

Water Resources Research®



RESEARCH ARTICLE

10.1029/2024WR038073

Copula Modeling and Uncertainty Propagation in Field-Scale Simulation of CO₂ Fault Leakage

Per Pettersson¹, Eirik Keilegavlen² , Tor Harald Sandve¹ , Sarah E. Gasda^{1,2}, and Sebastian Krumscheid³ 

¹NORCE Norwegian Research Centre, Bergen, Norway, ²University of Bergen, Bergen, Norway, ³Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Key Points:

- Framework for efficient stochastic upscaling, modeling, and uncertainty propagation for CO₂ storage, demonstrated on a North Sea test case
- Stochastic fault properties upscaled to two-phase flow functions with reduced complexity and a format suitable for uncertainty propagation
- Significant computational cost reduction for adaptive stratified sampling compared to Monte Carlo sampling in estimation of CO₂ leakage

Correspondence to:

E. Keilegavlen,
Eirik.Keilegavlen@uib.no

Citation:

Pettersson, P., Keilegavlen, E., Sandve, T. H., Gasda, S. E., & Krumscheid, S. (2025). Copula modeling and uncertainty propagation in field-scale simulation of CO₂ fault leakage. *Water Resources Research*, 61, e2024WR038073. <https://doi.org/10.1029/2024WR038073>

Received 16 JUN 2024
Accepted 18 SEP 2024

Abstract Subsurface storage of CO₂ is an important means to mitigate climate change, and the North Sea hosts considerable potential storage resources. To investigate the fate of CO₂ over decades in vast reservoirs, numerical simulation based on realistic models is essential. Faults and other complex geological structures introduce modeling challenges as their effects on storage operations are subject to high uncertainty. We present a computational framework for forward propagation of uncertainty, including stochastic upscaling and copula representation of multivariate distributions for a CO₂ storage site model with faults. The Vette fault zone in the Smeaheia formation in the North Sea is used as a test case. The stochastic upscaling method reduces the number of stochastic dimensions and the cost of evaluating the reservoir model. Copulas provide representation of dependent multidimensional random variables and a good fit to data, allow fast sampling and coupling to the forward propagation method via independent uniform random variables. The non-stationary correlation within the upscaled flow functions are accurately captured by a data-driven transformation model. The uncertainty in upscaled flow functions and other uncertain parameters are efficiently propagated to leakage estimates using numerical reservoir simulation of a two-phase system of CO₂ and brine. The expectations of leakage are estimated by an adaptive stratified sampling technique which effectively allocates samples in stochastic space. We demonstrate cost reduction compared to standard Monte Carlo of one or two orders of magnitude for simpler test cases, and factors 2–8 cost reduction for stochastic multi-phase flow properties and more complex stochastic models.

Plain Language Summary To limit global warming, greenhouse gases like CO₂ can be injected into large reservoirs of porous rocks below the bottom of the sea instead of being emitted to the atmosphere. CO₂ will slowly move in the reservoirs and may encounter faults, geological features that have properties that can either facilitate or stop the CO₂ from moving further in the underground. It is important that the CO₂ remains in the underground, and hence it is important to understand how it is affected by the fault, in particular when many physical rock properties are unknown due to very few or inexact measurements. We present methods to model the uncertainty in and surrounding the faults and show how more accurate computer simulations can be obtained by a combination of appropriate statistical models and adapted methods to investigate the effect of the fault uncertainty on the risk for leakage of CO₂.

1. Introduction

CO₂ capture and storage (CCS) is anticipated to play a major role in the target reduction of greenhouse gas emissions described in Krevor et al. (2023). By 2050, estimates indicate that CCS will contribute with 13% of the total cumulative emissions reduction (Ringrose & Meckel, 2019). The Norwegian part of the North Sea has an estimated storage potential of 70 gigatons of carbon dioxide (CO₂) (Halland et al., 2011). To put the storage capacity into perspective, the total European Union's CO₂ emissions in 2020 were 2.6 Gt (Crippa et al., 2021). Hence, the North Sea reservoirs are anticipated to be an important means to obtain permanent storage of substantial amounts of CO₂. To fully exploit the potential of the North Sea reservoirs, it is essential to perform numerical simulations to assess the viability of prospective storage sites. Complex physics and associated uncertainties should be taken into account, for instance regarding faults and fractures that are present in potential storage sites.

Faults are complex geological structures with heterogeneous properties that are associated with challenges in data acquisition. A fault zone consists of a fault core surrounded by a damage zone, where the former typically has sealing properties, and the latter exhibits conductive properties (Caine et al., 1996). Most of the deformation and

© 2025. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

displacement is localized to the core, which consists of slip surfaces, fault rock, fractures, cements and mineralizations, and rocks that have been trapped in the core. The damage zone can contain smaller faults, deformation bands, folds, and fractures. For further details of the fault architecture and properties, we refer to Torabi et al. (2020), and references therein. Faults can act as sealing barriers (Pei et al., 2015), but can also be associated with leakage risk (Wu et al., 2021), and the permeability within a fault can vary by orders of magnitude (Sperrevik et al., 2002). In particular, permeability can be enhanced during fault reactivation and induce seismic events that may have an effect on sealing properties, although the integrity of the entire fault may not be compromised (Rinaldi et al., 2014). In-situ stress measurements come with uncertainties that should be taken into account for reliable quantification of seismicity effects (Jeanne et al., 2016). Although direct pore pressure effects tend to dominate over indirect induced seismic effects, the combination of both is typically required for fault reactivation (Fan et al., 2019).

The sealing or conducting properties are largely determined by clay smear, that is, complex geological structures that result when clays from layered beds are entrained along a normal fault. Much effort has been devoted to describe the processes leading to clay smear as well as concepts to describe the amount of clay along the fault, including clay smear potential (Fulljames et al., 1997), shale gouge ratio (SGR) (Yielding et al., 1997), and shale smear factor (Lindsay et al., 1992). Within the current work, the probabilistic shale smear factor (Childs et al., 2007) is also relevant as it is based on probabilistic models to determine the clay distribution that will define the effective fault permeability. The structural properties, for example, fault thickness, are typically highly varying (Childs et al., 1990). Petrophysical properties are often not directly measured but estimated via other properties, for example, permeability as a function of SGR, estimated from well logs (Yielding et al., 1997).

The local fault properties are known only to a limited extent. Seismic imaging cannot resolve the smallest scales, and measurements typically come with uncertainty. In addition to these objective uncertainties, there are subjective uncertainties due to interpretation of data, as demonstrated in Faleide et al. (2021), where a test-panel of 20 geoscientists interpreted the same data with varying results. A framework for uncertainty quantification of seismic imaging, providing detailed maps with confidence indices measuring normalized standard deviations, was introduced in Barbosa et al. (2020). Recently, uncertainty quantification of the fault structure inferred from seismic data was performed by means of convolutional neural networks, providing probability distributions of the faults and demonstrating the partial irreducibility of the uncertainty (Feng et al., 2021). Such purely data-driven approaches that avoid assumptions about certain distributions and independence leads to general probabilistic model descriptions, amenable to the data-driven framework for dependent uncertainties to be presented here. Even if increased seismic resolution were possible, it may still not be possible to represent all small-scale features within the scale of the grid discretization possible to achieve in numerical simulation. An important source of uncertainty that is often not taken into account is the frequency of throw variability along a fault, leading to underestimated uncertainty effects (Manzocchi et al., 2008).

The Vette Fault Zone in the North Sea is a first-order fault zone, hard-linked with two distinct basement structures, and considered a suitable test case in this paper. Structural analysis indicates favorable juxtaposition properties for sealing under CO₂ injection (Mulrooney et al., 2020), and the sealing properties have been further investigated using seismic interpretation (Michie et al., 2021). A probabilistic fault system stability investigation with four clay smearing scenarios resulted in estimates for system failure probability between 10⁻⁴ and 10⁻³ (Rahman et al., 2021). A comparison between a stochastic geocellular model and a fault smear model for the Vette Fault Zone demonstrated limited potential for fluid flow (Bjørnara et al., 2022). Fault seal analysis results based on real SGR data derived from local shale values that in turn were estimated from gamma ray well logs were presented in Michie and Braathen (2024), where it was emphasized that SGR values come with uncertainties stemming from the gamma ray logs.

Comprehensive uncertainty quantification of fault properties is paramount; deterministic base-case simulations often come with order(s) of magnitude errors, and improved measurements of a single source of uncertainty have a limited effect on the accuracy of predictions (Freeman et al., 2008). While a stochastic fault model in itself does not reduce the uncertainty, it provides a quantitative measure and a means to improve predictions. Stochastic modeling of faults amounts to characterization and representation of both the fault core and the surrounding damage zone. First, a structural model needs to be established, for instance using a fault facies model where the fault is divided into discrete geometric objects, each with its own petrophysical properties and different dominating features (Braathen et al., 2009). The properties of the damage zone of the fault structure is typically

dominated by the spatial distribution of deformation bands, and any model of those should honor empirical data (Berge et al., 2022; Schueller et al., 2013). Simulation methods for the spatial distribution of lenses in the fault core and deformation bands in the damage zone are described in Kolyukhin and Tveranger (2015).

Unknown subsurface properties, large spatial and temporal scales, in combination with complex physics has motivated uncertainty quantification in simulation of underground CO₂ storage. A concept for data-driven polynomial chaos expansions to represent uncertainties in reservoir properties was demonstrated to outperform Monte Carlo simulation (Oladyshkin et al., 2011), but did not consider dependencies between input parameters. Correlation between homogeneous reservoir permeability and porosity was accounted for using the Nataf transformation to two independent Gaussian random variables, allowing for application of the classical polynomial chaos expansion combined with systematic basis selection, resulting in accurate CO₂ plume migration at competitive computational speed (Y. Zhang & Sahinidis, 2013). A vertical-equilibrium CO₂ migration model with discontinuous coefficients was combined with polynomial chaos and adaptive integration and showed consistent results, but is limited to a small number of input random variables due to the computational cost (Pettersson, 2016). A similar spectral expansion method in random variables with known and independent distributions was used together with Bayesian methods for uncertainty quantification in CO₂ storage (Jia et al., 2018). Polynomial chaos expansions in general are confined to a relatively small number of random variables, and it is not straightforward to accurately take dependencies between variables into account within this framework. In addition, in cases where the flow is affected by the presence of faults, the uncertainties in fault properties cannot easily be quantified in terms of a moderately sized set of independent random variables. Methods based on Monte Carlo sampling are, in general, less sensitive to the number of input random variables, as demonstrated for faults (Rahman et al., 2021) and long-term fate of CO₂ trapping (Alcalde et al., 2018). However, the statistical errors of these methods converge slowly in the number of samples unless techniques for acceleration are employed, for example, adaptivity or variance reduction.

The goal of this work is to provide a competitive uncertainty quantification framework that can handle and accurately represent dependent physical data from faults and reservoirs, honoring statistical and physical properties. This framework relies on stochastic adaptivity to improve significantly on standard Monte Carlo (SMC) sampling, while being less sensitive to the number of random dimensions than, for example, spectral expansions or grid-based stochastic methods. We aim to demonstrate the importance of stochastic models that honor the underlying physics, and provide geophysicists and geologists with tools to gain further physical insights, for example, by replacing the fault model with their own data. We consider uncertainty of fault properties with regard to pressure communication, fluid flow, and leakage. Geomechanical aspects will not be considered, although we recognize that fault mechanics is an equally important area of uncertainty to quantify.

2. Conceptual Approach

In the context of large-scale CO₂ storage simulation, we are faced by a large number of possibly dependent uncertainties interacting with significant physical complexity. We present a framework for handling this situation, illustrated in Figure 1 and explained below. To fix notation, we aim to perform uncertainty propagation (UP) for a physical model, denoted by M , subject to uncertainties. Specifically, the considered physical model M is a dynamic simulation model that depends on physical parameters $\eta = (\eta_1, \dots, \eta_d)$, for example, permeability and relative permeabilities. Furthermore, M will also depend on stochastic variables $Y = (Y_1, \dots, Y_n)$ to account for the uncertainties affecting the physical problem. These stochastic and physical parameters' dependency is highlighted by the notation $M(Y, \eta)$. This distinction indicates whether the current focus is on physical or stochastic models. However, the two are clearly intertwined in that a physical parameter η_i can be uncertain (say, a fault permeability is a function of unknown fault properties), which we write as $\eta_i = \eta_i(Y)$. For a particular realization of Y and η , the model $M(Y, \eta)$ is used to compute a prediction (e.g., leakage rates) of the physical system. Such an output quantity, denoted Q , also depends on Y and η . Our UP approach in this work focuses on computing statistical properties μ_Q , in particular the expected value, of $Q \equiv Q(Y, \eta)$, in short

$$\mu_Q = UP(M(Y, \eta)). \quad (1)$$

While dynamic simulation models for fluid flow through faults in theory can be coupled with reservoir flow simulation, there are three main reasons to refrain from this path. First, resolving the complex geometry of the

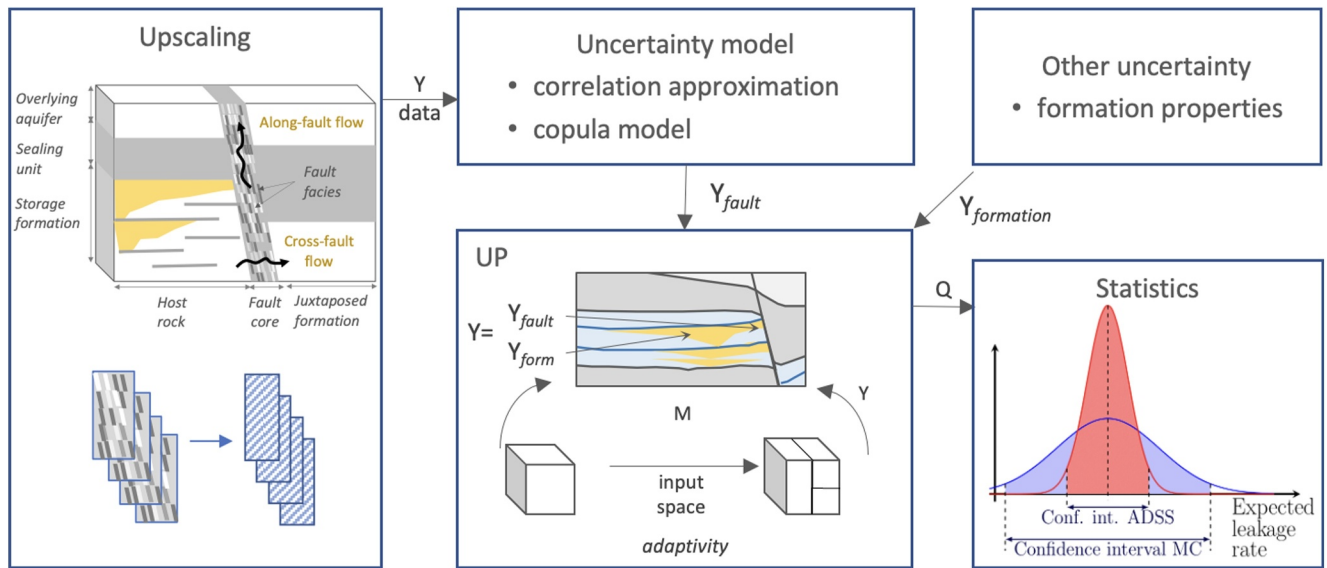


Figure 1. Overview of methodological framework. A stochastic geocellular fault model is upscaled (left panel) and a stochastic model is fitted to the upscaled samples (upper center panel). Together with models for other sources of uncertainty, the uncertainties are adaptively propagated through a physical model via random-space stratification (lower middle panel). Expectations of quantities of interest are anticipated to have lower variance compared to standard Monte Carlo estimates (lower right panel).

fault will lead to a prohibitively high computational cost of the dynamic model M . Second, instead of randomly sampling existing data as input to the dynamic model M , it is often more efficient to first assign or fit a stochastic model for the input distributions. Third, for scarce data, as will often be the case for geological uncertainty, it may be desirable to add further stochastic modeling assumptions that also need to enter the model, for example, an assumed but unobserved upper bound on data. Therefore, in this work we have chosen to treat uncertainty in fault parameters separately from those of the reservoir, as is indicated in Figure 1, where the two are represented through random variables Y_{fault} and $Y_{\text{formation}}$, respectively.

As will be explained in detail in Section 3, we express fault-related uncertainty through upscaled random variables Y_{fault} as illustrated in the left panel in Figure 1. While the upscaling itself follows a standard approach (Rabinovich et al., 2016), which is described in Section 3.2, the coarse-scale flow functions are represented by a stochastic model employing copulas, see Section 3.4 for details. Copulas comprise a statistical modeling framework that is suitable for representation and analysis of multivariate (in particular dependent) data. Since the seminal work in Joe (1996), building upon the results in Sklar (1959), copulas have been used in a wide range of applications, see Bhatti and Do (2019) for a list of examples. Copulas have gained popularity within the water resources community for analysis of, for example, rainfall data and flood forecasting, as demonstrated by entire monographs devoted to these topics (Chen & Guo, 2019; L. Zhang & Singh, 2019). To our knowledge, however, copulas have not yet been employed to parameterize uncertainties in fault properties to be used for forward propagation of subsurface flow properties. There are different parametric families of copulas, each with their own tail dependence, that is, modeling to what extent an extreme event in one variable is correctly paired with an extreme event in another variable. This is an important feature in the context of CO_2 storage where extreme behavior, for example, leakage, is of particular interest (Elsheikh et al., 2014; Pettersson et al., 2022).

The copula representation of Y_{fault} together with a model for $Y_{\text{formation}}$, which we will represent by simple approaches as detailed in Section 6, form the input to the uncertainty propagation (see Figure 1). One can distinguish between random sampling methods or statistical methods on the one hand, and deterministic uncertainty propagation methods on the other hand. The former includes variants of Monte Carlo methods, for example, Latin hypercube sampling (McKay et al., 1979), multilevel Monte Carlo (Giles, 2015), and importance sampling (Kloek & van Dijk, 1978), while the latter family of methods includes Quasi-Monte Carlo (L'Ecuyer & Lemieux, 2002) and generalized polynomial chaos via projection (Xiu & Karniadakis, 2002) or stochastic collocation (Babuška et al., 2007). Whether statistical or deterministic, there is typically a trade-off between robustness and theoretical accuracy. The computational complexity of the different groups of methods grows

more or less with the number of random dimensions and, to some extent, the dependencies between random inputs. For the current work, non-intrusive random sampling methods are attractive as they do not require modifications of existing physical solvers, and they typically scale well with the number of random dimensions. In addition, the UP method selected should not impose regularity requirements on the quantity to be approximated.

For uncertainty propagation from random inputs to the quantities of interest, that is, μ_Q introduced in Equation 1, we will use the adaptive stratified sampling method presented in Pettersson and Krumscheid (2022), where samples are adaptively concentrated to regions in input random space that has a large impact on the quantity of interest. The method is sampling-based and does not assume apriori knowledge about what inputs have the greatest impact. It performs best when the variability is localized to a small subset of random variables, or the variance is localized to a subdomain that may not have smaller dimension than the full stochastic space. Related methods that build upon adaptive sampling and stratification include (Étoré et al., 2011; Krumscheid & Pettersson, 2023; Shields, 2016; Shields et al., 2015). The adaptive stratified sampling method employed here is further described in Section 4.

For the dynamic model M in Equation 1, we utilize existing simulation tools that model large-scale CO₂ storage and thus treat the simulator as a black-box model. Commercial simulators like Eclipse and Intersect from SLB, and GEM from the CMG group, have dedicated CO₂ modules, as have open-source simulators like DuMuX (Flemisch et al., 2011), GEOX (Gross & Mazuyer, 2021), MRST (Lie, 2019), TOUGH2 (Pruess, 1991) and OPM-Flow (Alvestad et al., 2021). Developing simulation tools is not the focus of this study and we use the OPM-Flow simulator, which has a dedicated CO₂ module, supports standard industry input/output formats, and is fast and free to use (Sandve et al., 2021, 2022). The latter points are critical for uncertainty quantification, as we will need to run a large number of simulations with varying inputs η and Y . It is noteworthy that OPM-Flow can account for all relevant fluid flow processes involved in field-scale CO₂ storage in faulted reservoirs, and connected aquifers: CO₂ injection, migration and trapping within the storage formation, along and cross-fault pressure communication and fluid flow.

We test our methodology consisting of stochastic upscaling together with copula representation and further stochastic modeling, and adaptive stratified sampling, on the Smeaheia formation in the North Sea, where the framework will be applied to quantify leakage from the Vette fault zone. We base the physical model on the reservoir data set openly available from CO2Share (<https://co2datashare.org/>), but extend it to include aquifers that are connected to the reservoir through vertical and horizontal flow in the Vette fault zone. The upscaled stochastic two-phase flow functions specific to the Vette fault are constructed in Section 3. These functions depend on Y_{fault} , that is, they reflect uncertainty in the fine-scale distribution of facies within the fault based on available data and a standard fault facies model. Uncertainty in the formation parameters, $Y_{\text{formation}}$, is introduced in Section 6, where we also present simulations probing the performance of the stochastic framework. The objectives of this work include demonstration of how the upscaled stochastic flow functions can be included in a reservoir model, and that our sampling strategy gives substantial improvements compared to SMC methods. The methodological components are not limited to the models selected, and hence more widely applicable.

3. Stochastic Fault Model

Here we describe a stochastic modeling framework for the fault properties on the fine-scale and how these quantities can be upscaled for use as stochastic input for the reservoir simulator. This amounts both to derive parameters for a physically coarser scale, and simultaneously reduce the stochastic complexity by reducing the number of random parameters. While the derivation uses advanced stochastic modeling, the goal is effective stochastic flow functions where each realization is compatible with and has a format similar to what would have been expected in a deterministic model.

We split the stochastic fault modeling into three stages: First, the fine-scale geology within the fault is described in terms of a geocellular fault facies model with a stochastic description of the distribution of facies, as well as of their flow properties as we describe in Section 3.1. Next, we generate realizations of the fine-scale geometry and compute upscaled permeabilities, and capillary and relative permeability functions valid for the entire fault, see Sections 3.1 and 3.2, respectively. The final stage entails the representation of the upscaled quantities in a format suited for the stochastic sampling of the field-scale model described in Section 4. In practice, this means the

coarse-scale models should be expressed in terms of independent stochastic variables with known distributions, which is the assumed format for most uncertainty propagation methods.

The coarse-scale absolute permeability can relatively simply be brought into a suitable form, again see Section 3.1. For the two-phase flow functions the situation is more challenging, partly reflecting the additional complexity of stochastic functions compared to that of constants (e.g., absolute permeability): Upscaling from fine-scale realizations produces point clouds of possible two-phase flow properties, which can be represented by equivalent probability distributions that consider only coarse-scale properties, as shown in Section 3.3. This removes the link to the fine-scale model, however, the variables in the stochastic model are statistically dependent. These dependencies can be incorporated into the modeling framework by using so-called copulas, as is done in Sections 3.4 and 3.5.

While arguably complicated, the presented approach has two main advantages compared to applying a more direct coupling between the stochastic fine-scale geology and the field-scale simulations: First, properties such as dimension, variability, and heterogeneity of the stochastic fault facies model will not be seen directly by the field-scale reservoir model. Thus, advanced stochastic descriptions of fault facies can be applied and the added computational effort will only be seen in an offline stage, which is often cheap compared to field-scale simulations. Second, the modularization allows for applying different level of sophistication for the individual components. For instance, the analytical upscaling described in Sections 3.1 and 3.2 could have been replaced by numerical upscaling. Although this may also necessitate modifying the coarse-scale stochastic modeling to avoid incurring modeling errors in that step (a more complex stochastic model may be needed), there is still further significant flexibility in how our framework is implemented. For instance, a different fine-scale model can be used as long as it is possible to produce independent samples, and the coarse-scale parameters are not assumed to be independent or belong to any specific class of distributions. Indeed, we expect this flexibility to be particularly useful in the stochastic modeling, since the coarse-scale properties to be represented may vary significantly with the fine-scale model.

3.1. Stochastic Modeling of Fault Facies Distribution

As a simple non-trivial example of a stochastic fault-facies geocellular model, we consider a model based on the work presented and validated against a reference model (with less than 8% error using a 95% confidence interval) in Bjørnará et al. (2021). The model consists of facies, or geocells, that are stacked on top of each other. The geological properties are assumed constant within each cell, and can vary discontinuously between cells. The geometric distribution of facies is assumed to only vary in the vertical direction, while the facies flow properties can be expressed in terms of the SGR, which is constant within each facies. SGR is a function of fault throw and local clay volume fraction $V_{c,j}$ in the j th reservoir zone with thickness Δz_j , that is,

$$\text{SGR} = \frac{\sum_j V_{c,j} \Delta z_j}{\text{throw}} \times 100\%,$$

as defined in the seminal work in Yielding et al. (1997). For the Vette Fault Zone, SGR can be obtained from data as follows. The Vette fault throw and juxtaposition can be obtained from seismic interpretation (Michie et al., 2021; Mulrooney et al., 2020). A gamma ray log from a well located in the relative proximity of the hanging wall of the fault is then converted to clay contents values, as described in Bjørnará et al. (2023). The SGR values can then be obtained from the expression above. For details on how to obtain SGR data for the Vette Fault Zone, we refer to Michie and Braathen (2024). The range of SGR values is confined to the interval $[0, 100]$, where the endpoints correspond to pure sandstone (SGR = 0) and shale (SGR = 100). The model assumes that the total fault height H and the number of facies N_{facies} is known. The interior boundaries between the facies are assumed to be uniformly distributed around expected locations that are equidistant on $[0, H]$. The assumption of no variation in the horizontal direction means that a realization of the fault geometry can be expressed in terms of a one-dimensional model. The SGR distribution of a given facies is assumed to be independent of the other facies and to follow a Gaussian distribution with mean given by interpolation of the depth-dependent synthetic data shown in Figure 2, and an assumed standard deviation of 14 percentage points (Bjørnará et al., 2021; Freeman et al., 2008). Due to data protection policy, we do not have direct access to SGR data, but use the synthetic data shown in Figure 2 that are assumed to approximate those in Figure 5 (center) in Bjørnará et al. (2021) and obtained from real data as described in detail in Michie and Braathen (2024). We note that the distributions can be non-




Figure 2. (a) Mean shale gouge ratio (SGR) based on synthetic data modeling the Vette fault. Upscaled permeability distributions from Equations 2 and 3, and different clay permeabilities: (b) $k_c = 10^{-4}$ mD; (c) $k_c = 10^{-3}$ mD, and (d) $k_c = 1$ mD. SGR assumed to follow a Gaussian distribution with a standard deviation of 14% for all cases.

Gaussian and obey different distributions across the geocells without affecting the upscaling method. The probability of this model yielding results outside the interval $[0, 100]$ is assumed to be negligible. In total, the model has $2N_{\text{facies}} - 1$ random variables as inputs.

For each facies j ($j = 1, \dots, N_{\text{facies}}$), there are various functional relationships to obtain permeability k_j from the corresponding SGR, denoted SGR_j . For instance, in Bjørnå et al. (2021), where the permeability within facies j is assumed constant, the relationship

$$\log(k_j) = 0.01 \cdot \text{SGR}_j \cdot \log(k_c/k_s) + \log(k_s), \quad (2)$$

was given, where k_c is the permeability of clay-rich shale, and k_s is the permeability of sandstone. Other functional relationships between SGR and permeability are summarized in Jolley et al. (2007), and a general discussion of methods for approximating the effect of clay content on fault permeability is provided in Fisher and Jolley (2007). In this work, we make the modeling decision to consider $k_c = 1$ mD as in Bjørnå et al. (2021), as well as $k_c = 0.001$ mD and $k_c = 0.0001$ mD based on the observation that many SGR-permeability models in the literature yield substantially lower output values than Equation 2 with $k_c = 1$ mD. For other SGR models that have been tested, the results aligned well with the stochastic modeling framework to be presented in this work, in the sense that the models themselves were different but of similar complexity and similar goodness-of-fit.

The upscaled uncertain permeability K is computed from the weighted harmonic mean,

$$K = \frac{\sum_{j=1}^{N_{\text{facies}}} h_j}{\sum_{j=1}^{N_{\text{facies}}} \frac{h_j}{k_j}}, \quad (3)$$

where h_j, k_j are the random height and permeability of facies j for $j = 1, \dots, N_{\text{facies}}$. In the case of single-phase flow, the permeability can be represented by a single random variable that does not depend on any other random variable. To enable conditional sampling of the permeability (e.g., random sampling restricted to a subrange of the random variable), an empirical or parametric distribution (e.g., beta or lognormal) can be fitted to a large set of permeability samples. Conditional samples are then obtained from $F_{\text{fit.dist}}^{-1}(u)$, where $F_{\text{fit.dist}}^{-1}$ denotes the inverse cumulative distribution function (CDF) of the fitted distribution, and u is drawn uniformly from a subset of $[0, 1]$.

Empirical histograms, approximating the probability density functions (PDFs) for the permeability obtained from Equation 3 and the three SGR-permeability models are shown in Figures 2b–2d. An alternative to deterministic and fixed values of k_c and k_s is to account for the uncertainties in the composition by modeling them as random variables, but this approach will not be further investigated here.

3.2. Upscaling of Two-Phase Flow Functions

Similar to the absolute permeability, the fault relative permeability and capillary pressure functions will be stochastic due to their dependency on the stochastic fine-scale facies. However, since the modeling of the upscaled quantities must capture stochastic functions rather than just a number, as was the case for the absolute permeability, the task is more involved.

For the computations presented herein, the fine-scale capillary pressure function is computed from a Brooks-Corey function (Brooks, 1965). Assuming that the entry-pressure scales with the square root of permeability (Leverett, 1941) the entry-pressure in a geocell can be computed from the permeability in the geocell k_{cell} and the entry-pressure $p_{\text{s,entry}}$ and permeability k_{s} in the sand as

$$p_{\text{cell, entry}} = p_{\text{s,entry}} \cdot \sqrt{k_{\text{s}}/k_{\text{cell}}}. \quad (4)$$

Thus, as for the absolute permeability, the fine-scale capillary pressure is a stochastic function of the SGR. The relative permeabilities are not stochastic on the fine scale, but will be so on the coarse scale due to the dependency of the fine-scale geometry and permeability distribution. For the fine-scale simulations herein, we use $p_{\text{s, entry}} = 2.5$ kPa and Brooks-Corey exponent of 0.67. These are values computed to fit the capillary-pressure values in the Smeaheia data set (see Section 5.2).

To compute upscaled two-phase flow functions for a given facies distribution, here we assume that the phases are in capillary equilibrium. While this is a strong assumption that may be of questionable validity, for example, in cases where viscous forces have significant impact on the flow, it allows for analytical upscaling of relative permeabilities and capillary pressure. Specifically, we follow the approach described in Rabinovich et al. (2016), which for completeness is summarized in a form tailored to the setting considered here in Appendix A. The result is a set of discrete values for the flow functions, in the sense of point values for permeabilities, and the corresponding values for capillary pressures, saturations, and relative permeabilities parametrized via the deterministic parameter s_{d} . This is related, but not equivalent, to a fine-scale saturation at which the flow functions are sampled; see the appendix for an explanation.

3.3. Reduced Multivariate Stochastic Model for Flow Functions

We apply the upscaling described in Section 3.2 to a set of realizations of the fine-scale fault geometry, using the same realizations as those used to upscale permeability. Five coarse-scale flow functions are needed as input to the model M for the case of two-phase flow: permeability K , capillary pressure P_{c} , saturation S , wetting relative permeability K_{r}^{w} , and non-wetting relative permeability K_{r}^{nw} . A large number of random samples of the coarse-scale flow functions evaluated at various deterministic fine-scale saturation values s_{d} can be generated at negligible computational cost. The random flow functions are so far parameterized via the $2N_{\text{facies}} - 1$ random variables, but the upscaling implies that the stochastic dimensionality may be reduced to obtain a more succinct representation which is anyhow the format effectively accepted by the coarse-scale simulator.

The stochastic representation of coarse-scale properties deviates from what would be expected of two-phase flow functions in a traditional deterministic modeling framework. First, the model expresses the coarse saturation as a function of s_{d} on the same level as the capillary pressure and relative permeabilities. This is contrary to deterministic modeling, where the saturation is considered an independent variable, and the capillary pressure and relative permeability are functions of this saturation. This is simply a consequence of the combination of the fine-scale stochastic description and the upscaling procedure; as will be seen, the coarse-scale flow functions have the expected dependency of saturation, albeit in a stochastic sense, but deriving this relation is not straightforward. Second, although the absolute permeability is not a deterministic function of saturation, nor of the two-phase flow functions, there is a clear random dependence between these parameters, as will be shown below.

Without any further conditions, the appropriate a priori model of the parameter quintuple $K, P_{\text{c}}, S, K_{\text{r}}^{\text{w}}, K_{\text{r}}^{\text{nw}}$ would be a five-dimensional stochastic process in s_{d} , for instance, parameterized by a vector-valued Karhunen-Loeve expansion with suitable covariance structure (Perrin et al., 2013). Such a representation would, however, still be complex and with a substantial computational cost and may not guarantee physically relevant output. These considerations therefore motivate us to search for a different method. As we will demonstrate below, the data are indeed well represented by a multivariate random variable model with five distinct but dependent random

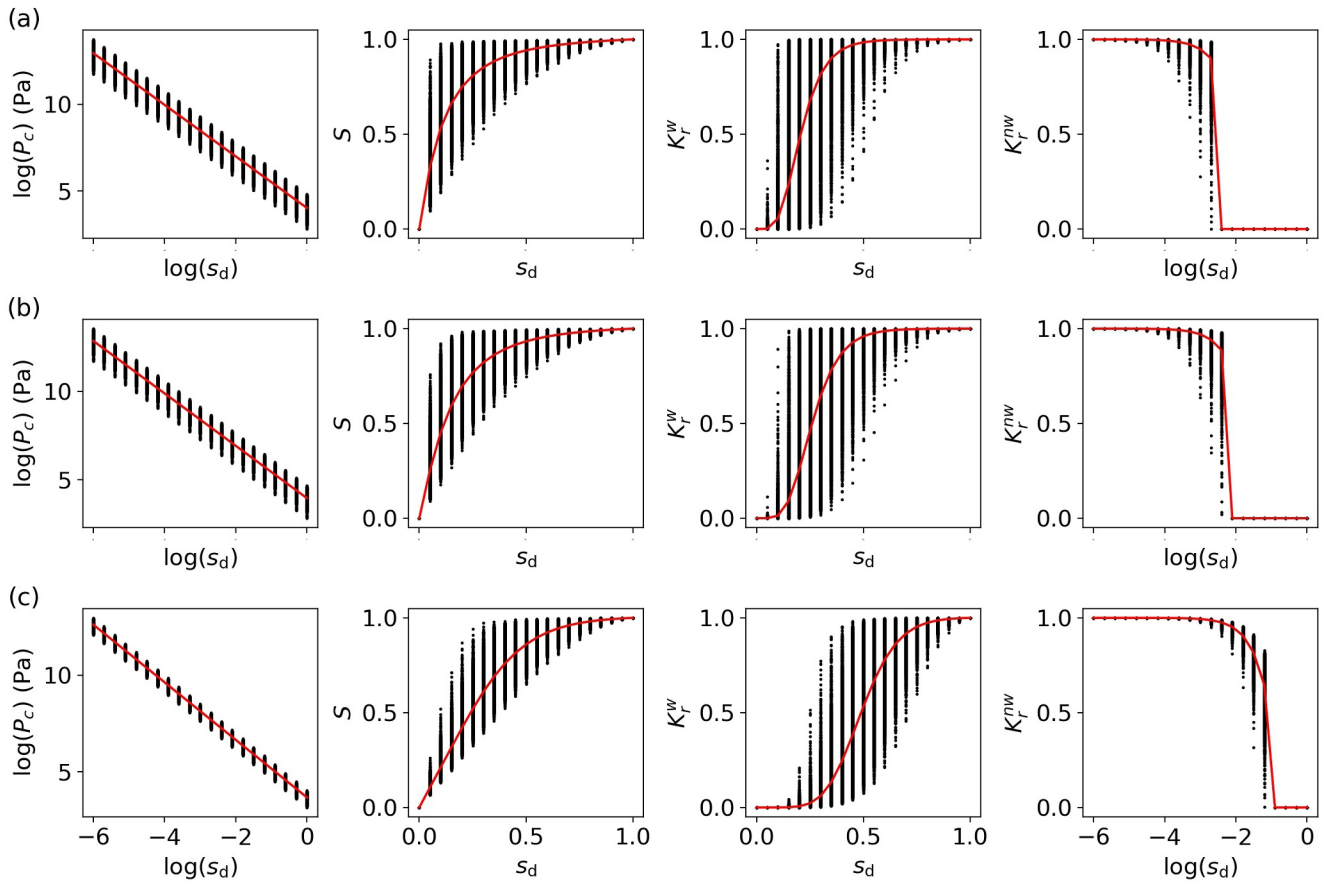


Figure 3. Visualization of random data (black dots) as support for stochastic modeling choices. Red curves indicate the means of the data. For ease of visualization, equidistant discretization of s_d is used for capillary pressure and non-wetting relative permeability, and a log-scale discretization of s_d for saturation and wetting permeability. (a) $k_c = 0.0001$ mD (upper row), (b) $k_c = 0.001$ mD (middle row), and (c) $k_c = 1$ mD (lower row).

variables, whose distributions are to be determined to match data at some user-defined value of s_d , henceforth denoted s_d^0 . We then seek functional relationships to make sure the model matches the data for all relevant values of s_d . Ultimately, such a model is fully and sufficiently motivated by presenting an acceptable goodness-of-fit and showing that it cannot be further simplified. Nevertheless, we find it instructive to illustrate the model choice via the data we seek to fit as follows. To that end, we generate $N_{\text{ref}} = 10,000$ samples of the coarse-scale flow functions evaluated on each of 21 values of s_d on the unit interval. Note that we have adapted the discretization to the scale of the s_d -axis, but in all numerical experiments we stick to the logarithmically spaced discretization that is denser for small values of s_d . The computation is done for three different values of k_c , and the results are shown in Figure 3.

The upscaled permeability is not a function of s_d , but so are the remaining four coarse-scale functions. We first seek to characterize the dependence on the deterministic parameter s_d , and subsequently the residual is parameterized by random variables. As can be seen from Figure 3, for all values of k_c there is a very good log-linear fit of the mean of the capillary pressure with respect to s_d . We note that the values taken by the capillary pressure at any two different values of s_d are perfectly correlated, that is, the same single random parameter models all randomness in the capillary pressure.

The upscaled saturation and wetting relative permeability share some characteristics that make the stochastic modeling challenging. The randomness in each of them exhibits positive but not perfect two-point correlation (where perfect correlation means that the random outcome at any given value of s_d completely determines the value at any other value of s_d). This lack of perfect correlation is due to crossing of sample paths, where a sample path here is to be understood as the flow function traced out by keeping the random variable it is dependent upon at a fixed value, and varying s_d only. To be more specific, saturation sample path crossing means that there are two

different values of s_d , say $s_d = s'$ and $s_d = s''$ (i.e., with $s' \neq s''$), and two distinct random samples $Y^{(m)} \neq Y^{(n)}$, for which we have for the saturation S that $S(s', Y^{(n)}) > S(s', Y^{(m)})$ and $S(s'', Y^{(n)}) < S(s'', Y^{(m)})$. A suitable parametric stochastic model should reflect the correlation while ensuring that the output is monotone in s_d to be compatible with the physical interpretation, and within physically relevant bounds, for example, the unit interval for saturation and relative permeabilities. Finally, the mean of the upscaled non-wetting relative permeability K_r^{nw} can be approximated by a piecewise linear function in s_d . There is nearly perfect correlation in the randomness of this function, and it can be parameterized with a single random variable and a small correction for unphysical values that occur with small probability. These properties would become more clear by plotting K_r^{nw} as a function of s_d , but we have chosen to plot K_r^{nw} as a function of $\log(s_d)$ in Figure 3 for visibility of the full range of s_d .

Based on the above considerations and the observed data, an appropriate stochastic model for the upscaled parameters is given by

$$K = \exp(Y_1), \quad (5)$$

$$\log(P_c(s_d)) = a_{P_c} + b_{P_c} \log(s_d) + Y_2, \quad (6)$$

$$\log(S(s_d)) = F_{Y_3, s_d}^{-1}(T_{s_d}^S F_{Y_3}(Y_3)), \quad (7)$$

$$\log(K_r^{nw}(s_d)) = F_{Y_4, s_d}^{-1}(T_{s_d}^{K_r^{nw}} F_{Y_4}(Y_4)), \quad (8)$$

$$K_r^{nw}(s_d) = \begin{cases} 1 + (a_{K_r^{nw}} + b_{K_r^{nw}} Y_5^{(s_d)}) s_d & \text{if } s_d \leq s_d^0 \\ 0 & \text{if } s_d > s_d^0 \end{cases} \quad (9)$$

with

$$Y_5^{(s_d)} = \begin{cases} Y_5 & \text{if } Y_5 \geq -(1 + a_{K_r^{nw}} s_d)/(b_{K_r^{nw}} s_d) \\ -(1 + a_{K_r^{nw}} s_d)/(b_{K_r^{nw}} s_d) & \text{if } Y_5 < -(1 + a_{K_r^{nw}} s_d)/(b_{K_r^{nw}} s_d) \end{cases}$$

where $\{Y_j\}_{j=1}^5$ is a set of five dependent random variables, to be estimated from data. Due to the orders-of-magnitude variability of permeability an exponential relationship has been employed, although this is not necessary in a conceptual setting. In addition, we note that the log transformations in Equations 7 and 8 are not necessary, but they do make sure that the output values are positive.

Here, F_{Y_3} (and F_{Y_4}) denotes the CDF of a given random variable Y_3 (and Y_4), where the dependence on $s_d = s_d^0$ has been suppressed for brevity of notation. The notation $T_{s_d}^S$ (and $T_{s_d}^{K_r^{nw}}$) denotes a mapping from the unit interval to the unit interval, to account for correlation between data of S (and K_r^{nw}) at s_d^0 and other values of s_d . The purpose of the latter is to account for the fact that sample paths of S or K_r^{nw} (as functions of s_d) can cross each other (i.e., a sample path of S crosses another path of S). The mapping $T_{s_d}^S$ is constructed as follows, aiming to reflect the underlying stochastic dependencies. Recall that we have chosen a fixed $s_d = s_d^0$ as reference for the parameterization of the random variables $\{Y_j\}_{j=1}^5$, that is, we will use data for this particular value s_d^0 only to define and estimate the joint distributions of $\{Y_j\}_{j=1}^5$. Any value of s_d can be chosen for this purpose, but it is natural to choose the smallest considered s_d as a starting point for a random process. The upscaled data at s_d^0 are then sorted to be used as a reference order. For any target $s_d > s_d^0$, the data are sorted locally with respect to the reference data at s_d^0 . If no crossing of sample paths has occurred, the reference and target indices will be identically ordered, and $T_{s_d}^S$ would be the identity operator. In general, the order of indices will differ, and a one-to-one mapping between them is given by pairing entry by entry in the order of appearance. The range and domain of the mapping can both be normalized to the unit interval to reflect the fact that we want to map between the ranges of cumulative distribution functions. It is however practical to consider an integer-index discretization as follows. Let I be the indices

permuting the reference data to the target data, and \mathcal{I}^{-1} the indices that define the inverse mapping. We then introduce the approximation

$$T_{s_d}^S F_{Y_3}(Y_3) \approx N_{\text{ref}}^{-1} \mathcal{I}^{-1} [N_{\text{ref}} F_{Y_3}(Y_3)],$$

where double brackets denote integer rounding. Thus, the model $T_{s_d}^S$ approximates the two-point correlation present in the data, which is not stationary, that is, the correlation depends on the actual values of s_d , and not just the difference between any two points.

The mapping F_{Y_3, s_d}^{-1} is given by the inverse of the empirical CDF at any given value of s_d , and ensures that the pointwise distribution approximates that of the underlying data. Hence, the accuracy of both the mappings $T_{s_d}^S$ and F_{Y_3, s_d}^{-1} grows with the number of reference samples N_{ref} . The discretization of the model for K_r^w is completely analogous to that of S .

To summarize the above considerations, the stochastic facies description and upscaling method applied here, results in an upscaled stochastic model that can be described by relatively simple functional forms. We emphasize that changing either the facies description or the upscaling procedure may necessitate a new analysis of the upscaled model, and the stochastic model may take a different form, but the basic modeling paradigm and steps in the calculation remain the same.

3.4. Copula Models and Transformations for Multivariate Distributions

The stochastic upscaled model Equations 5–9 provides a probabilistic description of the fault flow function in terms of coarse-scale parameters. However, the random variables $\{Y_j\}_{j=1}^5$ do not follow known distributions, nor are they independent. Both conditions are necessary for the model to be amenable to the computational uncertainty propagation method applied below. We therefore parameterize and transform our stochastic model, using a combination of two techniques: First, copulas are employed to derive a stochastic model which honors the data, and have the attractive features of separation of dependency structure and marginal distributions, and a parameterization of (albeit dependent) uniform distributions on the unit interval. Second, dependencies between the uniform variables are handled by transforming them to independent uniform variables using the inverse Rosenblatt transform (Rosenblatt, 1952). Thus, the model provides an accurate mapping between standard and independent uniform random variables suitable for uncertainty propagation methods, and physically accurate random flow functions in the correct format for input to a reservoir simulator. The application of both copulas and Rosenblatt transformations are described in some detail below, starting with an outline of copula construction of multivariate distributions.

According to Sklar's theorem (Sklar, 1959), any multivariate CDF $F_Y(y_1, \dots, y_n)$ of the random vector $Y = (Y_1, \dots, Y_n)$ with marginal CDFs $F_{Y_1}(y_1), \dots, F_{Y_n}(y_n)$ can be expressed as

$$F_Y(y_1, \dots, y_n) = C_Y(F_{Y_1}(y_1), \dots, F_{Y_n}(y_n)), \quad (10)$$

where C_Y is a so-called n -copula. Differentiating the above expression and using the notation $c_Y(u_1, \dots, u_n) \equiv \partial_{u_1} \dots \partial_{u_n} C_Y(u_1, \dots, u_n)$, the joint PDF can be written,

$$f_Y(y_1, \dots, y_n) = c_Y(F_{Y_1}(y_1), \dots, F_{Y_n}(y_n)) \prod_{j=1}^n f_{Y_j}(y_j),$$

where c_Y is the copula density function, and $f_{Y_j}(y_j)$ denote the marginal PDFs of the individual random variables Y_1, \dots, Y_n . The copula framework separates the multivariate dependence properties described by the copula from the modeling of the marginal distributions. There are several commonly used families of copulas with different properties suitable for different kinds of data. Having chosen a suitable copula family, one can use data to estimate its parameters, and then use the fitted copula model to generate new samples from a multivariate distribution that approximates the distribution of the original data.

Vine copulas offer a framework for multidimensional dependence modeling, relying on the use of repeated conditioning of up to two variables on subsets of all the remaining random variables. This leads to replacement of a multivariate function by a factorization into members of families of standard bivariate copulas (Bedford & Cooke, 2002). The literature on high-dimensional copulas is more limited than that on bivariate copulas, and existing high-dimensional copulas can fail to correctly capture tail dependence, that is, extreme behavior in several variables simultaneously. Conditioning on bivariate copulas also adds flexibility in the sense that different families of copulas can be used for different subsets of variables, allowing for a more accurate fit (Aas & Berg, 2009). For more details and an illustrative example of exact copula factorizations for three random variables, we refer to Czado and Nagler (2022). Note that the copula factorizations are exact, but in practice each bivariate distribution is estimated from data, introducing an approximation error. Using data, both an appropriate vine copula factorization, bivariate copulas, and marginal distributions can be estimated. Pair-copula modeling, to which the class of vine copulas belong, are suitable for relatively high-dimensional modeling with the number of pair-copulas to be estimated being equal to $n(n+1)/2$ (Aas & Berg, 2009). However, this can become computationally prohibitive if the number of degrees of freedom from a large number of candidate models become very large (Brechmann & Schepsmeier, 2013). To remedy the computational cost and avoiding overfitting in the absence of large sets of data for those cases, sparse copulas partially relying on the use of the parameter-free independence copula have been proposed (Nagler et al., 2019).

Once the copula model has been fitted to data, it can be sampled to provide inputs for uncertainty propagation through a physical model. The random variables $\hat{U}_j = F_{Y_j}(Y_j)$ are uniformly distributed on the unit interval (0,1), but they are, in general, not independent. For efficient coupling with many uncertainty propagation methods, a mapping between $\{\hat{U}_j\}_{j=1}^n$ and a set of independent uniformly distributed random variables $\{U_j\}_{j=1}^n$ is required. Such a mapping is offered by the inverse of the Rosenblatt transformation (Rosenblatt, 1952), from the conditional variables $\{U_j\}$ to the marginals $\{\hat{U}_j\}$ which are the arguments of the copula Equation 10. Specifically, we have the two sets of random variables:

$$\begin{aligned} \hat{U}_1 &= F_{Y_1}(Y_1), & U_1 &= F_{Y_1}(Y_1), \\ \hat{U}_2 &= F_{Y_2}(Y_2), & U_2 &= F_{Y_2|Y_1}(Y_2|Y_1), \\ & \vdots & & \vdots \\ \hat{U}_n &= F_{Y_n}(Y_n), & U_n &= F_{Y_n|Y_1, \dots, Y_{n-1}}(Y_n|Y_1, \dots, Y_{n-1}). \end{aligned}$$

The output from the inverse Rosenblatt transformation then gives a means to express the target random variables $\{Y_j\}_{j=1}^n$ as a function of the parametric input distributions:

$$\begin{aligned} Y_1 &= F_{Y_1}^{-1}(U_1), \\ Y_2 &= F_{Y_2|Y_1}^{-1}(U_2|Y_1), \\ & \vdots \\ Y_n &= F_{Y_n|Y_1, \dots, Y_{n-1}}^{-1}(U_n|Y_1, \dots, Y_{n-1}). \end{aligned}$$

3.5. Results: Copula Fit for Upscaled Flow Functions

The upscaling model described in previous sections is used to generate $N_{\text{ref}} = 10,000$ samples of the dependent random variables $\{Y_j\}_{j=1}^5$ given by Equations 5–9, and also as shown in Figure 3. These samples are treated as data to which a copula model is fitted to allow targeted sampling and application of various uncertainty propagation methods. The Python interface `pyvinecopulib` to the C++ `vinecopulib` package (Czado & Nagler, 2022; Nagler et al., 2022) is employed for this objective, resulting in copula models defined by the tree structures illustrated in Figure B1 in Appendix B. The tree structures provide efficient representations for storage and further sampling and show which pairs of random variables exhibit the strongest dependencies. For all three values of clay permeability, the dependence structure is the same, that is, the nodes and edges are identical. This particular tree structure is referred to as a D-vine (Direct vine) in the literature. The bivariate copulas describing the kind of dependence between the pairs of variables differ between the cases. For the majority of them, there are no suitable

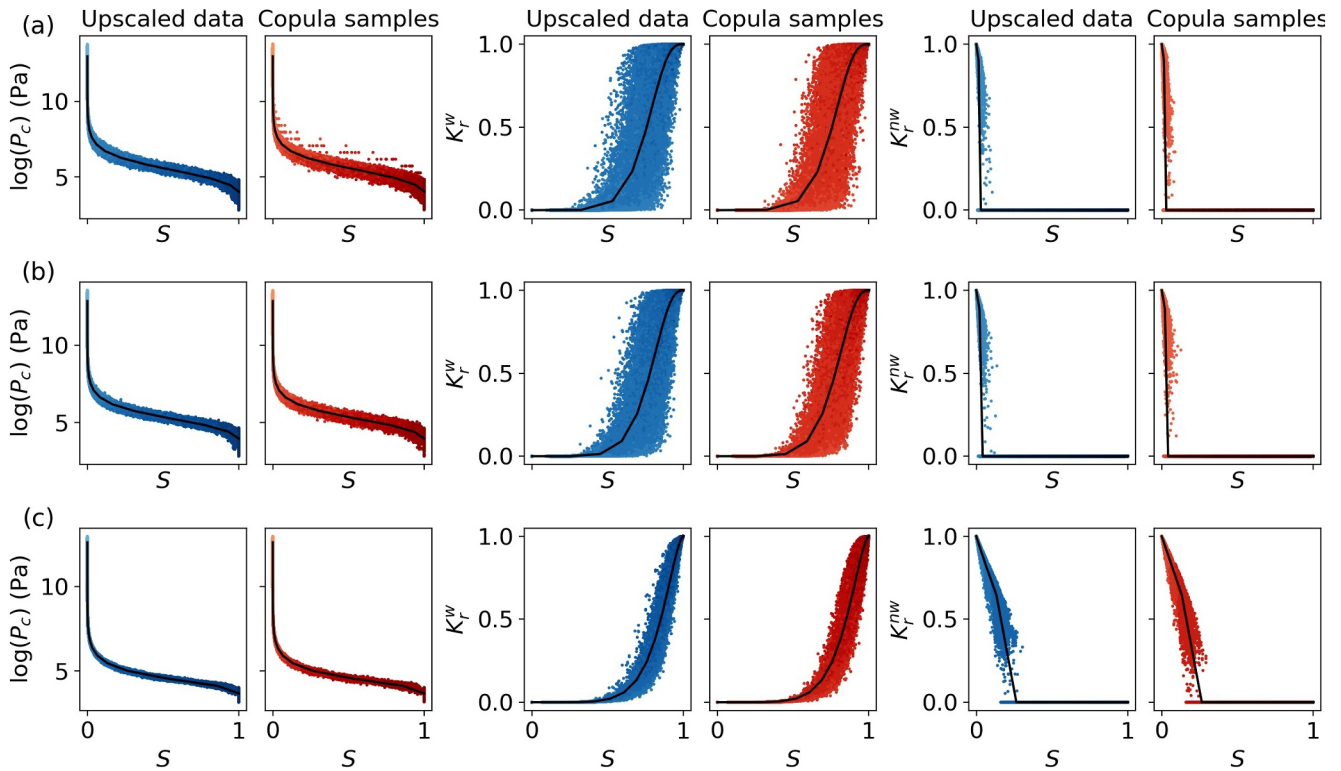


Figure 4. Upscaled and copula models of capillary pressure and relative permeabilities as function of saturation. The coloring of the data runs from low deterministic entry pressure (light blue/red) to high entry pressure (dark blue/red). The mean values of the data are shown in black. (a) $k_c = 0.0001$ mD (upper row), (b) $k_c = 0.001$ mD (middle row), and (c) $k_c = 1$ mD (lower row).

parametric copulas and instead a nonparametric transformation local likelihood (TLL) copula that uses a discretization based on data has been assigned. More details are provided in Appendix B.

It is of interest to illustrate the effect of representing the upscaled model in the form of copulas, and also to show the final stochastic flow functions in a form that is comparable with traditional deterministic modeling. Figure 4 plots the capillary pressure and relative permeability as functions of the saturation. The figure shows both data upscaled from fine-scale realizations and data sampled from the copula model, using different values of clay permeability k_c in Equation 2. We emphasize that since the samples are drawn independently from the two distributions (upscaled data vs. copula model), they will not exhibit a one-to-one relation. However, for a sufficiently large sample size, the generated point clouds should have a similar shape. We observe that, although there are some discrepancies, there is a good agreement between the results from the fine-scale and the copula model. The plots indicate that the stochastic flow functions take forms that are reminiscent of what would be expected from a deterministic model.

As a second illustration and as a measure of similarity between the underlying and the modeled distributions, we show estimated PDFs for the capillary pressure and wetting relative permeability, both as functions of saturation, for the three values of k_c . Contour plots of kernel density estimates of the PDFs of the upscaled data and equally many samples from the estimated copula and inverse Rosenblatt transformation of independent uniform samples are shown in Figure 5. The figure axes are scaled to include the whole range of observed data, which can be sparse in certain regions, leading to a localization of the estimated PDFs to a small subregion of the figure. Since the PDF is a point value, we consider a fixed value of $s_d = 0.5$, and note that a different value would imply a different PDF. For this value, the non-wetting relative permeability is zero, and PDFs are shown only for capillary pressure and wetting relative permeability as functions of saturation. Again, we observe a good agreement between the data generated from the fine-scale and the copula model. Moreover, a comparison of the results for the three k_c values shows that, both the capillary pressure and wetting relative permeability versus saturation functions have PDFs with relatively similar shape for the two smaller values of k_c . For $k_c = 1$ mD there are significant differences in the




Figure 5. Estimated probability density functions (PDFs) for $s_d = 0.5$ for the joint marginal distribution of upscaled capillary pressure and saturation, and upscaled wetting relative permeability and saturation, respectively. The PDFs are estimated using kernel density estimation from 10,000 samples. Expected values are marked with red crosses. (a) $k_c = 0.0001$ mD (upper row), (b) $k_c = 0.001$ mD (middle row), and (c) $k_c = 1$ mD (lower row).

shape of the PDFs and in their expected values (marked with red crosses in Figure 5). In particular, the wetting relative permeability versus saturation exhibits marked differences in the shape of the PDF for the different fine-scale facies properties, in particular due to the proximity to the upper bound $K_r^w \leq 1$ for $k_c = 0.001, 0.0001$ mD.

Figures 3–5 indicate good fits with respect to data of the copula model. Many UP methods take independent (and often uniform) random variables as inputs, and the properties of the output as a function of those inputs will typically determine the performance of the selected UP method. Before presenting the UP method, it is therefore instructive to visualize the five upscaled model outputs as functions of U_j , $j = 1, \dots, 5$. Figure 6 depicts all upscaled flow functions of all combinations of pairs of the first three independent uniforms $\{U_j\}_{j=1}^3$, where the remaining random variables are kept fixed at 0.5. The remaining seven combinations of all pairs of $\{U_j\}_{j=1}^5$ are exhibiting similar patterns, and are therefore not included. The results are shown for $k_c = 0.0001$ mD and $s_d = 0.5$ only. Other values of s_d would give different plots, but some general observations will be useful. The upscaled non-wetting relative permeability is identically zero for this value of s_d and hence displays no variance at all, whereas upscaled capillary pressure varies in all input variables. The wiggles in upscaled saturation are not spurious oscillations but a consequence of the nonlinear mapping $T_{s_d}^S$ in Equation 7. Although these qualitative observations do not make it apriori clear what an optimal UP method would be, they do indicate that a method that is insensitive to irregular features and that can exploit localization in random space may be a suitable choice for the problems of interest. Figure 6 indicates that there will be large regions of the input parameter space where the flow functions do not vary, and hence there would be no variation in an output quantity of interest. This can be exploited by adaptive sampling, placing fewer points in low-variability regions in favor of denser sampling in regions with large variability. Adaptive stratified sampling, to be described next, does not rely on regularity assumptions and is tailored to problems where adaptivity in stochastic space is possible, that is, with localization of variance. Hence, we believe it is an appropriate choice of method for the problems considered in the current work.

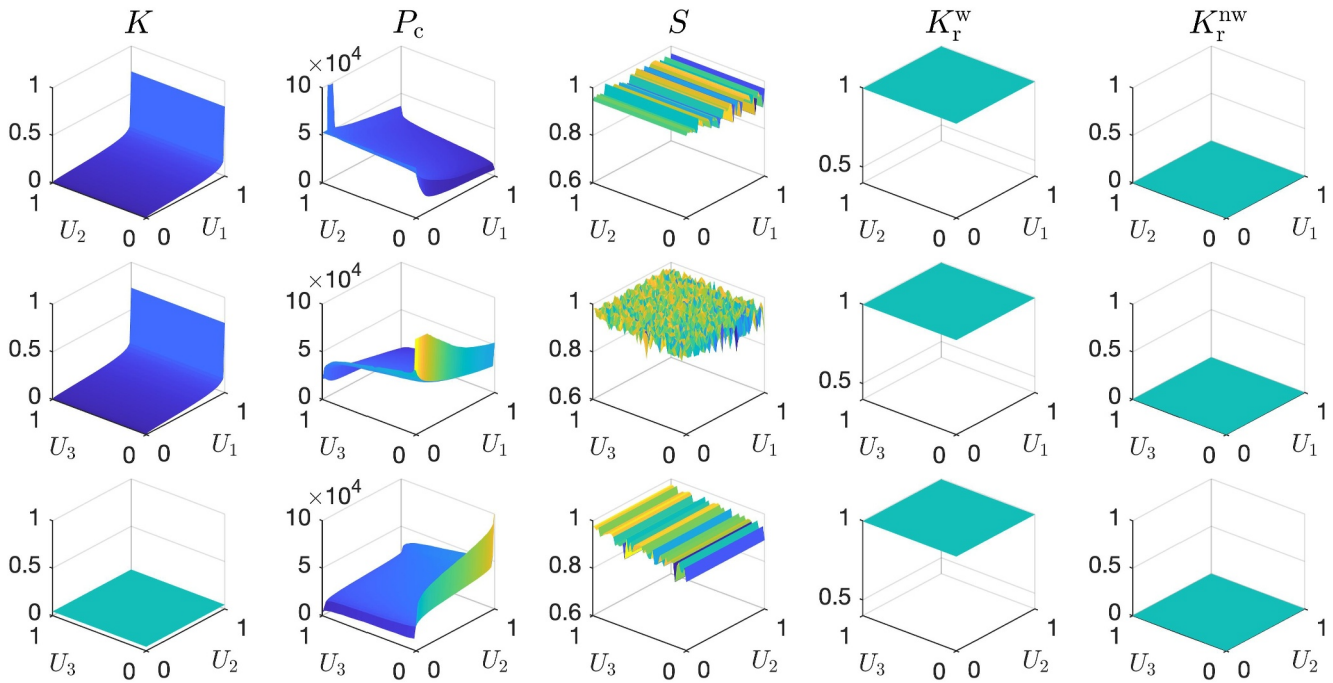


Figure 6. Upscaled flow functions displayed as functions of pairs from the set U_1, U_2, U_3 for $s_d = 0.5, k_c = 0.0001$ mD.

4. Adaptive Stratified Sampling for Uncertainty Propagation

We now briefly outline adaptive stratified sampling and refer the reader to Pettersson and Krumscheid (2022) for more details. Recalling Equation 1 and suppressing the dependence on physical deterministic parameters η for ease of notation, let $Q = M(Y(U)) \equiv \tilde{M}(U)$ denote the quantity of interest as a function of random input variables $U \in [0, 1]^n$ where we have assumed that all inputs have been transformed to the unit hypercube, for instance by means of the Rosenblatt transformation. In SMC sampling, the estimator of the expectation of a quantity of interest is approximated as

$$\mu_Q \equiv \mathbb{E}(Q) = \mathbb{E}(\tilde{M}(U)) \approx \hat{Q}_{MC} \equiv \frac{1}{N} \sum_{j=1}^N \tilde{M}(u^{(j)}), \quad (11)$$

where $\{u^{(j)}\}_{j=1}^N$ denotes a sample set of size N of standard pseudo-random samples in $[0, 1]^n$. The variance of the estimator Equation 11 can be reduced by partitioning the stochastic domain into disjoint strata, and sampling each stratum independently according to a suitable allocation rule to determine how the sample budget will be distributed between strata. Common allocation rules include proportional allocation, where the number of samples per stratum is proportional to the size of the stratum, and optimal allocation, that minimizes variance if the intra-stratum variances are known. A hybrid between proportional and optimal allocation provides a good compromise between robustness and optimality, and will be employed here.

Let S denote the stratification (partition) of the stochastic domain $[0, 1]^n$, with disjoint strata of probability measure p_S for all $S \in \mathcal{S}$. The probability measure or size is allowed to vary between strata. Then, if N_S^{prop} and N_S^{opt} denote respectively the number of samples allocated to stratum S with proportional and optimal sampling, hybrid allocation corresponds to the allocation rule

$$N_S^\alpha := (1 - \alpha)N_S^{\text{prop}} + \alpha N_S^{\text{opt}}, \quad (12)$$

for some choice of $\alpha \in [0, 1]$. The hybrid stratified sampling estimator is given by

$$\hat{Q}_\alpha := \sum_{S \in \mathcal{S}} \frac{p_S}{N_S^\alpha} \sum_{j=1}^{N_S^\alpha} Q_S^{(j)}, \quad (13)$$

for all $1 \leq j \leq N_S^\alpha$ and $S \in \mathcal{S}$.

The relative performance of two UP methods is measured by comparing their variances, assuming both methods are unbiased. The variance of the estimator \hat{Q}_α is given by

$$\mathbb{V}(\hat{Q}_\alpha) = \frac{1}{N} \sum_{S \in \mathcal{S}} \frac{p_S \sigma_S^2}{1 + \alpha(\bar{\sigma}_S - 1)}, \quad (14)$$

where $\bar{\sigma}_S \equiv \sigma_S / \sum_{T \in \mathcal{S}} p_T \sigma_T$. As a quantitative measure of the performance of adaptive stratified sampling relative to SMC, we define

$$\text{Speedup} \equiv \frac{\mathbb{V}(\hat{Q}_{\text{MC}})}{\mathbb{V}(\hat{Q}_\alpha)}. \quad (15)$$

The variance Equation 14 is only exact if the prescribed sampling rates Equation 12 are maintained, which may be difficult in practice. If that is not the case, for instance as a consequence of too few samples to compensate for previous non-ideal allocation, the resulting estimates may be inaccurate. As a remedy, the denominator in Equation 15 can be computed by simply computing the sample variance over repeated independent evaluations of the stratified sampling estimator. This procedure is robust but also expensive and is only motivated for demonstration purposes and cases where the model of interest is relatively cheap to sample from. In this work, we use the Python implementation of the adaptive stratified sampling method available at Krumscheid et al. (2023).

5. Physical Model

5.1. Numerical Reservoir Simulation Model

A simulator is used to compute the quantity of interest from a set of physical parameters and stochastic variables. The simulator solves a set of discretized equations for mass conservation for CO₂ and brine in a two-phase system of gas(g) and liquid(l) using a blackoil formulation

$$\begin{aligned} \frac{\partial \phi(b_l S_l)}{\partial t} + \nabla \cdot (b_l \mathbf{v}_b) &= 0, \\ \frac{\partial \phi(b_g S_g + r_s b_l S_l)}{\partial t} + \nabla \cdot (b_g \mathbf{v}_g + r_s b_l \mathbf{v}_l) &= 0, \end{aligned}$$

where ϕ is the rock porosity, r_s the dissolved gas-liquid ratio, and $b_{g,l}$ the inverse formation-volume factor (measuring the ratio between the bulk volumes of the phases occupied at surface and reservoir conditions). The saturation sum is $S_g + S_l = 1$ and the volumetric flux for phase $\alpha = g, l$ is given by the Darcy equation;

$$\mathbf{v}_\alpha = -K \frac{K_r^\alpha}{\mu_\alpha} (\nabla p_\alpha - \rho_\alpha \mathbf{g}),$$

where ρ_α is the fluid density, p_α is phase pressure, \mathbf{g} the gravity, μ_α the fluid viscosity. The phase pressures are related through the capillary pressure $P_c(S_l) = p_g - p_l$.

These equations are discretized on a grid using a two-point-flux approximation with upwinding for the spatial discretization, and backward Euler for the time discretization. We refer to Aziz and Settari (1979), Rasmussen et al. (2021) for details. The flux between two grid cells (a and b) can then be formulated in terms of a transmissibility T_{a-b} :

$$u_{\alpha, a-b} = -T_{a-b} \lambda_{\alpha}^{\text{up}} (p_{\alpha, a} - p_{\alpha, b} - \rho_{\alpha}^{\text{avg}} \mathbf{g}). \quad (16)$$

Here the mobility $\lambda_a^{\text{up}} = k_{ra}/\mu_a$ is evaluated in the upstream cell while the pressure is evaluated in the cell centers and an average value of the cells is used for the density ρ_a . The transmissibility between cell a and b is given by the harmonic mean of the one-sided transmissibilities such that

$$T_{a-b} = \frac{T_a T_b}{T_a + T_b}, \quad T_{a,b} = K_{a,b} A / D_{a-b}, \quad (17)$$

where $K_{a,b}$ is the permeability in cell a and b , respectively, and D_{a-b} is the distance from cell a to the center of the interface between the cells and from cell b to the center of the interface, respectively. The discretized equations are solved to an appropriate tolerance using a Newton-Raphson type method with iterative linear solvers. For further details on solvers and discretization, we refer to Rasmussen et al. (2021) and the OPM manual (Alvestad et al., 2021).

We need a set of boundary conditions and sink/source terms for the reservoir models, representing wells. For details on how wells are modeled in OPM Flow, we refer to the OPM manual (Alvestad et al., 2021). The default setup assumes no-flow boundaries, but flow through the boundaries can be included using various aquifer models, both based on simplified numerical models and analytical approximations. In our setup, we will use numerical aquifers to model the flow to aquifers connected through faults.

In the simulation model, dynamics within the numerical aquifers are simplified, and the aquifers are represented with only a few grid cells in the model that can be connected to any number of cells in the reservoir model using non-neighbor connections (NNC). Note that the framework itself allows for adding wells and more cells to the numerical aquifers if needed. Properties of the numerical aquifers like porosity, volume, permeability and connectivity can be set independently of the physical grid, which gives significant flexibility in usage. A very large porosity value can for instance be assigned to increase the pore volume of the aquifers, effectively disallowing pressure build-up in the aquifer.

The transmissibility used for the NNCs from the reservoir to the aquifer is given by the harmonic mean of the half transmissibility computed from the reservoir side, T_r and the transmissibility computed from the aquifer side, T_{aq} , as

$$T_{r-\text{aq}} = \frac{T_r T_{\text{aq}}}{T_r + T_{\text{aq}}}, \quad T_{\text{aq}} = 2K_f A_f / L_f. \quad (18)$$

The fault area, A_f , represents the area available for flow between the reservoir cell and the aquifer, and L_f , represents the length between the reservoir and the aquifer, while K_f is the effective permeability of the fault. Note that we here assume that the transmissibility computed for the aquifer side is given by the properties of the fault, which is valid as long as $K_f A_f / L_f \ll K_{\text{aq}} A_{\text{aq}} / L_{\text{aq}}$. For along-fault flow, this assumption is true for most range of permeability as $A_f \ll A_{\text{aq}}$. For cross-fault flow $A_f = A_{\text{aq}}$ and the assumption depends on $K_f \ll K_{\text{aq}}$.

Since the fault is modeled implicitly by the aquifer connections, we use the saturation function computed for the fault in the aquifer setup. As the relative permeability is evaluated in the upwind direction (Aziz & Settari, 1979), a second cell is added to the aquifer to assure the effect of the upscaled fault relative permeability is modeled in the setup. The aquifer thus consists of two cells, the first with pore volumes representing the fault core, and the second with a pore volume representing the aquifer itself.

5.2. Description of the Smeaheia Data Set

The starting point for the data set used in the simulation is the Smeaheia data set available at CO2Share (<https://co2datashare.org/>). To reduce the computational cost, we define a sector model representing a part of the reservoir by removing grid cells far from the injection point and the fault. This reduces the number of cells from 250,000 to 44,000, and leads to a simulation time of less than a minute for a single model realization of M . A typical grid cell in the model is 400 m times 400 m in the x and y direction, respectively, and between 1 and 20 m in the z direction. While this is very coarse, it is also typical for field-scale simulations. Using a coarse grid near the fault may introduce modeling error due to inaccurate representation of pressure gradients and upconing effects near the fault. Analytically derived model corrections or local grid refinement could remedy the modeling error (Kang

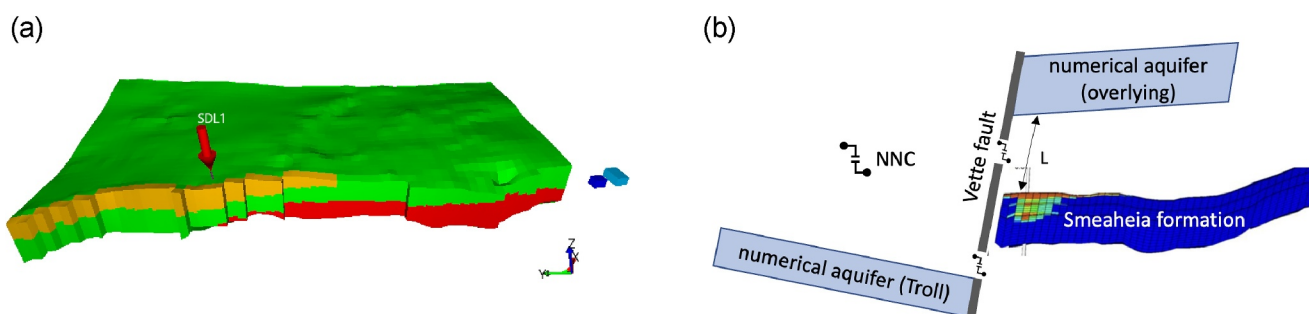


Figure 7. The section of the Smeaheia simulation model used in the study and the setup with non-neighbor connections to hydraulically connected zones. (a) Smeaheia section model with injector (red arrow). Orange shows the part of the Vette fault connected to the top aquifer, while red shows the connection to Troll. (b) Hydraulically connected zones are represented as numerical aquifers and coupled to the storage simulation via non-neighbor connections (NNC).

et al., 2014), but will not be considered in our approach. The reservoir consists of six horizontal layers with constant permeabilities. Three highly permeable sand layers with horizontal permeability 1,000 mD, 1,000 mD and 850 mD, respectively that are separated by two less permeable shale layers with horizontal permeability 50 mD and a less permeable bottom layer with horizontal permeability 25 mD. The vertical permeability is 1/10 of the horizontal permeability.

To include connections to nearby aquifers, the reservoir model is extended with two numerical aquifers, as illustrated in Figure 7. The first one, denoted Top aquifer (shown in orange in Figure 7), represents an aquifer above the Smeaheia model potentially connected vertically along the Vette fault, while the second (shown in red in Figure 7) represents the Troll field connected through Vette in the lower parts of the fault and is denoted Troll aquifer. To represent the highly complex dynamics of the Troll field with only one cell is clearly a crude approximation, which nevertheless should suffice to investigate the impact of Troll on the Smeaheia formation through uncertain fault parameters. An extra cell is added for the Top aquifer representing the fault in order to model the effect of having a different saturation function in the fault. For the Troll aquifer this is not needed as the connection is in the water zone and no two-phase effects are present.

There are two sets of saturation functions, that is, relative permeability and capillary pressure, one for the reservoir and one for the fault. Each is discretized and provided to the model as input tables. The reservoir saturation functions are computed using a standard Brooks-Corey function. For the fault, the tabulated values are obtained from realizations of the stochastic upscaled functions as computed in Sections 3.2 and 3.3, sampled with the adaptive stratified sampling approach described in Section 4.

The parameter set of the numerical aquifers are given in Table 1 and for details we refer to AQNUM in Alvestad et al. (2021). CO₂ is injected with a constant rate of 1.6 million tons per year for 59 years. The injection rate is the rate given in the Smeaheia data set. For simplicity, we do not consider CO₂ to dissolve into the brine phase. For the injection rate and timescale, this simplification is not expected to alter the results qualitatively.

6. Numerical Results

Numerical results will be presented using the stochastic model described in Section 3 with the following setup. A geocellular fine-scale model comprised of 20 geocells is employed for the Vette fault from 700 to 1,200 m depth. This model is represented by a total of 39 independent random parameters; 19 uniformly distributed uncertain

Table 1
Input Numerical Aquifers

Name	Area (m ²) (A_f)	Length [m] ($2L_f$)	Depth (m)	Poro (-)	Pres (bars)
Top Aquifer	40	2 · 250	*	1e12	*
Troll Aquifer	40	2 · 250	1,100	1e12	110

Note. The * indicates that the values are defaulted and thus computed internally in the simulator.

Table 2
Summary of the Test Cases

Case	Stochastic model	<i>n</i>
I	Fault permeability <i>K</i> from lognormal data fit	1
II	Fault permeability <i>K</i> as in Case I Layer permeabilities k_i ($i = 1, \dots, 6$) lognormal with parameters s.t. $\mathbb{E}(k_i) = \{1000, 50, 1000, 50, 850, 25\}$ mD, $\text{Std}(k_i) = 100$ mD.	7
III	$K, P_c(s_d), S(s_d), K_r^w(s_d), K_r^{nw}(s_d)$ from copula + Equations 5–9	5
IV	$K, P_c(s_d), S(s_d), K_r^w(s_d), K_r^{nw}(s_d)$ as in Case III K_{Troll} Lognormal data fit	6
V	Layer permeabilities as in Case II $K, P_c(s_d), S(s_d), K_r^w(s_d), K_r^{nw}(s_d)$ as in Case III	11
VI	Layer permeabilities as in Case II $K, P_c(s_d), S(s_d), K_r^w(s_d), K_r^{nw}(s_d)$ as in Case III K_{Troll} as in Case IV	12

Note. Only the stochastic parameters and functions are stated. All other quantities are assumed fixed at their deterministic reference values. The number of stochastic dimensions denoted by *n*.

geocell boundary locations, and 20 uncertain SGR values. Using the presented physical-stochastic upscaling framework results in coarse-scale two-phase flow parameters described by five dependent variables.

6.1. Test Cases

We consider six stochastic test cases, where we use the Smeaheia discretization described in Section 5.2. The test cases are presented in order of increasing complexity, but it should be noted that the simpler cases are really special cases of the more complex models, where certain parameters are kept fixed to the mean values. This allows us both to study the fault upscaling in detail, and then to introduce the interaction between fault and formation uncertainty.

In Case I and II, the permeability in the Vette fault is assumed uncertain and follows the distributions depicted in Figure 2. For practical sampling, lognormal distributions are fitted to the data for the different settings of k_c . In Case I all other parameters including the two-phase flow functions are assumed deterministic and set to their respective expected values (red curves in Figure 3), that is, nominal reference values. In Case II we assume that each of the six layers of the reservoir model has lognormal homogeneous permeability, given by $k_i \sim \text{Lognormal}(\mu_{k_i}, \sigma_{k_i})$ for $i = 1, \dots, 6$. The distribution parameters μ_{k_i}, σ_{k_i} are chosen so that the expectations of the layer permeabilities are, respectively, 1000, 50, 1000, 50, 850, 25 mD with the standard deviations for simplicity set to 100 mD for all layers. In total, this model has seven independent random input parameters.

For the remaining Cases III–VI, we assume stochastic two-phase flow functions (permeability, capillary pressure, saturation and wetting/non-wetting relative permeabilities) as represented by the copula model. Case III assumes all other parameters to be deterministic, while Case IV is supplemented by an upscaled model for the uncertain permeability in the connection with Troll. This permeability is assumed to be independent of the Vette fault parameters, and computed using the same SGR-permeability mapping Equation 2. The mean SGR values used to populate the geocellular model are given by the synthetic data in the depth range from 1,500 to 1,600 m shown in Figure 2, and the standard deviations are assumed to be 14% as for the Vette fault. Arithmetic averaging over the facies is performed to accommodate for flow across the fault.

In Case V we consider the upscaled flow functions with the copula model in combination with the six reservoir layers following the distributions described in connection with Case II. Finally, in Case VI, we combine the random models from Case V (upscaled flow functions and reservoir layers) with uncertainty in the Troll

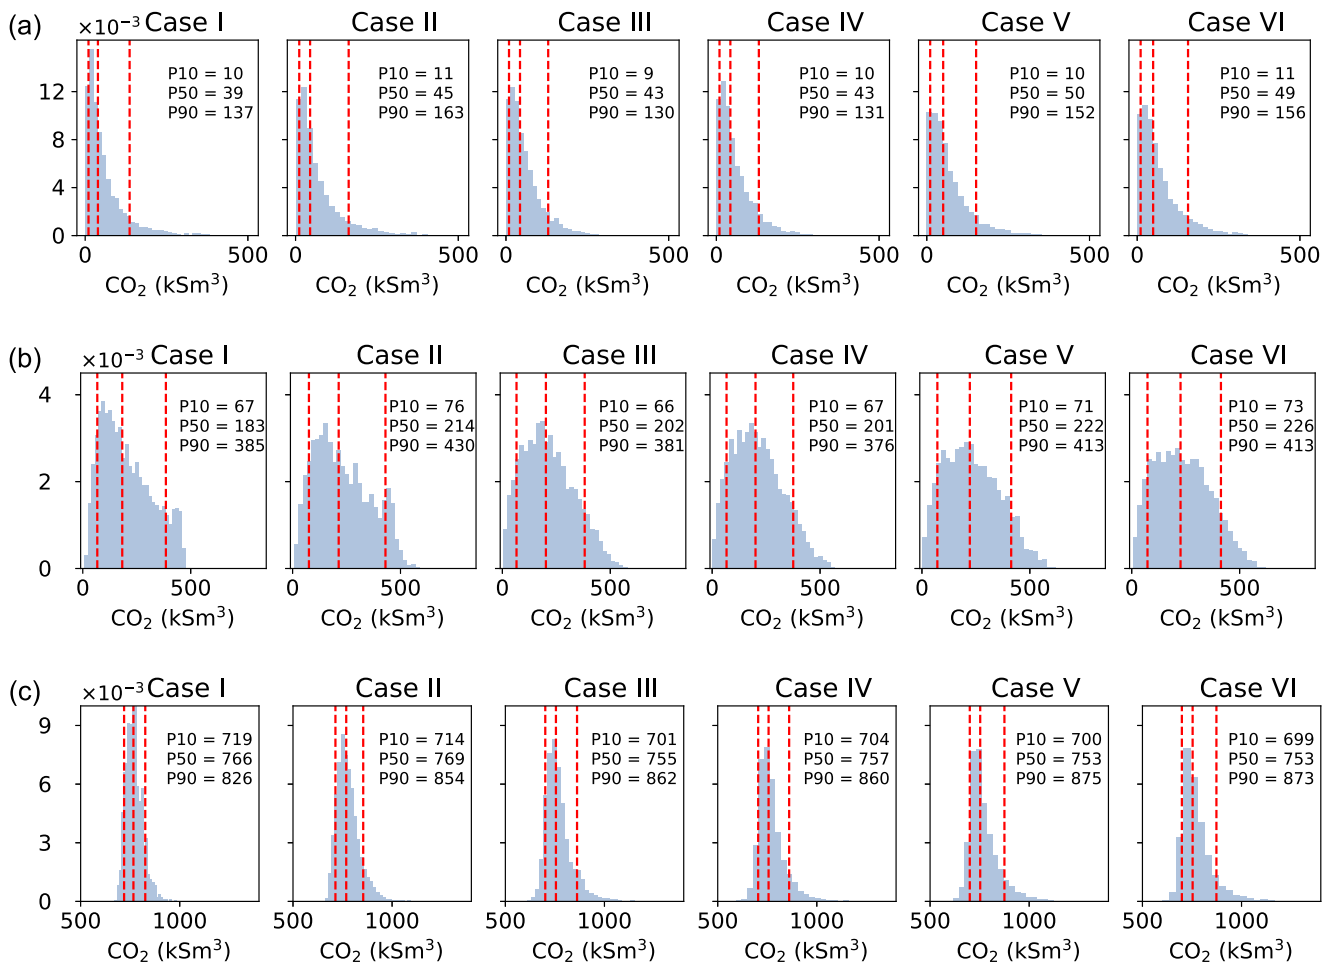


Figure 8. Histograms of leaked CO_2 for Case I–VI. The three vertical dashed lines in each subfigure indicate the 10th percentile (left), median (middle), and 90th percentile (right). (a) $k_c = 0.0001$ mD, (b) $k_c = 0.001$ mD, and (c) $k_c = 1$ mD.

connection permeability with the model described in Case IV. The setups of the test cases are summarized in Table 2.

6.2. Monte Carlo Simulation

The leakage distributions for the six test cases estimated from 5000 SMC samples and visualized as histograms are shown in Figure 8 for $k_c = 0.0001$ mD, $k_c = 0.001$ mD and $k_c = 1$ mD, respectively. As expected, the lower shale permeability leads to less amount of CO_2 leaked. The probability that there is no leakage at all is below 0.01 for the case $k_c = 0.0001$ mD, and identically zero for larger values of k_c . For all clay permeability cases, the setup Case I differs from Cases II–VI in capturing the upper tail of the probability distributions. As a quantitative characterization of the leakage distribution, each figure includes the percentiles P10, P50 (median), and P90. More extreme percentiles (e.g., P99) are not included, as they would not be accurately estimated with the current number of samples. The variability of the distributions is larger between the results for different k_c than between the six cases of different stochastic models. This suggests that improved SGR models, or modeling k_c as a random variable, may yield more accurate predictions of the leakage distributions. The limited variability between the six cases of different stochastic models indicate that the simpler models are sufficient to obtain the statistics investigated for the current problem setup. However, a physically more refined model, or some other choice of statistics (e.g., rare events), may benefit from the more complex models involving the uncertainty in two-phase flow functions.

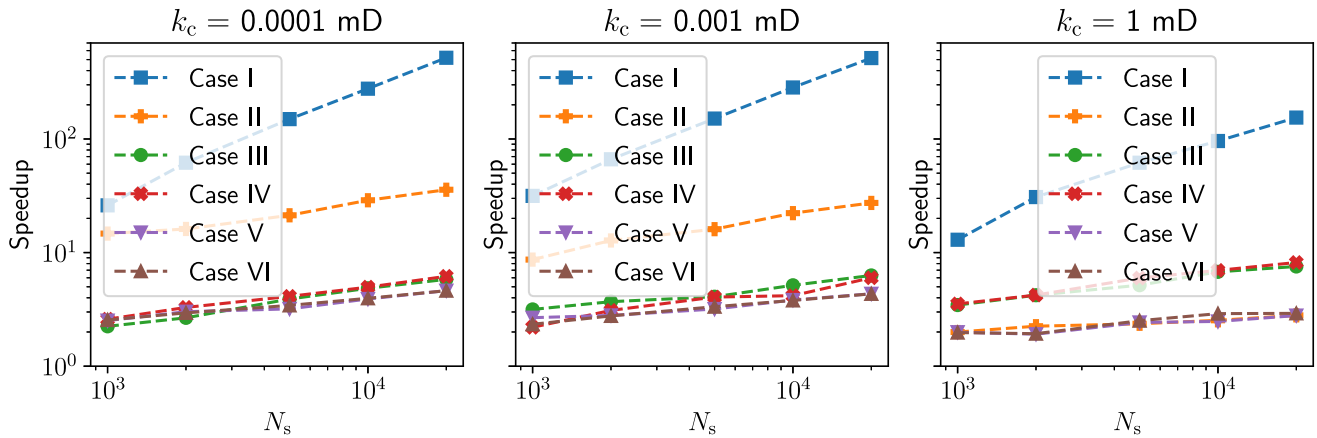


Figure 9. Estimated speedup of ADSS with respect to standard Monte Carlo.

To put the results into perspective, the emissions can be compared to the CO₂ emissions of a dairy cow. For a single dairy cow, we assume an average of 6 m³ daily emissions of CO₂, based on the measured average 6.137 m³ presented in Kinsman et al. (1995). The most extreme leakages observed in Figure 8 are around 1,000,000 m³ at standard surface conditions and correspond to a total leakage of 0.002% of the injected mass, or the emissions of less than 8 dairy cows for the same period of time, that is, 59 years.

6.3. Adaptive Stratified Sampling Results

Next we apply the ADSS method described in Section 4 to the six test cases to compute the expected leakage using Equation 13 with the hybrid sampling parameter $\alpha = 0.5$, and allocation of 50 samples per iteration. These design parameter settings were found to be suitable for different test problems in Pettersson and Krumscheid (2022), but have not been tailored to the test problems in the current paper. The speedups of the ADSS method compared to SMC as given by Equation 15 for all test cases and varying number of sampling budgets are shown in Figure 9. For the simplest Case I, we observe two orders of magnitude speedup. For Case II with seven independent random variables, the speedup is one order of magnitude for the SGR models with $k_c = 0.0001, 0.001$ mD, but significantly lower for $k_c = 1$ mD. For all other cases, the speedups vary from 2 to 4 for the smallest sample sizes, to 2–8 for the largest sample sizes. The lower speedups are observed for the Cases III–VI where the relatively complex two-phase model Equations 5–9 is employed, with its sample path crossings of S and K_r^w , and random variable dependence described by the vine copula model as discussed in Section 3.5. For $k_c = 1$ mD, we observe higher speedups for Cases III–IV than V–VI, which is to be expected given the difference in the number of random dimensions. Any localization of variance is more difficult to resolve for ADSS in high-dimensional spaces. Taken together, the speedup results indicate that the increased complexity of ADSS is motivated by the reduction in sample sizes possible while maintaining similar accuracy as SMC.

6.4. Discussion

We present a framework for uncertainty quantification for complex physical systems, for example, subsurface CO₂ storage, where the focus is on accurate and efficient representation of uncertainty from available data, and a subsequent uncertainty propagation through the physical model via adaptive sampling. If some input parameters to the physical model are considered uncertain but lacking in data, they should nevertheless be included as random variables, for example, uniform random variables on an interval including possible physical variation. The proposed uncertainty modeling and propagation framework has been applied to a test case from the Smeaheia formation in the North Sea, but should be applicable to a wider class of problems. Uncertainty in the physical model itself (in the current case the reservoir model) has not been considered in this work.

A geocellular fault model with uncertain properties modeled as random variables has been upscaled, resulting in stochastic effective flow functions. Faults acting as conduits or as barriers can equally well be represented

via the proposed upscaling approach. A conduit (barrier) fault will result from populating the fine-scale facies by samples from an SGR distribution in a higher (lower) regime, preferably fitted to empirical data. We emphasize that the assumed random distributions of the fine-scale fault model, that is, uniformly distributed sizes of geocells and Gaussian SGR values of each geocell, could and should be substituted with real data when possible, without requiring any further changes of the proposed algorithmic framework. The permeability of the clay-rich shale k_c should also be based on measurements to result in realistic upscaled fault permeability. Since the geocellular model can consist of several layers parallel to the fault, it is possible to explicitly model a low-permeability fault core surrounded by high-permeability damage zones. If data are available, a high degree of space-dependent variability can be incorporated into the model. However, such a model would require modification of the upscaling method to account for both vertical and horizontal flow. There are indeed limits on the performance of any semi-analytical upscaling method when it comes to accurately estimate effective properties.

All upscaled flow functions except fault permeability are functions of a deterministic parameter and need to be modeled as random fields, in addition to the copula model used to represent the dependencies at some fixed value of the deterministic parameter. Karhunen-Loeve expansions, commonly used for this purpose, fail to capture physical features such as monotonic sample paths. Instead, we use data to fit an empirical correlation model that does not assume stochastic stationarity, that is, the two-point correlation is allowed to change with the deterministic model parameter. The proposed model does not rely on any smoothness assumptions about the underlying random fields we seek to model. Indeed, it is data-driven and returns discrete values at user-selected points, whether these come from a smooth process or one with jumps. Hence, it is applicable to a wider class of data. All observed realizations of the resulting model are physically meaningful, although we cannot prove that the model excludes unphysical results.

Dependencies between the random variables in the models are due to physical relations between the underlying quantities, many of which can not easily be described by closed-form functional relationships. Correctly modeling dependencies is hence crucial to accurately capture the physics, and copulas can do exactly that. Capturing the dependencies present in the multivariate distributions representing the upscaled model flow functions relies on existing methodology (and, practically speaking, software) for multidimensional copula fitting, effectively decomposing the problem into separate treatment of marginal distributions and dependence structure. The considered copula approach offers flexibility, providing a good number of distinct pair copulas that can be fitted to high accuracy. In particular, PDFs concentrated in the vicinity of the physical boundary of a parameter range (relative permeabilities and saturation are all confined to the unit interval) are well represented without any visible spurious boundary effects. Indeed, the fitted copula models (for three different values of clay permeability) share the same tree structures, but the bivariate copula functions connecting the nodes differ somewhat, reflecting the differences in pairwise PDFs that have been observed. There are no intuitive and standardized measures for assessing proximity between multivariate distributions, but examining the 1D and 2D projections of the data and the model at least gives an idea of the model fit, and they do agree well. Data-availability may be a limitation for real-world copula models with significantly larger number of random variables than those considered here for the upscaled flow functions. Sparse copula models could be an option to tackle both the high numerical cost of fitting a large number of pair copulas, and avoid overfitting.

In addition to accurately representing empirical distributions of data, the upscaled model allows targeted model evaluation, that is, directed sampling within a given region of random space. In SMC sampling, it is sufficient to be able to draw any random samples as long as they follow prescribed distributions honoring data. The adaptivity of the adaptive stratified sampling method employed here, in addition requires that the stochastic input domain can be sampled from a given subdomain. Such subdomain sampling in a reduced and uniform space is made possible by the copula model and inverse Rosenblatt transformation. Many other groups of uncertainty propagation methods, including Quasi-Monte Carlo, Multi-Level Monte Carlo, and spectral expansions (generalized polynomial chaos) would either require or greatly benefit from these properties. Hence, the proposed stochastic model with copulas is a contribution to uncertainty quantification by being amenable to adaptivity, employing a low-dimensional stochastic representation, and capturing dependencies while providing the user the option to only use standard independent random variables as inputs. Such a model offers practitioners more accurate

uncertainty estimates of geological properties than a model that just assumes standard distributions instead of purely empirical distributions whenever possible.

Adaptive stratified sampling requires significantly fewer samples than SMC to reach a given error threshold and is relatively insensitive to the number of stochastic dimensions, in particular if the number of effective dimensions is small (i.e., the total variance is dominated by a small subset of random dimensions). The design parameters set by the user, that is, the number of samples allocated in every iteration and the proportion of samples optimally allocated (as opposed to the more robust proportional allocation), will affect the observed speedups. In this work, a relatively conservative approach has been used, with as many as on average 50 samples per stratum allocated every iteration, and only a fraction $\alpha = 0.5$ of samples optimally allocated. While any exact recommendation for these design parameters is hard to motivate, the chosen numbers appear suitable given the complex features of the Cases III-VI, which are illustrated in Figure 6. The QoI, that is, the expected leakage of CO₂, was chosen as a natural starting point for investigation of leakage risk. Other QoI could have been selected instead, for instance the probability of exceeding some predefined leakage threshold. The observed speedups are specific to the particular QoI, and one would need to rerun the algorithm for different QoIs, which is a common feature shared by most accelerated methods. In addition to providing a relative computational cost estimate compared to SMC, the speedup estimates can also be used to estimate the accuracy of the results. The mean-squared error of SMC decays linearly in the number of samples with an error constant given by the variance of the QoI, and since the effective number of samples of the ADSS method compared to SMC is the speedup multiplied with the actual number of samples, it provides a very direct error control.

Accelerated methods including adaptive stratified sampling provide improvement over SMC with respect to some user-defined metric, for example, a mean value. SMC is hence more amenable to general analysis of quantities of interest, and was chosen to produce estimates of the distributions of CO₂ leakage. The histograms of leaked CO₂ obtained by SMC indicate the possibility of long upper tails, that is, small but nonzero probability for relatively large leakage. If these probabilities are of interest, a rare-event simulation framework may replace the current method for UP. The number of samples required is otherwise inversely proportional to the probability of the outcome of interest, so very unlikely scenarios are correspondingly difficult to correctly estimate.

The numerical results show that, with overwhelmingly large probability, the leakage of CO₂ is less than 0.002% of the injected volume. This negligible amount of CO₂ (corresponding to less than the emissions of 8 dairy cows (Kinsman et al., 1995)) obtained for all 18 test cases, indicate that refined models are not crucial due to the low risks involved. Surely, other prospective storage formations with increased risk for leakage may require more elaborate models.

The fault-related uncertainties have a higher impact than the reservoir uncertainties, as demonstrated by the fact that the empirical distributions do not change much when the stochastic reservoir model is added to the fault model. The high variability in distributions across the different values of clay permeability strengthens the conclusion that the fault uncertainties dominate the total leakage distribution. This suggests that a first stochastic and physical model improvement would be to base the value of clay permeability on actual data, preferably modeled as an empirical stochastic distribution to account for uncertainties due to limited seismic resolution.

Accurate uncertainty quantification relies on a similar level of accuracy of the physical model describing the complex system of interest. In the current work, this includes both the fault model and the model for the surrounding reservoir. The fault model considered in Section 3 assumed a 1D description of the fine-scale geometry of the fault. In reality, there can be significant horizontal variations both along and across the fault, thus in the general case it could be necessary to represent the fine-scale geometry as truly 3D. In such a setting, the fault core and damage zones could be explicitly represented. While we expect that the overall stochastic modeling framework to be still applicable, ingredients of the workflow will need to change: The upscaling must be based on numerical rather than analytical methods and since the upscaled stochastic ensembles may take different forms, the stochastic modeling approach could also need adaptations. The effect of the capillary equilibrium assumption needs to be quantified by taking into account the relative effect of viscous forces on the flow compared to the effect of capillary forces. Also, the effect of neglecting potential fault reactivation due to induced stress should be investigated quantitatively for the fault zone under consideration. Otherwise, the CO₂ leakage may be systematically underestimated.

The simulation framework successfully produced meaningful estimates of expected fault-related leakage for a prospective storage site in the Norwegian North Sea with relatively little computational overhead. It should thus be feasible to implement this approach in fault de-risking workflows used in an industrial context, that is, it is not simply a theoretical exercise. One notable observation is that the leakage estimates obtained in the 18 simulated cases of the Smeaheia test case range from 80 to 1,500 tons cumulative leaked CO₂ over 60 years. This amounts to very tiny amounts of CO₂ leakage that would be below the detection limits of modern subsea sensing equipment. We expect that these leakage estimates along the Vette fault are indeed representative given that the underlying SGR map, reservoir model, and other parameters are based on actual data. However, more work needs to be done to ensure that the simplifying approximation in the dynamic simulations (i.e., NNC connections and grid resolution) gives a reasonable first-order approximation of the fluid flow in the fault. Refining the fault model by adding more virtual (aquifer) cells and connections (NNCs) is possible, but needs to be balanced by the availability of the underlying geological data. In addition, the results may be sensitive to the simplicity of the fault model (already discussed above) or require other stochastic simulation tools to capture the upper tails of the leakage distributions. A final point to consider is the relative coarseness of the grid in the used Smeaheia reservoir model. Refining the grid close to the fault and the well is expected to improve the accuracy of the solution, but at the cost of higher simulation time. In summary, the exact histograms and mean values derived in this study are for demonstration purposes and should not be taken out of context.

7. Summary and Conclusions

We have presented a numerical framework for upscaling of uncertain fault properties, uncertainty modeling including complexity reduction, and adaptive uncertainty propagation, and demonstrated numerical results for a real-world CO₂ storage test case. The combined physical and stochastic upscaling leads to decreased deterministic model evaluation times and smaller number of random dimensions. To exploit the stochastic model reduction, the reduced set of random variables are parameterized using an accurate copula model that captures both their dependence and their marginal distributions. The copula model is then combined with an empirical (data driven) model for the non-stationary correlation in deterministic parameter space, where standard methods such as Karhunen-Loeve expansions fail to produce physically relevant flow function models.

The numerical results show that the leakage of CO₂ is negligible in all cases investigated. The most significant uncertainties with respect to expected leakage are those related to the Vette fault properties, which dominate over those related to reservoir properties and the connection to Troll. In particular, the large effect from the value of the clay permeability suggests that a natural first improvement of the suggested model is to include a data-driven distribution for the clay permeability. This, and the same for other variables for which data become available, should be done within the copula framework to properly account for dependencies and marginal distributions.

To reduce the number of physical model evaluations to obtain mean values of leaked CO₂, ADSS has been used instead of direct SMC. Compared to SMC, for the sample size ranges that have been investigated, the computational budget for ADSS can be reduced with up to a factor of, respectively, 10 and 100 for the simpler stochastic cases, and 2–8 times for the copula models with nonlinear transformations. The speedups are sufficient to motivate ADSS instead of SMC. In summary, the successful demonstration of the methodology is promising for future implementation of the framework in a real setting with more site-specific data available.

Appendix A: Two-Phase Upscaling

The two-phase upscaling applied in Section 3.2 following Rabinovich et al. (2016) can be summarized as follows:

1. Calculate the upscaled absolute permeability K from Equation 3.
2. Define a set of capillary pressure values, to be used to calculate the fluid distribution under capillary equilibrium, see comments below for how to choose this set.
3. Pick a value from the set of capillary pressure values, assign this to all fine-scale facies (this is the assumption of capillary equilibrium).

4. Use an inverted fine-scale capillary pressure function to calculate a fine-scale saturation. If the fine-scale capillary pressure function is equal in all facies, the fine-scale saturation will be uniform.
5. Calculate an upscaled saturation using arithmetic averaging.
6. From the fine-scale saturation distribution, compute the upscaled effective phase-wise permeabilities, K^α , $\alpha \in \{w, nw\}$, as the harmonic average of the product of the fine-scale absolute and relative permeability (respectively k_s and k_r^α) using Equation 3.
7. Obtain the upscaled relative permeability as $K_r^\alpha = K^\alpha / K$, $\alpha \in \{w, nw\}$.

Points 3–7 are repeated over the full set of capillary pressure values defined in point 2; we see that points 3 and 5 together give an upscaled capillary pressure curve, while points 4 and 6 produce upscaled relative permeabilities. In practice, point number 2 is a bit tricky since, first, the Brooks-Corey capillary pressure is not bounded at zero saturation, and second, the fine-scale capillary function is stochastic through its entry pressure, see Equation 4, and varies between the fine-scale facies. To achieve a robust implementation, the equilibrium values for capillary pressure are chosen as follows: We sample a set of saturation values, denoted $s_d \in [1 \times 10^{-6}, 1]$ where the cut-off is set slightly above 0 to avoid sampling an unbounded capillary pressure. The sampling is done in logarithmic space to achieve good coverage in the steep region of the Brooks-Corey curve. The equilibrium capillary pressures in the set defined in point number 2 is calculated from these values of s_d , using the minimal entry pressure of the facies in the current realization, thereby ensuring that the equilibrium pressure can be reached in all facies.

Appendix B: Copula Modeling

B1. Copula Tree Fit

The tree structures defining the fitted copulas are shown in Figure B1. The nodes in the trees denote random variables, where the random variables are represented by their indices (e.g., 1,2 is short for Y_1, Y_2). The edge connecting two adjacent nodes denotes the joint distribution conditioned on the subset of repeated variables of the two nodes. The edges of Tree t are the nodes of Tree $t + 1$, so in this sense the structure of all subsequent trees can be inferred from Tree 1. The choice of copula function, among a suitable candidate set of copula families, for any pairs of variables is inferred by data. The PDF of the five-dimensional distribution is the product of all Equation 10 bivariate distributions represented by the edges of the trees in Figure B1 and the (five) marginal distributions. The bivariate copulas that appear in Figure B1 are the nonparametric TLL copula (Geenens et al., 2017), and common parametric function copulas (BB1 (Joe, 1997), Gumbel, Student-t, and Gaussian) that may in some cases be rotated by 90° , 180° , or 270° for best fit with data. The copula model of the case $k_c = 0.0001$ mD has the same tree structure as shown in Figure B1, but with bivariate copulas (from left to right) as follows. Tree 1: TLL, BB7 180, Gaussian, TLL; Tree 2: TLL, TLL, TLL; Tree 3: TLL, BB8; Tree 4: BB7 270° .

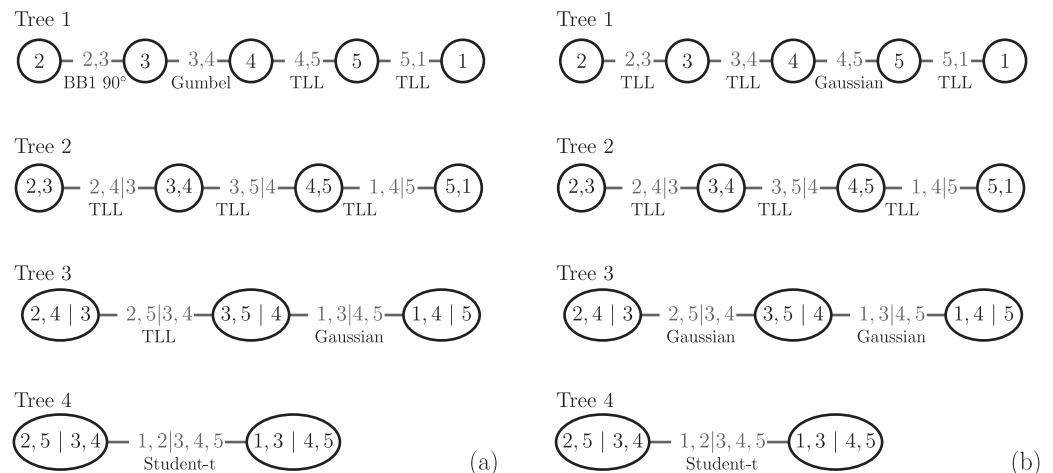


Figure B1. Tree structure of the fitted copulas. This tree structure is known as a D-vine copula. (a) $k_c = 1$ mD, and (b) $k_c = 0.001$ mD.

Data Availability Statement

The software used to generate the data, numerical methods, and figures are available from Zenodo (Pettersson et al., 2023). This includes scripts for stochastic upscaling using synthetic data, copula fits, random field parameterizations, and forward uncertainty propagation through OPM Flow using SMC and ADSS. Python scripts using Matplotlib are available to generate Figures 2–5 and 8, and 9, and a Matlab script to generate Figure 6 (tested with Matlab R2021a).

Acknowledgments

The work was funded by the Research Council of Norway through the projects Quantification of fault-related leakage risk (FRISK) under project number 294719, and Expansion of Resources for CO₂ Storage on the Horda Platform (ExpReCCS) under project number 336294.

References

- Aas, K., & Berg, D. (2009). Models for construction of multivariate dependence – A comparison study. *The European Journal of Finance*, 15(7–8), 639–659. <https://doi.org/10.1080/13518470802588767>
- Alcalde, J., Flude, S., Wilkinson, M., Johnson, G., Edlmann, K., Bond, C. E., et al. (2018). Estimating geological CO₂ storage security to deliver on climate mitigation. *Nature Communications*, 9(1), 2201. <https://doi.org/10.1038/s41467-018-04423-1>
- Alvestad, J., Baxendale, D., Bao, K., Blatt, M., Hove, J., Lauser, A., et al. (2021). OPM flow reference manual (2021-10) [Computer Software Manual]. https://opm-project.org/?page_id=955
- Aziz, K., & Settari, A. (1979). *Petroleum reservoir simulation*. Springer. <https://doi.org/10.2118/9781613999646>
- Babuška, I., Nobile, F., & Tempone, R. (2007). A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 45(3), 1005–1034. <https://doi.org/10.1137/100786356>
- Barbosa, C. H. S., Kunstmann, L. N. O., Silva, R. M., Alves, C. D. S., Silva, B. S., Filho, D. M. S., et al. (2020). A workflow for seismic imaging with quantified uncertainty. *Computers & Geosciences*, 145, 104615. <https://doi.org/10.1016/j.cageo.2020.104615>
- Bedford, T., & Cooke, R. M. (2002). Vines—a new graphical model for dependent random variables. *Annals of Statistics*, 30(4), 1031–1068. <https://doi.org/10.1214/aos/1031689016>
- Berge, R. L., Gasda, S. E., Keilegavlen, E., & Sandve, T. H. (2022). Impact of deformation bands on fault-related fluid flow in field-scale simulations. *International Journal of Greenhouse Gas Control*, 119, 103729. <https://doi.org/10.1016/j.ijggc.2022.103729>
- Bhatti, M. I., & Do, H. Q. (2019). Recent development in copula and its applications to the energy, forestry and environmental sciences. *International Journal of Hydrogen Energy*, 44(36), 19453–19473. <https://doi.org/10.1016/j.ijhydene.2019.06.015>
- Bjørnå, T. I., Haines, E. M., & Skurtveit, E. (2021). Upscaled geocellular flow model of potential across-and along-fault leakage using shale gouge ratio. In *Proceedings of TCCS-11, Trondheim conference on CO₂ capture, transport and storage*.
- Bjørnå, T. I., Skurtveit, E., Michie, E., & Smith, S. A. (2022). Overburden fluid migration along the Vette Fault Zone, North Sea, using different fault permeability models. In *Sixth international conference on fault and top seals* (pp. 1–5). European Association of Geoscientists & Engineers. <https://doi.org/10.3997/2214-4609.202243033>
- Bjørnå, T. I., Skurtveit, E., Michie, E. A. H., & Smith, S. A. (2023). Characterizing along-and across-fault fluid-flow properties for assessing flow rates and overburden fluid migration along faults: A case study from the North Sea. *Petroleum Geoscience*, 29(3), petgeo2023-033. <https://doi.org/10.1144/petgeo2023-033>
- Braathen, A., Tveranger, J., Fossen, H., Skar, T., Cardozo, N., Semshaug, S., et al. (2009). Fault facies and its application to sandstone reservoirs. *AAPG Bulletin*, 93(7), 891–917. <https://doi.org/10.1306/032309081116>
- Brechmann, E. C., & Schepsmeier, U. (2013). Modeling dependence with C- and D-vine copulas: The R package CDVine. *Journal of Statistical Software*, 52(3), 1–27. <https://doi.org/10.18637/jss.v052.i03>
- Brooks, R. H. (1965). *Hydraulic properties of porous media*. Colorado State University.
- Caine, J. S., Evans, J. P., & Forster, C. B. (1996). Fault zone architecture and permeability structure. *Geology*, 24(11), 1025–1028. [https://doi.org/10.1130/0091-7613\(1996\)024<1025:FZAAPS>2.3.CO;2](https://doi.org/10.1130/0091-7613(1996)024<1025:FZAAPS>2.3.CO;2)
- Chen, L., & Guo, S. (2019). *Copulas and its application in hydrology and water resources*. Springer. <https://doi.org/10.1007/978-981-13-0574-0>
- Childs, C., Walsh, J. J., Manzocchi, T., Strand, J., Nicol, A., Tomasso, M., et al. (2007). Definition of a fault permeability predictor from outcrop studies of a faulted turbidite sequence, Taranaki, New Zealand. *Geological Society, London, Special Publications*, 292(1), 235–258. <https://doi.org/10.1144/SP292.14>
- Childs, C., Walsh, J. J., & Watterson, J. (1990). A method for estimation of the density of fault displacements below the limits of seismic resolution in reservoir formations. In A. T. Buller, E. Berg, O. Hjelmeland, J. Kleppe, O. Torsæter, & J. O. Aasen (Eds.), *North Sea oil and gas reservoirs—ii* (pp. 309–318). Springer. https://doi.org/10.1007/978-94-009-0791-1_26
- Crippa, M., Guizzardi, D., Solazzo, E., Muntean, M., Schaaf, E., Monforti-Ferrario, F., et al. (2021). *GHG emissions of all world countries*. Publications Office of the European Union. <https://doi.org/10.2760/074804>
- Czado, C., & Nagler, T. (2022). Vine copula based modeling. *Annu. Rev. Stat. Appl.*, 9(1), 453–477. <https://doi.org/10.1146/annurev-statistics-040220-101153>
- Elsheikh, A. H., Oladyshkin, S., Nowak, W., & Christie, M. (2014). Estimating the probability of CO₂ leakage using rare event simulation. ECMOR XIV-14th European conference on the mathematics of oil recovery, 2014, 1–9. <https://doi.org/10.3997/2214-4609.20141876>
- Étoré, P., Fort, G., Jourdain, B., & Moulines, E. (2011). On adaptive stratification. *Annals of Operations Research*, 89(1), 127–154. <https://doi.org/10.1007/s10479-009-0638-9>
- Faleide, T. S., Braathen, A., Lecomte, I., Mulrooney, M. J., Midtkandal, I., Bugge, A. J., & Planke, S. (2021). Impacts of seismic resolution on fault interpretation: Insights from seismic modelling. *Tectonophysics*, 816, 229008. <https://doi.org/10.1016/j.tecto.2021.229008>
- Fan, Z., Eichhubl, P., & Newell, P. (2019). Basement fault reactivation by fluid injection into sedimentary reservoirs: Poroelastic effects. *J. Geophys. Res. Solid Earth*, 124(7), 7354–7369. <https://doi.org/10.1029/2018jb017062>
- Feng, R., Grana, D., & Balling, N. (2021). Uncertainty quantification in fault detection using convolutional neural networks. *Geophysics*, 86(3), M41–M48. <https://doi.org/10.1190/geo2020-0424.1>
- Fisher, Q. J., & Jolley, S. J. (2007). Treatment of faults in production simulation models. *Geological Society - Special Publications*, 292(1), 219–233. <https://doi.org/10.1144/SP292.13>
- Flemisch, B., Darcis, M., Erbertseder, K., Faigle, B., Lauser, A., Mosthaf, K., et al. (2011). Dumux: Dune for multi-{phase, component, scale, physics,...} flow and transport in porous media. *Advances in Water Resources*, 34(9), 1102–1112. <https://doi.org/10.1016/j.advwatres.2011.03.007>
- Freeman, S. R., Harris, S. D., & Knipe, R. J. (2008). Fault seal mapping – Incorporating geometric and property uncertainty. *Geological Society - Special Publications*, 309(1), 5–38. <https://doi.org/10.1144/SP309.2>

- Fulljames, J. R., Zijerveld, L. J. J., & Franssen, R. C. M. W. (1997). Fault seal processes: Systematic analysis of fault seals over geological and production time scales. In P. Møller-Pedersen & A. Koestler (Eds.), *Hydrocarbon seals* (Vol. 7, pp. 51–59). Elsevier. [https://doi.org/10.1016/S0928-8937\(97\)80006-9](https://doi.org/10.1016/S0928-8937(97)80006-9)
- Geenens, G., Charpentier, A., & Paindaveine, D. (2017). Probit transformation for nonparametric kernel estimation of the copula density. *Bernoulli*, 23(3), 1848–1873. <https://doi.org/10.3150/15-BEJ798>
- Giles, M. B. (2015). Multilevel Monte Carlo methods. *Acta Numerica*, 24, 259–328. <https://doi.org/10.1017/S096249291500001X>
- Gross, H., & Mazuyer, A. (2021). GEOSX: A multiphysics, multilevel simulator designed for exascale computing. In *SPE reservoir simulation conference*. <https://doi.org/10.2118/203932-MS>
- Halland, E. K., Johansen, W. T., & Riis, F. (2011). CO₂ storage atlas Norwegian North Sea. *Norwegian Petroleum Directorate*. PO Box 600.
- Jeanne, P., Rutqvist, J., Wainwright, H. M., Foxall, W., Bachmann, C., Zhou, Q., et al. (2016). Effects of in situ stress measurement uncertainties on assessment of predicted seismic activity and risk associated with a hypothetical industrial-scale geologic CO₂ sequestration operation. *Journal of Rock Mechanics and Geotechnical Engineering*, 8(6), 873–885. <https://doi.org/10.1016/j.jrmge.2016.06.008>
- Jia, W., McPherson, B., Pan, F., Dai, Z., & Xiao, T. (2018). Uncertainty quantification of CO₂ storage using Bayesian model averaging and polynomial chaos expansion. *International Journal of Greenhouse Gas Control*, 71, 104–115. <https://doi.org/10.1016/j.ijggc.2018.02.015>
- Joe, H. (1996). Families of m -variate distributions with given margins and $m(m-1)$ bivariate dependence parameters. In *Distributions with fixed marginals and related topics*. L. Rüschendorf, B. Schweizer, & M. D. Taylor (Eds.), (Vol. 28, pp. 120–141). : Institute of Mathematical Statistics. <https://doi.org/10.1214/lnms/1215452614>
- Joe, H. (1997). *Multivariate models and multivariate dependence concepts*. CRC Press. <https://doi.org/10.1201/9780367803896>
- Jolley, S. J., Dijk, H., Lamens, J. H., Fisher, Q. J., Manzocchi, T., Eikmans, H., & Huang, Y. (2007). Faulting and fault sealing in production simulation models: Brent province, northern North Sea. *Petroleum Geoscience*, 13(4), 321–340. <https://doi.org/10.1144/1354-079306-733>
- Kang, M., Nordbotten, J. M., Doster, F., & Celia, M. A. (2014). Analytical solutions for two-phase subsurface flow to a leaky fault considering vertical flow effects and fault properties. *Water Resources Research*, 50(4), 3536–3552. <https://doi.org/10.1002/2013WR014628>
- Kinsman, R., Sauer, F. D., Jackson, H. A., & Wolynetz, M. S. (1995). Methane and carbon dioxide emissions from dairy cows in full lactation monitored over a six-month period. *Journal of Dairy Science*, 78(12), 2760–2766. [https://doi.org/10.3168/jds.S0022-0302\(95\)76907-7](https://doi.org/10.3168/jds.S0022-0302(95)76907-7)
- Kloek, T., & van Dijk, H. K. (1978). Bayesian estimates of equation system parameters: An application of integration by Monte Carlo. *Econometrica*, 46(1), 1–19. <https://doi.org/10.2307/1913641>
- Kolyukhin, D., & Tveranger, J. (2015). Statistical modelling of fault core and deformation band structure in fault damage zones. In *77th EAGE conference & exhibition 2015*. European Association of Geoscientists & Engineers. <https://doi.org/10.3997/2214-4609.201413043>
- Krevor, S., de Coninck, H., Gasda, S., Ghaleigh, N. S., de Gooyert, V., Hajibeygi, H., et al. (2023). Subsurface carbon dioxide and hydrogen storage for a sustainable energy future. *Nature Reviews Earth & Environment*, 4(2), 102–118. <https://doi.org/10.1038/s43017-022-00376-8>
- Krumscheid, S., & Pettersson, P. (2023). Sequential estimation using hierarchically stratified domains with Latin hypercube sampling. In A. Hinrichs, P. Kritzer, & F. Pillichshammer (Eds.), *Monte Carlo and Quasi-Monte Carlo Methods. MCQMC 2022. Springer Proceedings in Mathematics & Statistics*, (Vol. 460). Springer. https://doi.org/10.1007/978-3-031-59762-6_19
- Krumscheid, S., Yuan, F., & Pettersson, P. (2023). Package ‘ADSS’. <https://doi.org/10.5281/zenodo.7660513>
- L’Ecuyer, P., & Lemieux, C. (2002). Recent advances in randomized Quasi-Monte Carlo methods. In M. Dror, P. L’Ecuyer, & F. Szidarovszky (Eds.), *Modeling uncertainty: An examination of stochastic theory, methods, and applications* (pp. 419–474). Springer. https://doi.org/10.1007/0-306-48102-2_20
- Leverett, M. (1941). Capillary behavior in porous solids. *Transactions of the AIME*, 142(01), 152–169. <https://doi.org/10.2118/941152-g>
- Lie, K.-A. (2019). *An introduction to reservoir simulation using MATLAB/GNU Octave: User guide for the MATLAB reservoir simulation toolbox (MRST)*. Cambridge University Press. <https://doi.org/10.1017/9781108591416>
- Lindsay, N. G., Murphy, F. C., Walsh, J. J., & Watterson, J. (1992). Outcrop studies of shale smears on fault surface. In *The geological modelling of hydrocarbon reservoirs and outcrop analogues* (pp. 113–123). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781444303957.ch6>
- Manzocchi, T., Heath, A. E., Palanathakumar, B., Childs, C., & Walsh, J. J. (2008). Faults in conventional flow simulation models: A consideration of representational assumptions and geological uncertainties. *Petroleum Geoscience*, 14(1), 91–110. <https://doi.org/10.1144/1354-079306-775>
- McKay, M. D., Beckman, R. J., & Conover, W. J. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2), 239–245. <https://doi.org/10.1080/00401706.1979.10489755>
- Michie, E. A. H., & Braathen, A. (2024). How displacement analysis may aid fault risking strategies for CO₂ storage. *Basin Research*, 36(1), e12807. <https://doi.org/10.1111/bre.12807>
- Michie, E. A. H., Mulrooney, M. J., & Braathen, A. (2021). Fault interpretation uncertainties using seismic data, and the effects on fault seal analysis: A case study from the Horda Platform, with implications for CO₂ storage. *Solid Earth*, 12(6), 1259–1286. <https://doi.org/10.5194/se-12-1259-2021>
- Mulrooney, M. J., Osmond, J. L., Skurtveit, E., Faleide, J. I., & Braathen, A. (2020). Structural analysis of the smeaeia fault block, a potential CO₂ storage site, northern Horda Platform, North Sea. *Marine and Petroleum Geology*, 121, 104598. <https://doi.org/10.1016/j.marpetgeo.2020.104598>
- Nagler, T., Bumann, C., & Czado, C. (2019). Model selection in sparse high-dimensional vine copula models with an application to portfolio risk. *Journal of Multivariate Analysis*, 172, 180–192. <https://doi.org/10.1016/j.jmva.2019.03.004>
- Nagler, T., Schepsmeier, U., Stoeber, J., Brechmann, E. C., Graeler, B., Erhardt, T., et al. (2022). Package ‘vinecopulib’. <https://doi.org/10.5281/zenodo.4288292>
- Oladyshkin, S., Class, H., Helmig, R., & Nowak, W. (2011). A concept for data-driven uncertainty quantification and its application to carbon dioxide storage in geological formations. *Advances in Water Resources*, 34(11), 1508–1518. <https://doi.org/10.1016/j.advwatres.2011.08.005>
- Pei, Y., Paton, D. A., Knipe, R. J., & Wu, K. (2015). A review of fault sealing behaviour and its evaluation in siliciclastic rocks. *Earth-Science Reviews*, 150, 121–138. <https://doi.org/10.1016/j.earscirev.2015.07.011>
- Perrin, G., Soize, C., Duhamel, D., & Funfschilling, C. (2013). Karhunen–Loeve expansion revisited for vector-valued random fields: Scaling, errors and optimal basis. *Journal of Computational Physics*, 242, 607–622. <https://doi.org/10.1016/j.jcp.2013.02.036>
- Pettersson, P. (2016). Stochastic Galerkin formulations for CO₂ transport in aquifers: Numerical solutions with uncertain material properties. *Transport in Porous Media*, 114(2), 457–483. <https://doi.org/10.1007/s11242-015-0575-9>
- Pettersson, P., Keilegavlen, E., Sandve, T. H., Gasda, S. E., & Krumscheid, S. (2023). Scripts for copula modeling and uncertainty propagation in field-scale simulation of CO₂ Fault Leakage first release. *Zenodo*. <https://doi.org/10.5281/zenodo.10362253>
- Pettersson, P., & Krumscheid, S. (2022). Adaptive stratified sampling for nonsmooth problems. *International Journal for Uncertainty Quantification*, 12(6), 71–99. <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2022041034>

- Pettersson, P., Tveit, S., & Gasda, S. E. (2022). Dynamic estimates of extreme-case CO₂ storage capacity for basin-scale heterogeneous systems under geological uncertainty. *International Journal of Greenhouse Gas Control*, *116*, 103613. <https://doi.org/10.1016/j.ijggc.2022.103613>
- Pruess, K. (1991). TOUGH2 - A general-purpose numerical simulator for multiphase fluid and heat flow. Retrieved from <https://escholarship.org/uc/item/0wx8q119>
- Rabinovich, A., Li, B., & Durlinsky, L. J. (2016). Analytical approximations for effective relative permeability in the capillary limit. *Water Resources Research*, *52*(10), 7645–7667. <https://doi.org/10.1002/2016wr019234>
- Rahman, M. J., Choi, J. C., Fawad, M., & Mondol, N. H. (2021). Probabilistic analysis of Vette fault stability in potential CO₂ storage site Smeaheia, offshore Norway. *International Journal of Greenhouse Gas Control*, *108*, 103315. <https://doi.org/10.1016/j.ijggc.2021.103315>
- Rasmussen, A. F., Sandve, T. H., Bao, K., Lauser, A., Hove, J., Skaflestad, B., et al. (2021). The open porous media flow reservoir simulator. *Computers & Mathematics with Applications*, *81*, 159–185. <https://doi.org/10.1016/j.camwa.2020.05.014>
- Rinaldi, A. P., Rutqvist, J., & Cappa, F. (2014). Geomechanical effects on CO₂ leakage through fault zones during large-scale underground injection. *International Journal of Greenhouse Gas Control*, *20*, 117–131. <https://doi.org/10.1016/j.ijggc.2013.11.001>
- Ringrose, P. S., & Meckel, T. A. (2019). Maturing global CO₂ storage resources on offshore continental margins to achieve 2DS emissions reductions. *Scientific Reports*, *9*(1), 17944. <https://doi.org/10.1038/s41598-019-54363-z>
- Rosenblatt, M. (1952). Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, *23*(3), 470–472. <https://doi.org/10.1214/aoms/1177729394>
- Sandve, T. H., Gasda, S. E., Rasmussen, A., & Rustad, A. B. (2021). Convective dissolution in field scale CO₂ storage simulations using the OPM flow simulator. In *TCCS-11. CO₂ capture, transport and storage. Trondheim 22nd–23rd June 2021 short papers from the 11th international Trondheim CCS conference*.
- Sandve, T. H., Rustad, A. B., Thune, A., Nazarian, B., Gasda, S., & Rasmussen, A. F. (2022). Simulators for the gigaton storage challenge. A benchmark study on the regional Smeaheia model. In *EAGE geotech 2022 sixth EAGE workshop on CO₂ geological storage* (Vol. 2022, pp. 1–5). <https://doi.org/10.3997/2214-4609.20224033>
- Schueller, S., Braathen, A., Fossen, H., & Tveranger, J. (2013). Spatial distribution of deformation bands in damage zones of extensional faults in porous sandstones: Statistical analysis of field data. *Journal of Structural Geology*, *52*, 148–162. <https://doi.org/10.1016/j.jsg.2013.03.013>
- Shields, M. D. (2016). Refined Latinized stratified sampling: A robust sequential sample size extension methodology for high-dimensional Latin hypercube and stratified designs. *International Journal for Uncertainty Quantification*, *6*(1), 79–97. <https://doi.org/10.1615/int.j.uncertaintyquantification.2016011333>
- Shields, M. D., Teferra, K., Hapij, A., & Daddazio, R. P. (2015). Refined stratified sampling for efficient Monte Carlo based uncertainty quantification. *Reliability Engineering & System Safety*, *142*, 310–325. <https://doi.org/10.1016/j.ress.2015.05.023>
- Sklar, M. (1959). *Fonctions de répartition à n dimensions et leurs marges* (Vol. 8, pp. 229–231). Publications de l'Institut de statistique de l'Université.
- Sperrevik, S., Gillespie, P., Fisher, Q., Halvorsen, T., & Knipe, R. (2002). Empirical estimation of fault rock properties. In A. G. Koestler & R. Hunsdale (Eds.), *Hydrocarbon seal quantification* (Vol. 11, pp. 109–125). [https://doi.org/10.1016/S0928-8937\(02\)80010-8](https://doi.org/10.1016/S0928-8937(02)80010-8)
- Torabi, A., Ellingsen, T. S. S., Johannessen, M. U., Alaei, B., Rotevatn, A., & Chiarella, D. (2020). Fault zone architecture and its scaling laws: Where does the damage zone start and stop? *Geological Society - Special Publications*, *496*(1), 99–124. <https://doi.org/10.1144/SP496-2018-151>
- Wu, L., Thorsen, R., Ottesen, S., Meneguolo, R., Hartvedt, K., Ringrose, P., & Nazarian, B. (2021). Significance of fault seal in assessing CO₂ storage capacity and containment risks - An example from the Horda Platform, northern North Sea. *Petroleum Geoscience*, *27*(3), petgeo2020-102. <https://doi.org/10.1144/petgeo2020-102>
- Xiu, D., & Karniadakis, G. E. (2002). The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, *24*(2), 619–644. <https://doi.org/10.1137/S1064827501387826>
- Yielding, G., Freeman, B., & Needham, D. T. (1997). Quantitative fault seal prediction. *AAPG Bulletin*, *81*(6), 897–917. <https://doi.org/10.1306/522B498D-1727-11D7-8645000102C1865D>
- Zhang, L., & Singh, V. P. (2019). *Copulas and their applications in water resources engineering*. Cambridge University Press. <https://doi.org/10.1017/9781108565103>
- Zhang, Y., & Sahinidis, N. V. (2013). Uncertainty quantification in CO₂ sequestration using surrogate models from polynomial chaos expansion. *Industrial & Engineering Chemistry Research*, *52*(9), 3121–3132. <https://doi.org/10.1021/ie300856p>