groningen field to reduce Seismicity",” 6/4/2017.

[7] SodM, “Beoordeling toepasbaarheid resultaten rapport "Optimisation of the Production Distribution over the Groningen field to reduce Seismicity",” 20/2/2018.

[8] Minister of Economic Affairs, “Letter to Parliament: Gaswinning Groningen”,” 29/3/2018.

[9] J. Van Elk and D. Doornhof, Assessment of Hazard, Building Damage and Risk, Assen: NAM, November 2017.

[10] H. Van Oeveren, P. Valvatne and L. Geurtsen, “Groningen Dynamic Model Updates 2017,” NAM, Assen, 2017.

[11] S. Bourne and S. Oates, Extreme threshold failures within a heterogeneous elastic thin-sheet and the spatial-temporal development of induced seismicity within the Groningen gas field, Journal of Geophysical Research: Solid Earth 122, 10,299--10, 2017.

[12] S. Bourne, S. Oates and J. Van Elk, The exponential rise of induced seismicity with increasing stress levels in the Groningen gas field and its implications for controlling seismic risk, Geophysical Journal International 213, 1693—1700, 2018.

[13] J. Van Elk, A. Mar-Or, L. Geurtsen and P. Valvatne, “Seismic Risk Assessment for Production Scenario “Basispad Kabinet” for the Groningen field,” NAM, Assen, June 2018.

[14] Ministry of Economic Affairs, “Letter to NAM: Verwachtingenbrief aanvulling winningsplan Groningenveld 2016,” 2/5/2018.

[15] SodM, “Advies Groningen-gasveld n.a.v. aardbeving Zeerijp van 8 januari 2018,” 1/2/2018.

[16] D. Wald, V. Quitoriano, T. Heaton and H. Kanamori, Relationships between peak ground acceleration, peak ground velocity, and modified Mercalli intensity in California, Earthquake Spectra 15, 557-564., 1999.

[17] D. Wald, K. Jaiswal, K. Marano, D. Bausch and M. Hearne, PAGER — Rapid assessment of an earthquake’s impact., U.S. Geological Survey Fact Sheet 2010–3036, 2010.

[18] T. I. Allen, D. Wald, P. S. Earle, K. D. Marano, A. Hotovec, K. Lin and M. Hearne, An Atlas of ShakeMaps and population exposure catalog for earthquake loss modeling., Bull. Earthq. Eng., v. 7, DOI: 10.1007/s10518-009-9120-y., 2009.

[19] Q. De Zeeuw and L. Geurtsen, “Groningen Dynamic Model Update 2017 – V5,” NAM, Assen, September 2017.

[20] “Seismic Risk Assessment for a selection of Gas Production Scenarios for the Groningen field,” NAM, Assen, March 2018.

[21] J. Uilenreef, J. van Elk and A. Mar-Or, “Exposure Database (EDB) voor het gebied van het Groningen veld -- Stand van Zaken september 2018,” 2018.

[22] “Bouwstenen voor Operationele Strategie Groningenveld 2018/2019,” NAM, 15/06/2018.

[23] T. Mitchell, Machine Learning, McGraw Hill. ISBN 0-07-042807-7., 1997.

[24] T. Hastie, R. Tibshirani and J. Friedman, The Elements of Statistical Learning: Data Mining, Inference and Prediction, Springer. ISBN 978-0-387-84857-0, 2008.

[25] L. Breiman, “Random Forests,” January 2001. [Online]. Available: https://www.stat.berkeley.edu/~breiman/randomforest2001.pdf.

[26] A. Liaw, “The randomForest package for R - Breiman and Cutler's Random Forests for Classification and Regression,” 7th October 2015. [Online]. Available: https://cran.r-project.org/web/packages/randomForest/randomForest.pdf.

[27] M. Kursa and W. Rudnicki, “Feature Selection with The Boruta Package,” *Journal of Statistical Software,* vol. 36, no. 11, 2010.

[28] J. Friedman, “Greedy Function Approximation: A Gradient Boosting Machine,” *Annals of Statistics,,* vol. 29, pp. 1189-1232, 2001.

[29] L. Devroye, “Non-Uniform Random Variate Generation,” Springer-Verlag, New-York, 1986.

[30] S. Bourne and S. Oates, “Extreme threshold failures account for earthquakes induced by fluid extraction,” Shell, 15/2/2017.

[31] “Assessment of Hazard, Building Damage and Risk,” NAM, Assen, November 2017.

[32] S. Bourne and S. Oates, “An activity rate model of seismicity induced by reservoir compaction and fault reactivation in the Groningen gas field,” Shell, 22 June 2015.

[33] S. Bourne and S. Oates, “An activity rate model of induced seismicity within the Groningen field,” 30 September 2014.

[34] “Bouwstenen voor Operationele Strategie Groningenveld 2018/2019,” NAM, Assen, 15/6/2018.

[35] H. Van Oeveren, P. Valvatne and L. Geurtsen, “Groningen dynamic model updates 2017,” NAM, Assen, September 2017.

[36] K. Jaiswal and D. Wald, “An Empirical Model for Global Earthquake Fatality Estimation,” 2010.

[37] “Groningen Meet- en Regelprotocol,” NAM, 29/5/2017.

Appendix A Model analysis using Machine Learning

A.1 What is meant by Machine Learning

In Reference [23] a popular definition of a Machine Learning algorithm is given. It is said to:

> *"learn from experience E with respect to some class of tasks T and performance measure M if its performance at tasks in T, as measured by M, improves with experience E".*

In the context of this report, this would translate as:

experience E	control response combinations from the experimental design,
task T	predicting the output for the objective of choice (population weighted PGV, max PGV, max PGA, number of tremors, or hybrid thereof)
performance measure M	error metric, such as the Mean Square prediction Error (MSE).

A.2 Measuring the Predictive Performance of a Model

Mathematical Notation and Definitions:

The mathematical notations and abbreviations used are as follows:

n	number of controls, i.e. the number of cluster groups that can contribute to the total production ($n > 0$)
p	number of available control response data points ($p > 0$)
c_{ij}	fractional contribution of cluster i for branch j to the normalised production of 1, with $1 \leq i \leq p$ and $1 \leq j \leq n$; $c_{ij} \in [0,1]$

Furthermore, we denote a response of the Mores-HRA coupling (e.g. pwPGV) by:

$$r_i \in \mathbb{R}$$

if the corresponding controls $c_i \in [0,1]^n$ are used as input. Note that for the sake of simplicity the actual units are omitted since they should be clear from the context.

To abbreviate our notation, we will refer to the i-th data point, consisting of control settings and a Mores-HRA response, as:

$$x_i = (c_i, r_i) \in [0,1]^n \times \mathbb{R}, \qquad \text{with } 1 \leq i \leq p.$$

We introduce the following notations :

$$R = (r_i, \dots, r_p), \text{ and } X = (x_1, \dots, x_p).$$

Let $T \subseteq X$ be a subset of data points in X. We denote by $|T|$ the number of points in the set. Furthermore, we denote a proxy model that has been trained/calibrated on the points in X that maps the controls to a response of the Mores-HRA coupling by

$$f_T : [0,1]^n \mapsto \mathbb{R}$$

For a control point $c_i \in [0,1]^n$ we denote the response of the proxy model, that was trained on T, by:

$$\tilde{r}_i = f_T(c_i).$$

In analogy to the previous abbreviations, we let

$$\tilde{R} = (\tilde{r}_i, \dots, \tilde{r}_p).$$

Finally, we call a function that allows to compare true Mores-HRA responses with predicted Mores-HRA responses (through a proxy model) as an error metric:

$$m: \mathbb{R}^p \times \mathbb{R}^p \mapsto \mathbb{R}$$

General Remarks:

In order to calibrate the proxy model, it is "trained" on a subset of X, which we will call the training set:

$$T \subseteq X.$$

The performance of the model is measured on a separate subset of X, the so-called validation set:

$$V \subseteq X.$$

When measuring the predictive performance of any model in an error metric m, it is important to have a methodology in place that can estimate the performance of the model on as-of-yet unseen data.

If the performance of a model is assessed *in sample*, there is overlap between the training and test data sets, such that:

$$T \cap V \neq \emptyset,$$

For an in-sample assessment, the error metrics are to a certain extent only statements about how well the existing training data can be reproduced.

For models that provide sufficient degrees of freedom in their internal representations, there is always the danger of over-fitting to T. In general, with increasing model flexibility and complexity the prediction error measured in sample tends to get smaller. However, this does not imply that the generalisation error of the model on data that was not used to train the model (such as points in $X \setminus T$) will decrease as well. Hence, one commonly accepted way of measuring and reporting the predictive power of a model is to test it only on unseen data. This is termed *out of sample validation*.

Practically speaking, for out of sample validation the data set X is partitioned into training and validation data sets T and V, normally in a ratio of around 80:20 or 90:10. The model is then trained on the respective training data set T and used to predict the data in the corresponding validation data set V. The performance metrics are then computed on the known responses of the validation data (from the Mores-HRA coupling) against the prediction by the model f_T. Additionally, uncertainties in the error measure itself are quantified and reported in terms of the standard error, which is the associated standard deviation.

It is good practice to not only randomly partition the data set once in training and validation set, but to repeat this procedure several times. A common scheme for subdividing the data into equally sized partitions is *cross-validation*, which is described in more detail in the following.

Estimating the Generalisation Error and Associated Standard Error via Cross Validation:

The general objective, as outlined above, is to compute the generalisation error and associated confidence interval of a model f_T on unseen data V in the chosen error metric(s). It is important to realise that the choice of the partition in which the given data is split into training and test set will have an impact on both the trained model f_T and on the error metrics computed on the test set. In order to reduce this effect related to the partitioning of the input data, a common strategy is to run several model training and prediction experiments on different partitions of the data and to report the mean of the error metrics and the associated standard errors. If the experiment is performed l times, the uncertainty in the estimates is reduced by a factor of $\frac{1}{\sqrt{l}}$.

A common partitioning technique is to split the given data into l (almost) equally sized bins which contain $\left\lceil \frac{p}{l} \right\rceil$ elements, except potentially the last bin which may have less elements. We abbreviate the data points in the i-th bin with B_i. Then consequently l training and test sets are formed where the i-th test set is given by B_i and the i-th training set is given by $X \setminus B_i$. In total l models are trained and evaluated. In the end, the mean of the individual error metrics is reported and the associated standard error is computed. This procedure is known as *cross-validation*. A nice property of this scheme is that it ensures that every point in X is predicted exactly once out of sample. A schematic overview of 5-fold cross-validation is contained in Figure A-1.

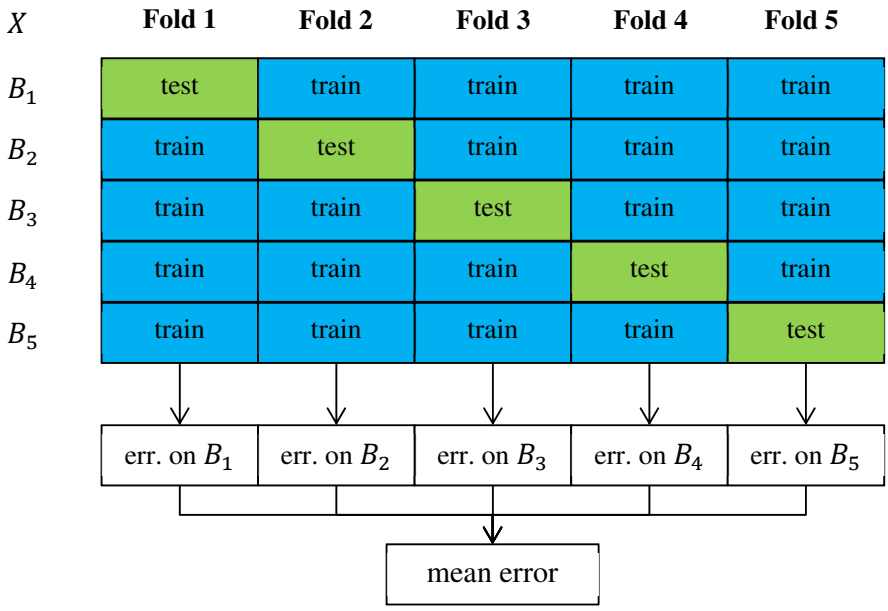

Figure A-1 Stylised example of 5-fold cross-validation

Error Metrics

The following error metrics are used related to the estimation of the predictive performance of the proxy models:

- The Mean Squared Error of a model f_T, with respect to training set T and validation set V, is defined as:

$$\mathrm{MSE}(f_T, V) = \frac{1}{|V|}\sum_{i=1}^{|V|}(r_i - \tilde{r}_i)^2,$$

- The Root Mean Squared error is defined as:

$$\mathrm{RMSE}(f_T, V) = \sqrt{\frac{1}{|V|}\sum_{i=1}^{|V|}(r_i - \tilde{r}_i)^2}\,.$$

Note that in these error measures larger deviations in predictions from the true values get penalised more than smaller deviations.

The associated Standard Errors are given by:

- For MSE:

$$\mathrm{SE}_{\mathrm{MSE}(f_T,V)} = \frac{\mathrm{sd}\left(\left(R - \tilde{R}\right)^2\right)}{\sqrt{|V|}}$$

- For RMSE:

$$\mathrm{SE}_{\mathrm{RMSE}(f_T,V)} \approx \frac{\mathrm{SE}_{\mathrm{MSE}(f_T,V)}}{2 \cdot \mathrm{RMSE}(f_T, V)} = \frac{1}{2}\frac{\mathrm{sd}\left(\left(R - \tilde{R}\right)^2\right)}{\sqrt{\sum_{i=1}^{|V|}(r_i - \tilde{r}_i)^2}}$$

where $\mathrm{sd}(\cdot)$ denotes the usual standard deviation and $\left(R - \tilde{R}\right)^2 = \left((r_1 - \tilde{r}_1)^2, \ldots, \left(r_{|V|} - \tilde{r}_{|V|}\right)^2\right)$.

A.3 CART Trees and Random Forests

A Classification and Regression Tree (CART) is a tree-based ML method. Tree based methods partition the feature space into a set of (high-dimensional) rectangles and fit a simple model for each rectangle, Reference [24]. We mention this method herein because it is important for understanding the Random Forest algorithm that is the primary ML algorithm we use for the analysis of the Mores-HRA coupling.

To construct a CART decision tree, first a root node is created. From the list of features, which in our situation are the cluster offtake splits, an optimal feature for branching and an associated splitting value are determined based on the minimisation of the sum of the square errors between the prediction of the resulting model and the actual value of the prediction target. This branching feature and value is used to partition the training data into a “left” and “right” subset. The left and right

branches correspond to samples that have a value less than or equal to the splitting value and greater than the splitting value, respectively, for the feature upon which the splitting is applied. For each of the branches, a new child node is defined and the previous procedure is repeated, based on the subset of the data that remains in the branches in question. It is generally considered to be sensible to use a binary split, at each branch, primarily since it is easier to implement but also because in some situations not enough data would remain in each of the branches, in the next level down in the tree. The splitting procedure is applied, recursively, until only one data point remains in each of the sub-branches. A terminal node is then defined, such that the value of the prediction target is stored. This is the value that the tree would predict for a new sample that follows the same route through the tree. In practice, for a single CART tree, it would be common to apply pruning, a method in which the terminal nodes contain more than one data point, for which an average would be taken. This avoids the tendency to over-fit to the training data by trading-off bias and variance. Note that the CART trees which are used inside the Random Forest algorithm are not pruned per-se, but they are also not grown to the point at which each termination node is a single data point. An ensemble of trees is used to reduce the chance of over-fitting. An illustrative example of an individual CART tree can be seen in Figure A-2. Internal nodes are represented by white circles and terminal nodes by red circles. Suppose that you have 10 controls and $c_3 = 0.1$, then the prediction for pwPGV from this individual tree is 0.06.

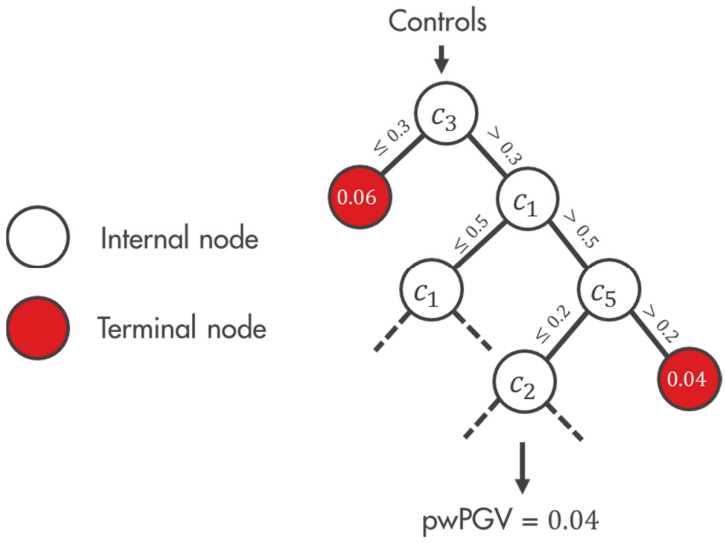

Figure A-2 *Illustrative example of a CART tree that relates controls to a Mores-HRA response*

The Random Forest algorithm uses an ensemble of CART trees. It was originally proposed by Breiman in 2001 and has since been cited more than 30000 times [25]. This illustrates its widespread use and acceptance throughout academia and in many industrial applications. Each of the CART trees is constructed on a random sub-sample of the training data and using only a subset of the features, chosen, again, at random. The prediction of the model is obtained by traversing each individual CART tree with the test data and in the context of regression, taking the mean of the predictions of each individual CART tree. An illustration of a Random Forest model, consisting of several hundred individual CART trees, can be seen in Figure A-3. The main tuning parameters are the number of trees and the number of features that are randomly sampled. The Random Forest is typically seen to be insensitive to hyper-parameter tuning and works well with default settings. It is good at capturing non-

linear relationships in small to medium sized data sets but is considered too computationally expensive for "Big Data". It can handle both categorical and numerical data and is essentially invariant to monotonic transformations of the features. All predictions it makes remain within the range of the observations in the training data. We use the implementation available in the '*randomForest'* package for R, Reference [26]. Note that the programming language R is the de facto standard for statistical computing and contains many implementations that have been checked numerous times for statistical rigor.

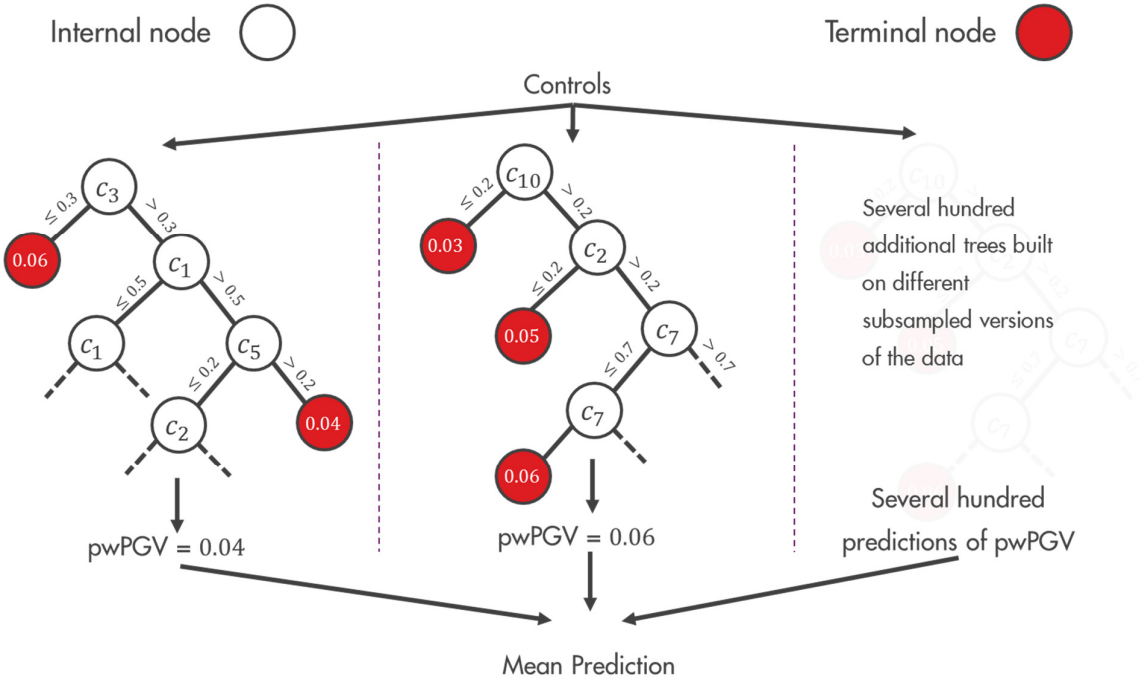

Figure A-3 Illustrative example of a random forest model, consisting of an ensemble of CART trees.

A.4 Relative Variable Importance Analysis

It is not our sole intention to build a proxy model for the relationship between the cluster offtake splits and the prediction target of choice, e.g. population weighted PGV. We note however, that by using the Random Forest algorithm, it is possible to construct a proxy model with an out of sample coefficient of determination of over 0.95. This would generally be considered sufficiently accurate such that a variable importance assessment, performed on the proxy model, would be informative.

Our primary objective is to use the proxy model to understand the relative impact of changing each of the cluster offtake splits on the prediction target in question. In the simplest case, one could consider constructing a Pareto plot by systematically varying each of the cluster offtake splits in a univariate manner, adjusting each of the other offtake rates, proportionally. However, this analysis does not account for the interactions and subsequent impact on reservoir response of a multivariate setup. There may be combinations of cluster offtake splits that result in more extreme behavior. The univariate analysis could result in one developing an overoptimistic view of how well one understands the model response over the entire parameter space.

Some Machine Learning algorithms, such as the Random Forest method, have built-in functionality that allows one to capture the relative importance of features in the multivariate context. One such technique, involves systematically switching the information for a feature between different samples, to remove the information that it contains, Reference [25]. A Random Forest model is first built for the prediction target of choice, using the training data, with performance metrics assessed, out-of-sample using test data (i.e. performing a blind test). For each of the features in turn, the values are permuted and re-run through the trained Random Forest model. For some of the features, there will be a detrimental impact on the accuracy of the predictions, as measured by the mean square prediction error (mse). If the feature in question is important, one would expect a large impact on the performance. If a relatively small impact on the performance is observed, then it can be inferred that the feature is not particularly important. The analysis can be repeated for each of the features in turn.

Whether the estimated importance in terms of the permutation based increase in mse of a feature is statistically significant can be determined by applying a variant of Student's t-test. The procedure is implemented for random forests in the Boruta algorithm and the associated R package. Further details can be found in Reference [27].

A.5 Partial Dependence Analysis

The relative variable importance analysis does not indicate whether an increase in the feature in question would cause an increase or decrease to the prediction target considered. For simple methods like a linear regression model, it is easy to understand the underlying relationship between the features and the prediction target because one can look at the sign and magnitude of the model coefficients. A simple parametric description cannot be examined in a similar fashion for the Random Forest.

Instead we can use the concept of a partial dependence plot, introduced by Friedman in Reference [28], that can be used for any so-called "black-box" style method. A Random Forest model is first built using the training data. The partial dependence analysis is then carried out, based on this model. Note that the partial dependence can be estimated for any subset of variables, such that interaction effects can be better understood. Let C denote the set of n controls on which the model response depends. We consider a subset C_S of C, such that $S \subset \{1, \dots, n\}$, and the complement set of controls is defined as $C_{\bar{S}}$, such that $\bar{S} = \{1, \dots, n\} \setminus S$. For a set of controls C, the response of the Random Forest model is given by:

$$f(C) = f(C_S, C_{\bar{S}}).$$

The partial dependence of the model on the controls specified in S is then defined as:

$$\overline{f}_s(C_S) = \frac{1}{p}\sum_{i=1}^{p} f(C_S, c_{i\bar{S}}),$$

where the $c_{i\bar{S}}$ are the values of the controls $C_{\bar{S}}$ of the i-th point of the training data on which $f(C)$ was trained.

In order to graph the relationship, we evaluate the partial dependence function $\overline{f}_s(C_S)$ on a regular grid spanned between the minimally and maximally observed values of C_s in the training set. Compare

Reference [24] for further details. The procedure can be performed for each of the subsets of the features, in turn. This can be a computationally expensive process but for decision tree based methods, the partial dependence can be rapidly computed, directly from the trees, Reference [24].

Note that this partial dependence analysis provides more than merely the marginal dependence between the feature and prediction target. It represents the average effect of the feature in question on the predictions of the Random Forest model after controlling for the effects of all other features on the Random Forest model.

Appendix B Random Forest Sampling

When creating a "proxy model", it is important that a sufficient number of control and output pairs is available to which the proxy model is fit. To be able to make statements about the overall quality of the proxy model, it is necessary that the proxy model is based on a representative set of control output pairs that capture enough of the dynamic behaviour of the original HRA model.

B.1 Workflow

The control space was uniformly randomly sampled and it was tested whether the out-of-sample prediction performance of the proxy model exceeded a certain threshold that is deemed sufficient for the modelling purposes. This threshold is chosen to be close to the irreducible uncertainty level that is due to the aleatoric uncertainty of the HRA model (section3.4.2). Once this threshold is exceeded, no additional samples need to be generated.

The epistemic uncertainty in the HRA model is captured by the various branches of the uncertainty tree. One proxy model is created for each of the HRA branches. A conceptual overview of the workflow for one individual branch of the HRA is shown in Figure B-1.

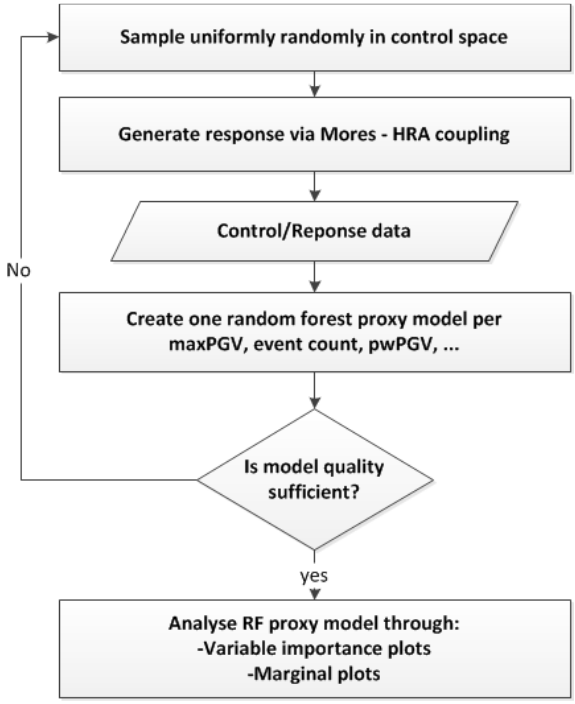

Figure B-1 Sketch of proxy workflow for one HRA branch

B.2 Mathematical Notation and Definitions

The mathematical notations and abbreviations used are as follows:

m	number of HRA branches ($m > 0$)
n	number of controls, i.e. the number of cluster groups that contribute to the total production ($n > 0$)
p_j	number of samples for branch j, with $1 \leq j \leq m$ and $p_j > 0$
c_{ij}	fractional contribution of cluster i for branch j for total normalised production of 1, ($0 \leq c_{ij} \leq 1$)

Furthermore, for $1 \leq j \leq m$ the equality $\sum_{i=1}^{n} c_{ij} = 1$ holds. For the sake of simplicity, in the following, the focus is on one individual branch at a time, and the corresponding branch index is omitted in the notation.

B.3 Uniformly Sampling from a Convex Polyhedron

The total control space is span by all possible combinations of production fractions for the various controls, whereby the sum of the production fractions is 1. Due to this constraint of $\sum_{i=1}^{n} c_i = 1$, the admissible control strategies all lie on a unit $(n-1)$-simplex. A simplex is the generalisation of the concept of triangles and tetrahedra to higher dimensions. In Figure B-2 an illustration is given for a system of 3 controls, where a multitude of admissible combinations make up a triangle.

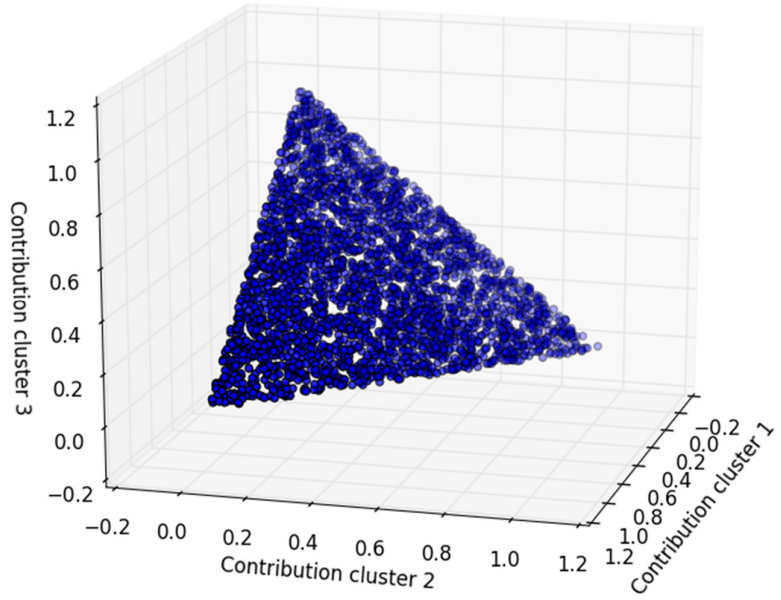

Figure B-2 Admissible combinations for 3 controls lie on 2-simplex

The theory for sampling uniformly from a unit simplex is well established and for instance described in Reference [29], page 568. A rather simple algorithm for drawing uniformly random samples that is also intuitive works as follows:

1. Draw $n-1$ samples uniformly from the interval $[0,1]$. Call the resulting set S.
2. Add the numbers 0 and 1 to the set S.
3. Order the set S in increasing order and denote its ordered elements by s_i with $1 \leq i \leq n+1$.
4. Let $c_i = s_{i+1} - s_i$ for $1 \leq i \leq n$.

One can show that the algorithm described above will create uniformly random samples in the unit $(n-1)$-simplex, see [29] p. 568 Theorem 2.1. The computational cost to create p_j samples for a particular branch j is in $O(p_j \log p_j)$, which is negligible compared to the time it takes to run the associated Mores model in combination with the HRA. A set of 3000 samples that have been generated using this approach is contained in Figure B-2.

Remark:
The strategy of drawing n times uniformly from the interval $[0,1]$ and then normalizing by the sum to obtain c_i such that $\sum_{i=1}^{n} c_i = 1$ and repeating this procedure p_j times will yield samples that are not uniformly distributed in the control space. See Figure B-3 for an illustration. Extreme combinations of parameters are sampled with a lower probability which may mean that the control space is not properly covered. This may be even more of an issue in a high dimensional space, as with $n = 10$, and/or if the model response is complex towards the extreme points of the control space.

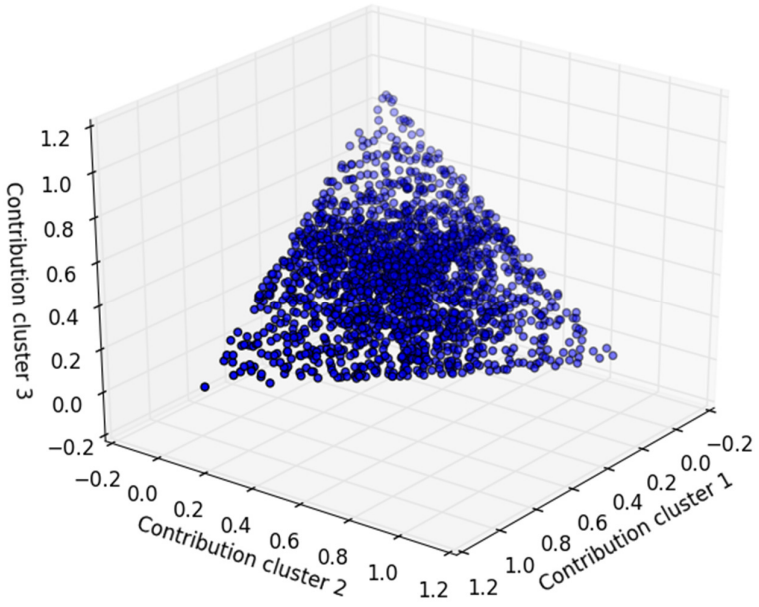

Figure B-3 *Non-uniform sampling of a 2-simplex*

B.4 Sampling strategy for calibration of the Random Forest proxy model

The actual situation at hand is slightly more involved than just sampling from a unit simplex, since some of the controls themselves are subject to boundary conditions, which are based on the maximal contribution that individual production cluster can make over the forecasting horizon for which the contribution of the cluster is kept constant. These upper boundaries have been determined empirically based on a set of several prediction runs. The lower boundaries are in place to ensure the minimal operational constraints of the individual clusters are met.

The following boundary conditions were applied:

$0.01 \leq c_1 \leq 0.25$	$0.01 \leq c_2 \leq 0.25$	$0.01 \leq c_3 \leq 0.51$	$0.01 \leq c_4 \leq 0.31$	$0.01 \leq c_5 \leq 0.38$
$0.01 \leq c_6 \leq 0.35$	$0.01 \leq c_7 \leq 0.16$	$0.01 \leq c_8 \leq 0.30$	$0.01 \leq c_9 \leq 0.36$	$0.01 \leq c_{10} \leq 0.60$

Schematic illustration

The resulting control space that one obtains if those boundary conditions are in place is a convex polyhedron. For illustrative purposes consider the case with 3 controls which have the following boundary conditions:

$$0.2 \leq c_1 \leq 0.60 \qquad 0.2 \leq c_2 \leq 0.60 \qquad 0.10 \leq c_3 \leq 0.90$$

A visualisation of the associated control space is contained in Figure B-4.

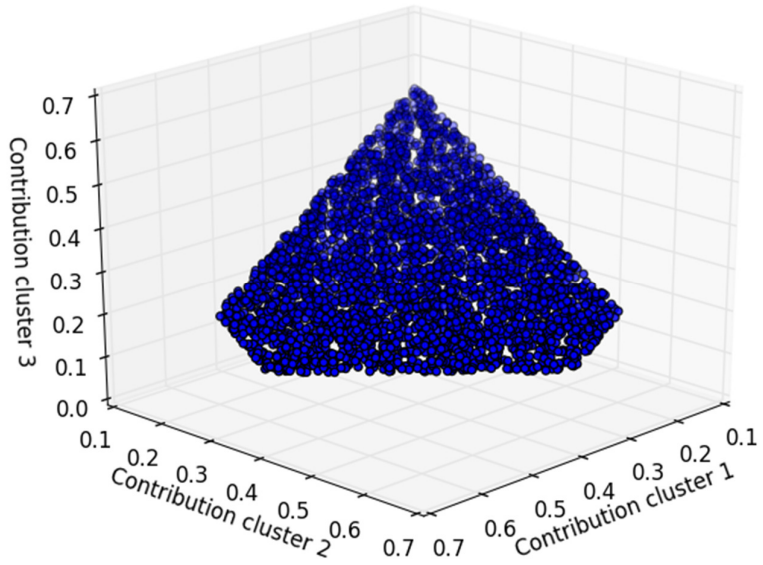

Figure B-4 *Control space for 3 controls with constraints is a convex polyhedron*

Applied sampling strategy

The strategy that is used to sample uniformly from a convex polyhedron is to sample uniformly from the associated unit simplex and to discard all samples which do not satisfy the boundary conditions on the c_i as specified above. This rejection sampling strategy is straightforward to implement. Even though more direct techniques would exist that are computationally less costly, the cost to generate a few 1000 samples for 10 controls is still negligible (in the order of seconds) compared to the cost of running the coupled Mores and HRA models. The sampling depicted in Figure B-4 was obtained using this rejection sampling strategy.

Appendix C Comparison of 2018 to 2017 optimisation study

The results as presented in this 2018 update of the optimisation study were derived using the updated V5 model suite. There are some significant differences in the optimised production fractions when compared to the 2017 optimisation study [4], which was using the V4 model suite. This chapter provides a more detailed investigation of the differences and their causes.

C.1 V4 to V5 differences

The update from the V4 to the V5 version of the model suite (dynamic reservoir model, seismological model, and ground-motion-prediction-equations) is discussed in Chapter 3. The differences in response between both model suites were assessed using the optimisation results of the 2017 study. The associated production rate (21.6 N.Bcm/y) and production distributions across the production clusters were evaluated with the V5 Mores model for the 5 year period from 1/1/2018 to 31/12/2022. Three production scenarios were considered (Figure C-2):

- Reference case (Business Plan 2017)
- 2017 optimisation for Events
- 2017 optimisation for population weighted PGV

The history matched pressure at the start of the period is shown in Figure C-2, while the associated pressure depletion for each scenario is given in Figure C-3 to Figure C-5.

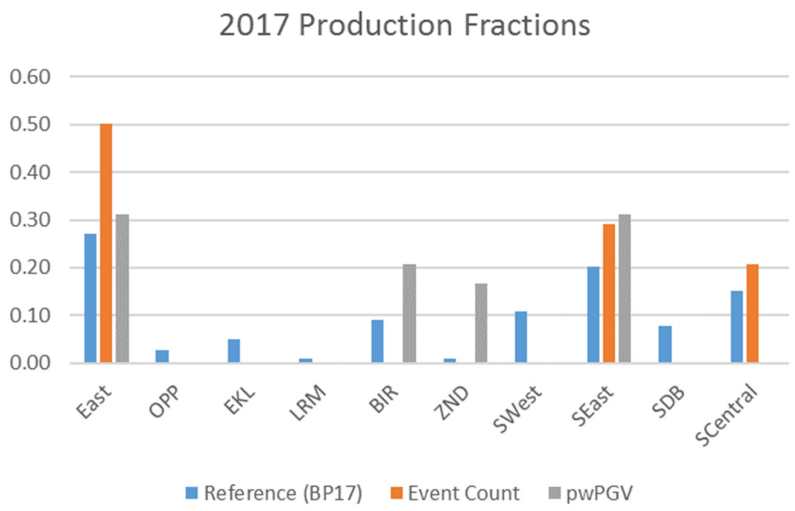

Figure C-1 2017 production distribution across the model controls

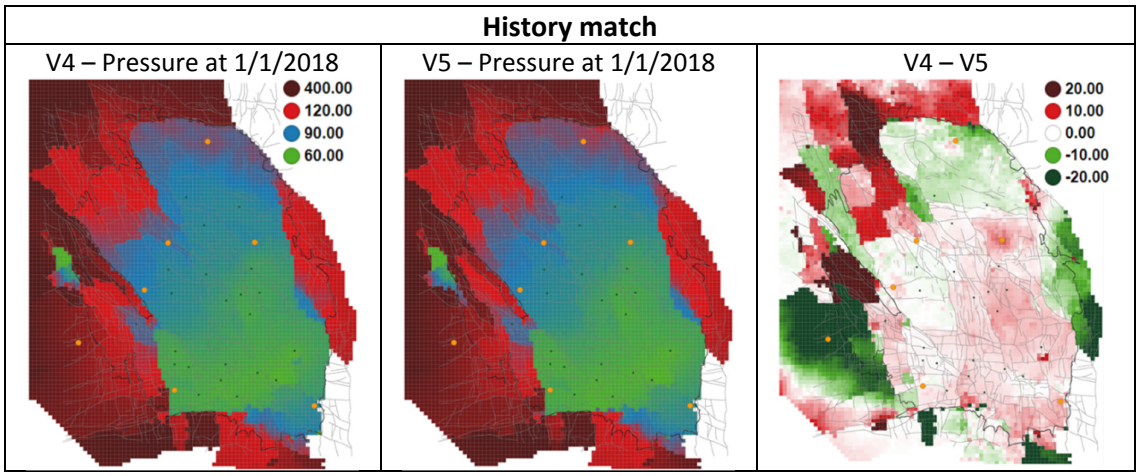

Figure C-2 *V4 versus V5 pressure match at 1/1/2018*

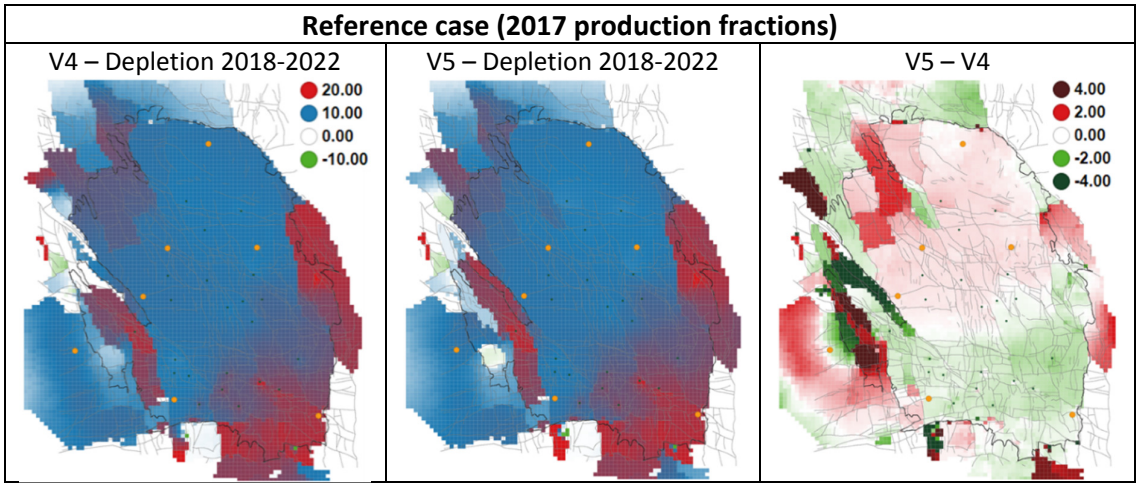

Figure C-3 *Depletion forecast from 1/1/2018 to 31/12/2022 for the Reference scenario.*

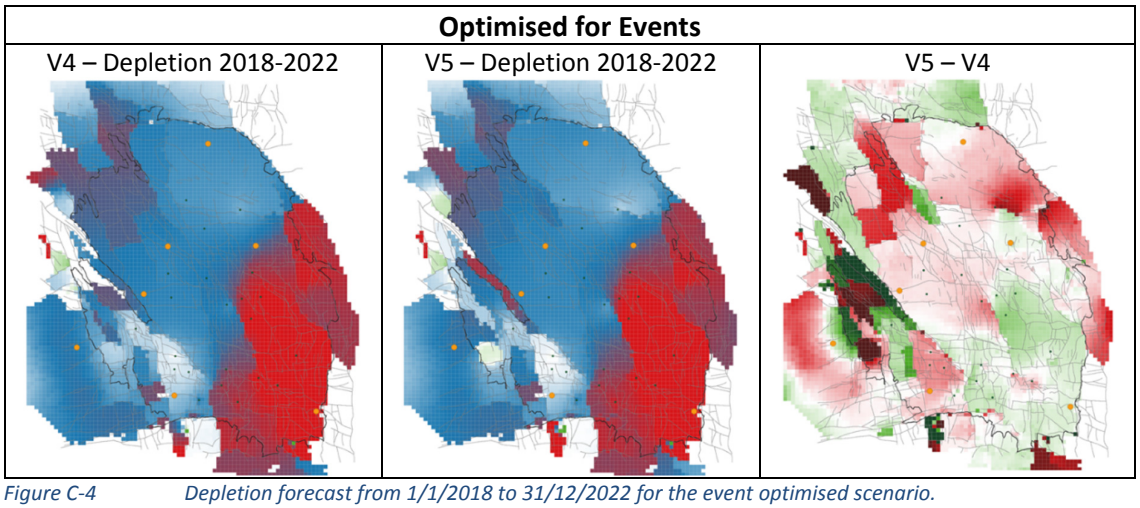

Figure C-4 *Depletion forecast from 1/1/2018 to 31/12/2022 for the event optimised scenario.*

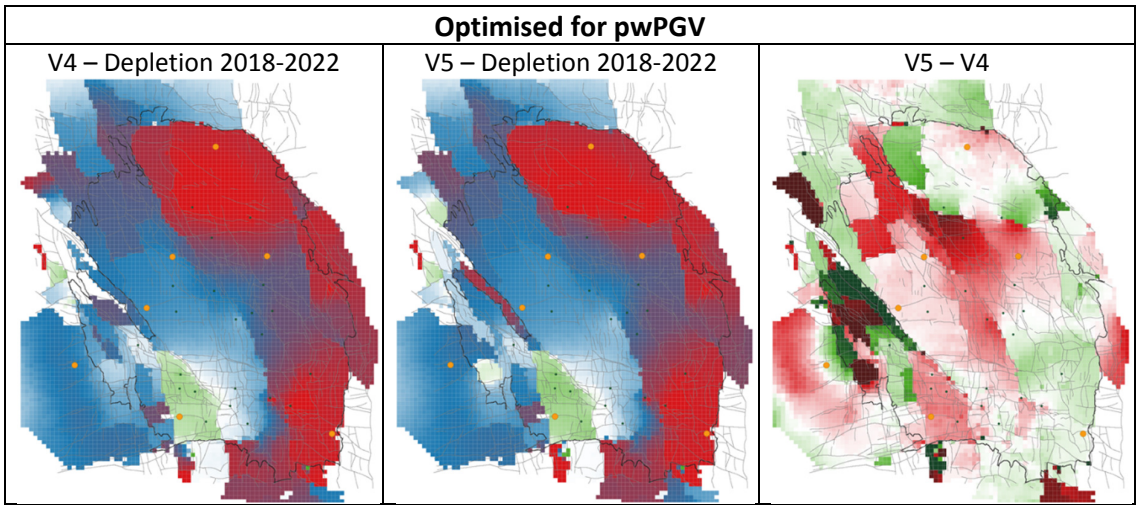

Figure C-5 *Pressure depletion between 1/1/2018 and 31/12/2022 for pwPGV optimised scenario.*

It can be observed from Figure C-2 to Figure C-5 that the largest difference in depletion, both during history match and forecast, is at the edge of the modelled area. This is outside the area of well control, resulting in lack of data to constrain the subsurface model. Within the bulk of the field, where most of the seismicity occurs, the difference between the models is relatively modest, typically within 1 bar.

As a next step, for each pressure depletion scenario, the difference in response between the V4 and V5 model suite can be assessed

The estimated earthquake density (Figure C-6) and hazard (Figure C-7) resulting from the reference case depletion (Figure C-3) shows that the V5 model has a systematic shift in earthquake density and hazard towards the East.

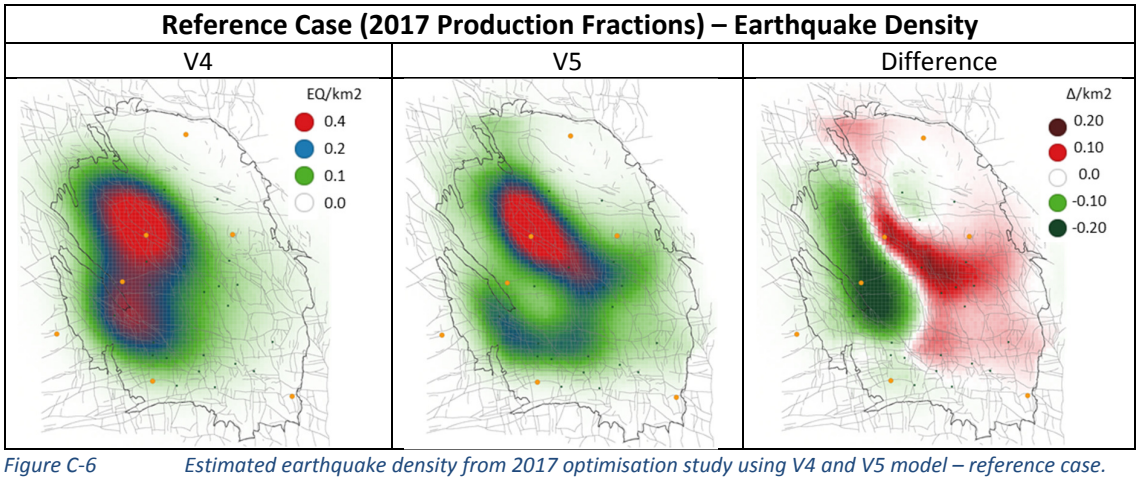

Figure C-6 *Estimated earthquake density from 2017 optimisation study using V4 and V5 model – reference case.*

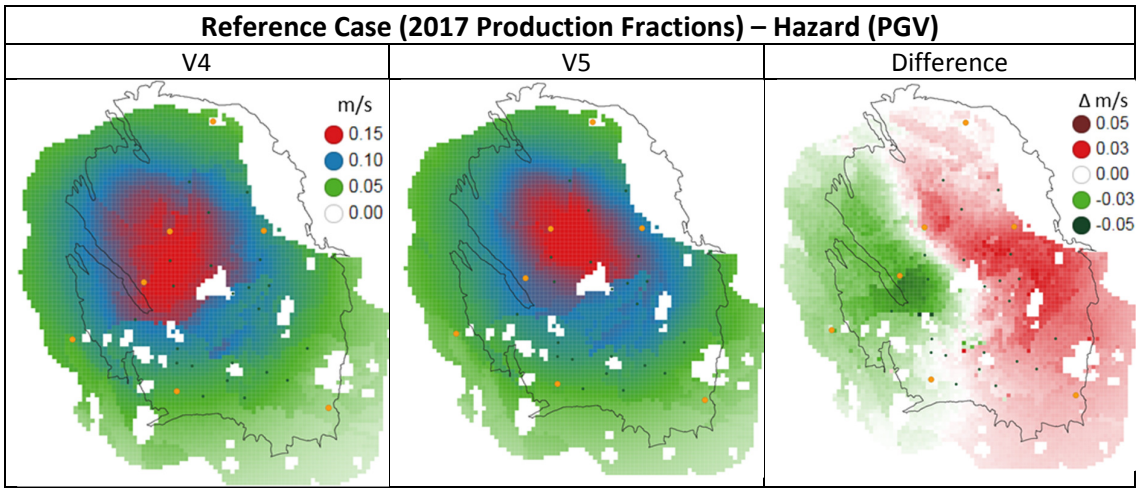

Figure C-7 *Estimated PGV hazard from 2017 optimisation study using V4 and V5 model – reference case.*

When optimising for event rate, production is mainly moved away from the South-West, including Eemskanaal, towards the South-East and East. The V4 model predicts a corresponding reduction in event rate in the West of the field, with only a limited increase in the East, Figure C-8. Using the V5 model, the same production distribution results in less pronounced reduction in the West, while at the same time a larger increase in the East, Figure C-9. The result is that the optimisation gain from the 2017 study is significantly reduced when the same production distribution is re-evaluated with the V5 model. In Table C-1 and Table C-2, the calculated hazard metrics are summarised for each production scenario, respectively for the V4 and V5 model suite. For example, the improvement in event rate has dropped from 15% (V4) down to 4% (V5).

Objective	Hazard Metric (V4)				Improvement (V4)			
	Events	maxPGA	maxPGV	pwPGV	Events	maxPGA	maxPGV	pwPGV
Reference (BP17)	105	0.265	0.166	0.061				
Event Count	89	0.245	0.156	0.054	-15%	-8%	-6%	-11%
pwPGV	89	0.247	0.163	0.053	-16%	-7%	-2%	-13%

Table C-1 *Outcome from 2017 optimisation study (results generated with V4 models).*

Objective	Hazard Metric (V5)				Improvement (V5)			
	Events	maxPGA	maxPGV	pwPGV	Events	maxPGA	maxPGV	pwPGV
Reference (BP17)	96	0.195	0.167	0.055				
Event Count	92	0.200	0.162	0.051	-4%	2%	-3%	-6%
pwPGV	96	0.198	0.178	0.050	-1%	1%	6%	-9%

Table C-2 2017 optimisation outcome re-evaluated with V5 models

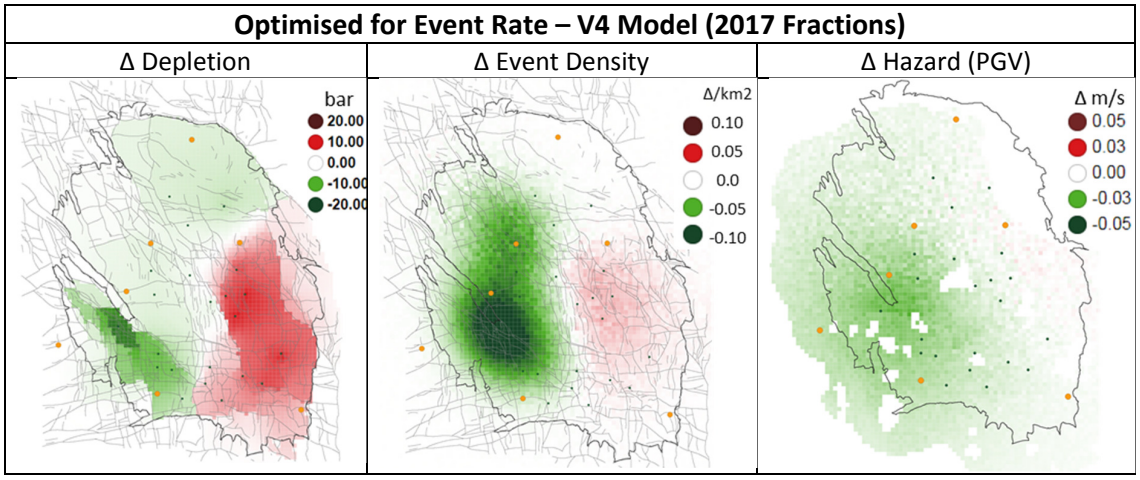

Figure C-8 *Event rate optimised case from 2017 optimisation study – original V4 based results.*

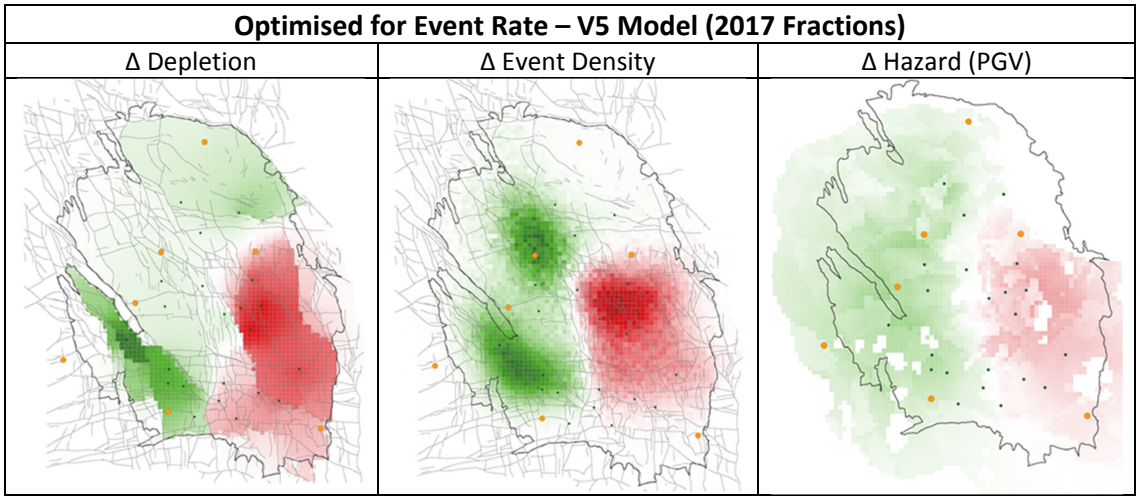

Figure C-9 *Event rate optimised case from 2017 optimisation study – original results re-run with V5 model.*

Overall, a systematic shift in hazard is found between the V4 and V5 model suites from West to East. However, as illustrated in Figure C-10, the shift is well within the epistemic confidence intervals.

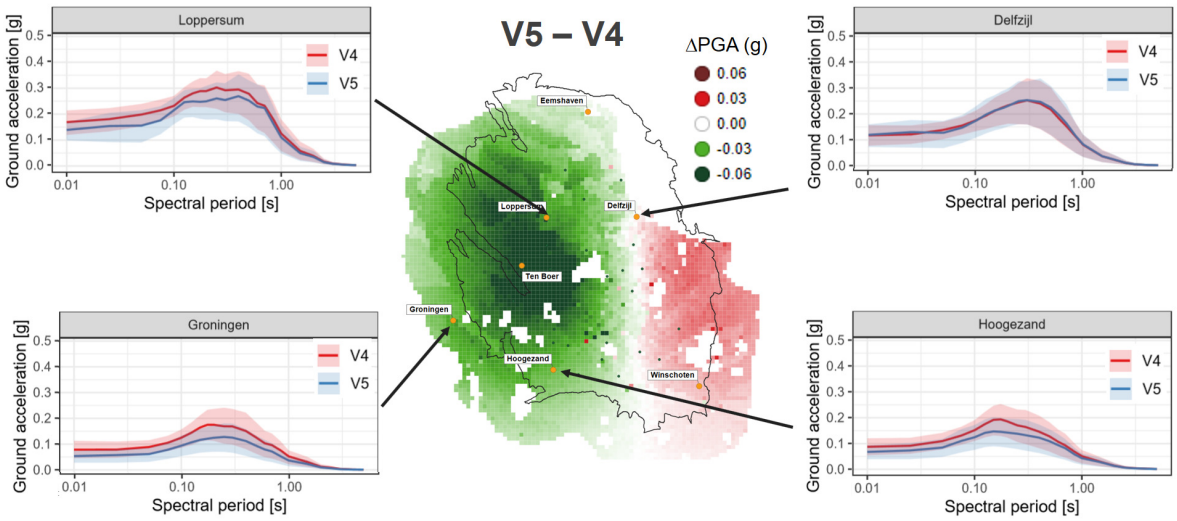

Figure C-10 PGA difference between V4 and V5 model in context of the epistemic uncertainty

C.2 Comparing Random Forest analysis between V4 and V5

The changes between the 2017 study and this update were also investigated for the Random Forest models. Figure C-11 compares the partial dependence summary map charts. It can be observed that the results for the V5 model are fairly similar between the Basispad Kabinet production profile and 21.6 Bcm/y (note that the latter was evaluated using the 2017 model controls). When in turn comparing the 21.6 Bcm/y production profile between V5 and V4, there are significant differences. Most notably there is a polarity flip at the `[EKL]` control, and the relative impact of the `[PAU_POS_OVS]` control is substantially lower.

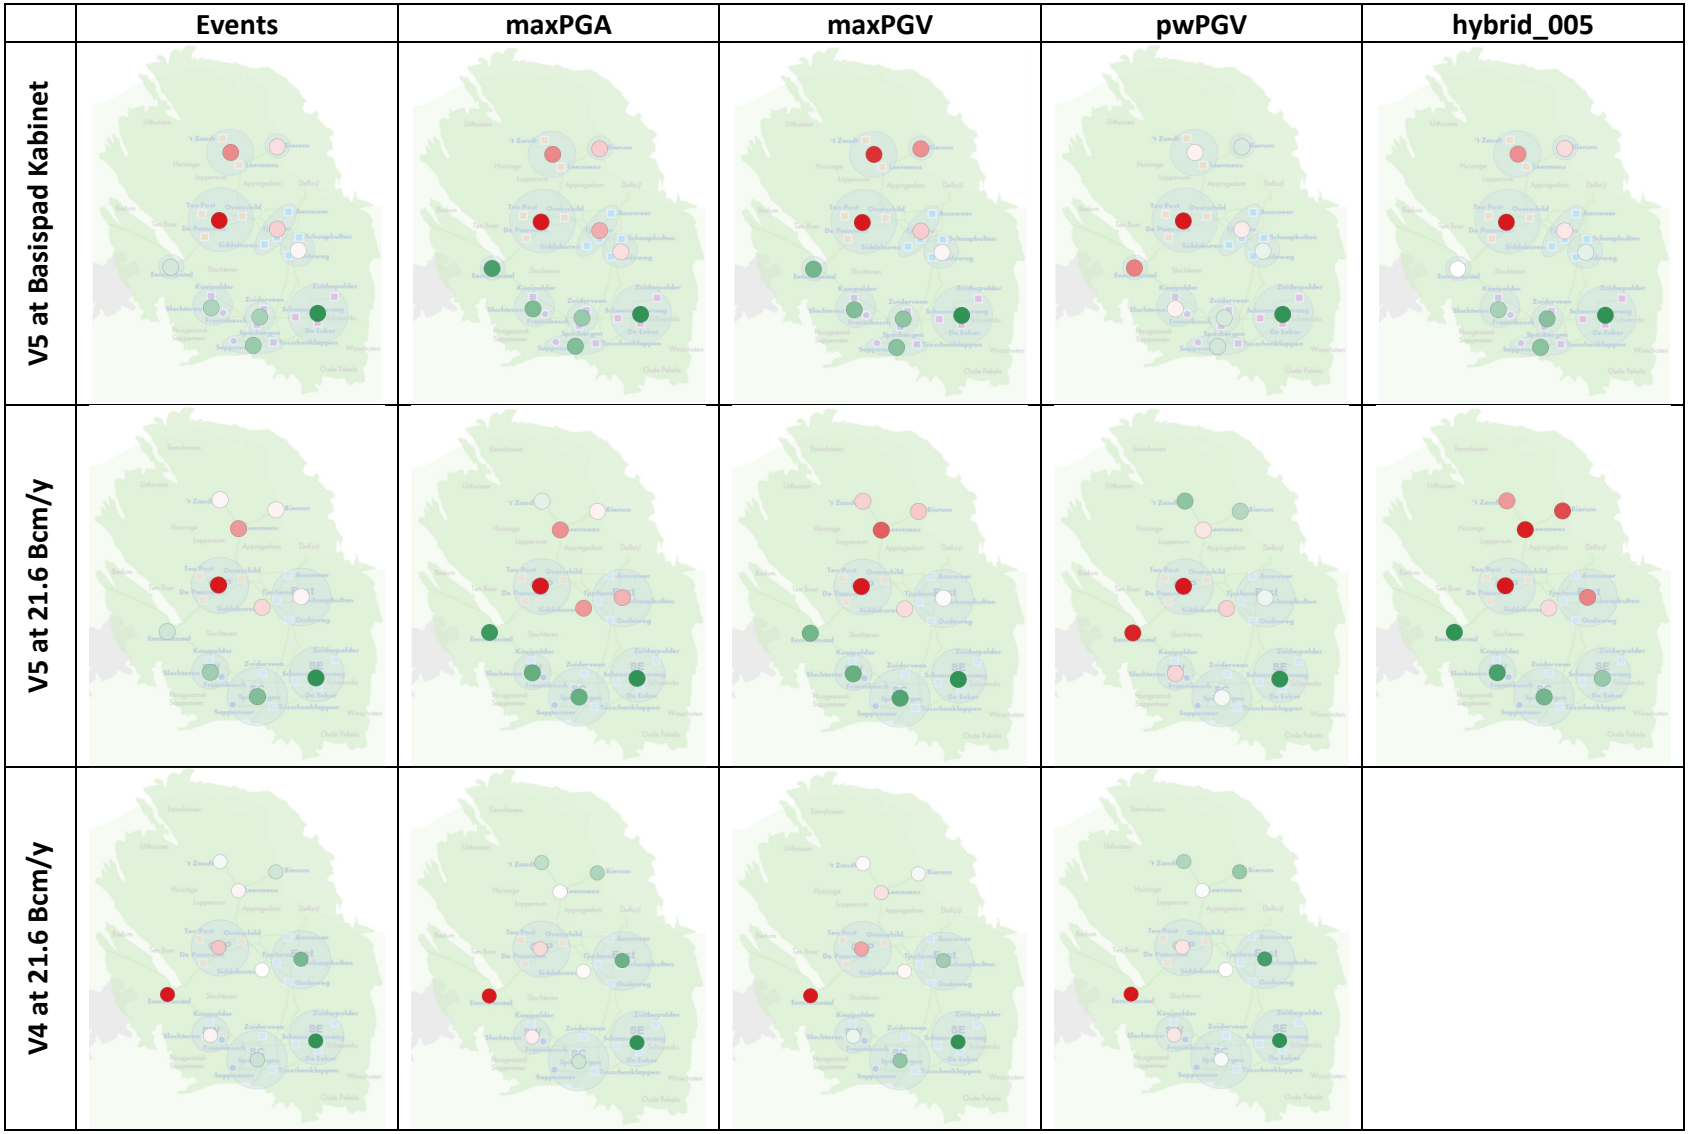

Figure C-11 *Comparison between [V5 Basispad Kabinet], [V5 at 21.6 Bcm/y] and [V4 at 21.6 Bcm/y] for various objective function*

Appendix D List of Abbreviations

Bcm	N.Bcm refers to a volume of a billion normal cubic meters. Normal means the volume is measured at a standard temperature (0 degree C) and pressure (1 bar)
BHP	Bottom Hole Pressure
BP17	Business Plan 2017
DCS	Distributed Control System
EZK	Ministerie van Economische Zaken en Klimaat
GPS	Global Positioning System
Gron	Groningen
GEM	Global Earthquake Model (Global Science Forum of the OECD)
GenREM	Generalised Reservoir Evaluation Model
GMPE	Ground Motion Prediction Equation
GMM	Ground Motion Model
GPS	Global Positioning System
GTS	Gasunie Transport Services B.V.
GWC	Gas water contact
HRA	Hazard and Risk Assessment
InSAR	Interferometric Synthetic Aperture Radar
IPSM	Integrated Production System Model
JT valve	Joules Thompson valve
KNMI	Koninklijk Nederlands Meteorologisch Instituut
KO drum	Knock Out Drum
M	Earthquake Magnitude
ML	Local Earthquake Magnitude
M/Mw	Moment magnitude
MEA	Minister of Economic Affairs
Mmax	Maximum Earthquake Magnitude
MoReS	Modular Reservoir Simulator
MSE	Mean Square prediction Error
NAM	Nederlandse Aardolie Maatschappij B.V.
LTS	Low Temperature Separation
OV	Overslagen (gas custody transfer stations)
PGA	Peak Ground Acceleration
PGV	Peak Ground Velocity
pwPGV	Population Weighted PGV
SAC	Scientific Advisory Committee (Winningsplan 2016)
SodM	Staatstoezicht op de Mijnen (also SSM State Supervision of Mines)
SPMI	Simultaneous Perturbation and Multivariate Interpolation
THP	Tubing Head Pressure
TNO	Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek, Netherlands Organisation for Applied Scientific Research
UGS	Underground Gas Storage