Fluid Substitution Modelling for Seismic Detectability of Underground Hydrogen Storage

Joris Stekelenburg

MSc Applied Geophysics
Delft University of Technology / RWTH Aachen University / ETH Zürich

Carried out as a graduation internship at
Shell Global Solutions International B.V.

August 3, 2026

Abstract

Underground hydrogen storage in porous formations will require seismic monitoring, and the interpretation of that monitoring rests on fluid-substitution modelling that assumes pore pressure equilibrates within a wave period. Where hydrogen and brine coexist as mesoscopic patches, that assumption fails, and the rock response becomes frequency dependent. We investigate whether including that frequency dependence materially changes the detectability interpretation that a conventional low-frequency description would produce, under conditions representative of Dutch sandstone storage targets.

We define three reference cases spanning the Dutch parameter envelope, a constant-frame Control case and two pressure-sensitive-frame cases in the style of the Rotliegend and Bunter sandstones, and establish their low-frequency Gassmann reference response. We model the frequency-dependent response with the Johnson branching-function description of mesoscopic wave-induced fluid flow, validate the implementation against exact reference solutions and the White concentric-sphere benchmark, and map how permeability and patch radius place the relaxation peak relative to the seismic band. We then embed each case at its true reservoir depth in a laterally invariant layer-cake column, generate synthetic prestack records, and migrate them to depth, comparing an all-brine state, a relaxed low-frequency state, and a dispersive state under uniform, gravity-segregated, and randomised per-layer saturation topologies, beneath both a single-block and a realistic layered Dutch overburden, and under band-limited noise at prescribed signal-to-noise ratios.

Our results show that the answer to the question we pose is conditional and case-dependent. Where the relaxation peak falls inside the seismic band, which among our cases means the larger-patch Bunter-style geometry, the dispersive correction survives imaging and noise at field-attainable signal-to-noise ratios, reaches ten percent of the reference reflector amplitude, and opposes the bulk fluid-substitution response, so that a low-frequency description overpredicts the imaged monitoring signal by a factor of between 1.3 and 2.0 across 15 to 50 Hz. Elsewhere the low-frequency description is accurate to within a few percent, and the full monitoring signal remains recoverable under noise in every case, so hydrogen injection itself is seismically visible throughout. Our results describe the fluid effect alone, under a dispersion-only, multiple-free, laterally invariant treatment, and are an upper bound on what a field survey could recover; within that scope, we identify where in Dutch-reservoir parameter space frequency-dependent rock physics is worth carrying, and where a conventional low-frequency workflow is sufficient.

Acknowledgments

This work was sponsored by Shell Global Solutions International B.V.

I'd like to send my thanks to my supervisors Deyan Dragonov, Christopher Willacy, Ali Reza Bagheri and Eric Verschuur for their guidance and support throughout the project. My gratitude also goes to Karin de Borst for her mentorship and support during my internship. I would also like to thank the members of the Shell internship team for their companionship and involvement during my internship experience at the Shell Energy Transition Campus Amsterdam, with an extended thanks to Evren Unsal for her management and involvement of this vibrant community.

Contents

List of Figures

Chapter 1

Introduction and Motivation

1.1 The energy transition and the need for large-scale storage

The transition from fossil fuels to renewable energy requires energy storage at scales that short-duration technologies cannot provide. Wind and solar generation vary on diurnal, weather-driven, and seasonal timescales, while demand remains high during periods of low renewable output. In Northwestern Europe this mismatch is particularly strong because winter demand coincides with reduced solar production and variable wind conditions. In the Netherlands, offshore wind capacity is planned to reach 21 GW by 2032 and 70 GW by 2050 (Reijnen-Mooij et al., 2025). Seasonal balancing at that scale requires storage capacities far beyond batteries or surface tanks.

Hydrogen is a leading candidate for long-duration energy storage because surplus electricity can be converted to H_2, stored, and later reused for power, heat, or industrial feedstock (Heinemann et al., 2021). For storage requirements on the order of several terawatt-hours, the relevant option is geological storage rather than surface infrastructure. Recent estimates place the Dutch need for large-scale hydrogen storage at 7–20 TWh by 2050 (Reijnen-Mooij et al., 2025). That scale makes underground hydrogen storage (UHS) in porous formations a serious monitoring problem as well as an energy-system opportunity. The monitoring requirement is concrete: hydrogen injection into a porous reservoir must be tracked for plume migration, leakage out of the storage formation, and pressure changes that could compromise the seal. Time-lapse seismic methods are the established geophysical tool for this class of monitoring problem in CO_2 storage, where they image plume evolution at injection sites such as Sleipner and Ketzin (Kazemeini et al., 2010; Eid et al., 2015), and the same logic carries over to hydrogen.

1.2 UHS in the Dutch subsurface

The Dutch subsurface is well suited to a thesis on porous-media UHS because it contains a large inventory of depleted gas fields with proven trapping structures, existing data, and operational history (Reijnen-Mooij et al., 2025). A recent national screening indicates that roughly 140 depleted fields show no identified showstoppers for UHS and that many more remain candidates subject to further assessment (Reijnen-Mooij et al., 2025). In western Netherlands, where industrial hydrogen demand is expected to develop first, depleted gas fields are particularly relevant because salt cavern options are limited.

Large-scale porous-media UHS remains only lightly tested in practice (Heinemann et al., 2021; Zivar et al., 2021). That raises a straightforward monitoring requirement: if H_2 is injected into a depleted reservoir, can seismic methods detect the resulting fluid substitution, and under what assumptions do the predicted signals remain trustworthy?

Two reservoir groups are especially relevant. Rotliegend sandstones provide many of the onshore candidates and commonly fall in the porosity and permeability range that is favourable for seismic monitoring studies. Bunter sandstones are also important in the wider Dutch and southern North Sea setting and bring stronger layering and depositional heterogeneity into view (Loyola et al., 2026). These formations define the practical context for our research but we limit our work to the rock-physics and 1D monitoring scale rather than to a site-specific field model. Figure 1.1 shows where these candidate fields and reservoir groups are located across the Netherlands.

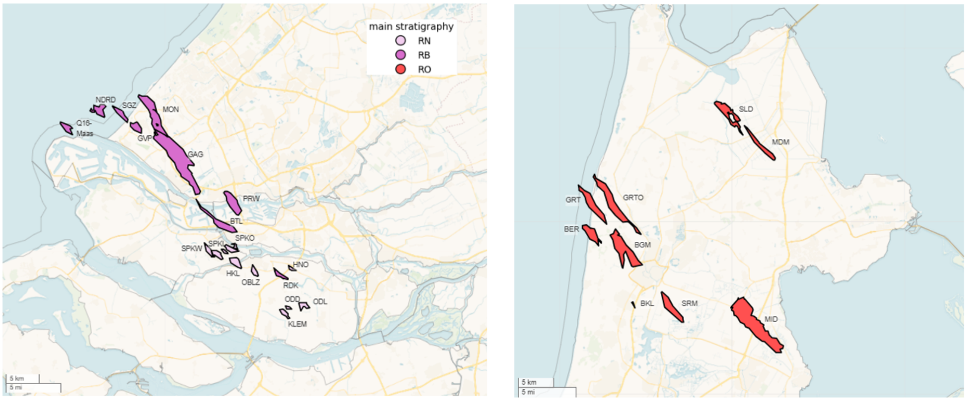

Figure 1.1: Locations and main reservoir units of candidate underground hydrogen storage fields in the Netherlands. The Rotliegend and Bunter sandstone candidates are concentrated in Noord-Holland and Zuid-Holland. Adapted from Reijnen-Mooij et al. (2025).

1.3 Seismic monitoring as a detectability problem

Time-lapse seismic monitoring is the established geophysical tool for tracking subsurface fluid changes. In CO_2 storage it has been used to image plume evolution, infer saturation changes, and quantify travel-time and reflectivity shifts (Kazemeini et al., 2010; Eid et al., 2015; MacBeth et al., 2024). The same monitoring logic is relevant to UHS: replacing brine with a lighter and more compressible pore fluid changes acoustic impedance, reflectivity, and travel time.

For hydrogen, the fluid contrast is particularly strong. Under reservoir conditions, H_2 remains much less dense and much more compressible than brine (Batzle & Wang, 1992; Gei et al., 2024). That suggests favourable detectability, but it does not by itself guarantee that low-frequency rock-physics predictions are sufficient. Hydrogen also has very low viscosity, and porous-media injection is likely to create sub-seismic saturation heterogeneity. Those ingredients make frequency-dependent patchy-saturation effects a plausible part of the monitoring response rather than a secondary detail.

Bagheri et al. (2025) already showed, within a low-frequency Gassmann workflow, that hydrogen substitution in sandstone can produce useful seismic signatures. Bagheri et al. (2026) extended this to a heterogeneous depleted gas reservoir with a methane cushion gas under cyclic injection and production. The open question in our current study is narrower and more specific: does adding frequency-dependent patchy-saturation physics materially change the interpretation that would be drawn from the established low-frequency baseline?

1.4 From the low-frequency baseline to the current step

The low-frequency reference model has already been established in earlier work (Stekelenburg, 2026). That stage used Batzle–Wang brine properties (Batzle & Wang, 1992), Peng–Robinson gas properties (Peng & Robinson, 1976), Gassmann substitution (Gassmann, 1951), and static partial-saturation mixing (Wood, 1955; Brie et al., 1995) to show that brine-to-hydrogen substitution should produce coherent changes in acoustic impedance, reflectivity, and travel time.

Here, we do not re-open the full baseline study as a co-equal contribution. Instead, we take the next modelling step. Throughout this thesis, “baseline” refers to the low-frequency Gassmann reference model against which the frequency-dependent results are compared; where the pre-injection state of a time-lapse survey is meant, we say “pre-injection” explicitly. The static model assumes complete pressure equilibration within a wave period, an assumption made here for the seismic monitoring band of roughly 10 to 100 Hz. When brine-rich and gas-rich regions coexist at mesoscopic scale, that assumption can fail within this band: incomplete pressure diffusion between patches then produces velocity dispersion and attenuation through wave-induced fluid flow (White, 1975; Johnson, 2001; Müller et al., 2010), the mechanism developed in Chapter 2. Our practical question is therefore whether this frequency-dependent mechanism is large enough to change the seismic detectability interpretation inherited from the low-frequency baseline.

This thesis is organised as a modelling ladder. The low-frequency Gassmann result is retained as the baseline reference. The active extension is the frequency-dependent patchy-saturation stage. That effect is first isolated in a small idealised comparison domain and is then inserted back into the broader 1D monitoring workflow. More realistic reservoir architecture, saturation topology, and depleted-field initial conditions remain later steps rather than completed parts of the present deliverable (Figure 1.2).

1.5 Research aims

We aim to determine whether frequency-dependent patchy saturation materially changes the seismic-monitoring interpretation for UHS, relative to the low-frequency reference model already established. The emphasis is not on multiplying models. It is on establishing a clear hierarchy among them, and on testing whether the added physics changes a conclusion that a simpler description would have reached.

Three commitments shape that hierarchy, and stating them here keeps the questions below compact.

The low-frequency reference model of Stekelenburg (2026) remains the comparison point throughout. Every frequency-dependent result is reported as a change relative

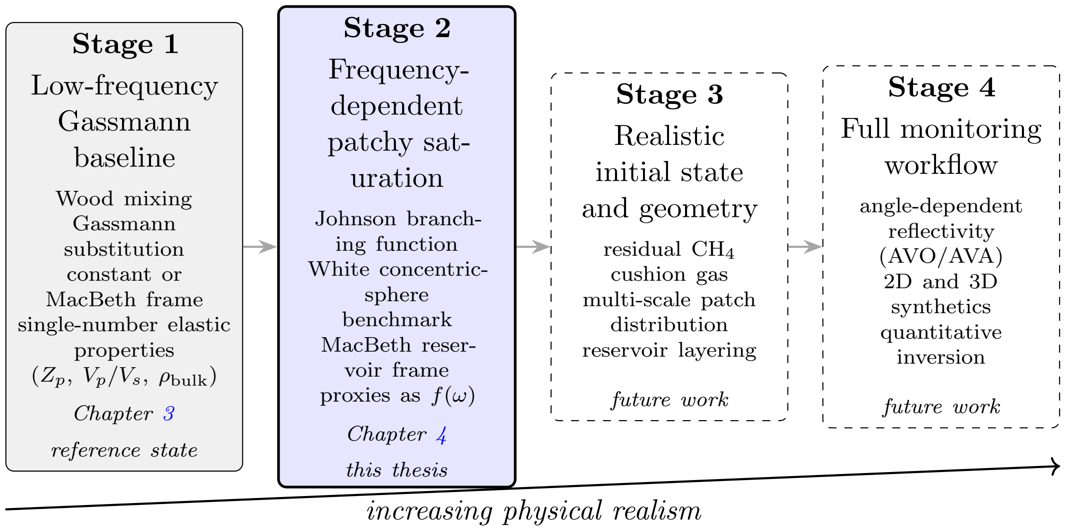

Figure 1.2: The four modelling stages and their increasing physical realism are shown for seismic monitoring of UHS in Dutch sandstone reservoirs. Stage 1 is the low-frequency Gassmann reference and Stage 2 is the frequency-dependent patchy-saturation model. Stages 3 and 4 show subsequent realism steps not implemented here.

to it, so that the question is always whether the additional physics matters rather than what it predicts in isolation.

The Johnson branching-function model (Johnson, 2001) is the frequency-dependent model we carry through the work, because its two analytically derived parameters make it evaluable across the parameter space without a geometry-specific boundary-value solution for each case. That scalability is what makes the later scenario modelling possible. The White concentric-sphere model (White, 1975) and the exact reference solutions of Johnson (2001, Appendix B) serve only as benchmarks, within the concentric-sphere gas-inner geometry where all three descriptions overlap; the hierarchy is set out in Chapter 2.

Detectability, finally, is assessed relative to the fluid effect alone. We ask whether a fluid-substitution change survives the modelling chain as a resolvable change in the seismic observables, not whether it would be recoverable from a real time-lapse survey, which depends in addition on acquisition geometry, noise, and processing repeatability. The three time-lapse proxies used throughout are the fractional acoustic-impedance change $\Delta Z/Z$, the normal-incidence reflectivity change ΔR, and the two-way travel-time shift Δt_{2way}.

Against those commitments, we address four linked questions.

1. **Frequency-dependent effect.** How does patchy saturation modify P-wave velocity and attenuation relative to the low-frequency reference model, and under what combinations of permeability and patch size does the relaxation move into the seismic band?

2. **Model hierarchy in the shared benchmark domain.** How closely does the Johnson branching-function model follow the exact reference solution and the White concentric-sphere benchmark in the shared concentric-sphere gas-inner domain?

3. **Detectability relevance.** When the frequency-dependent bulk modulus is carried into the monitoring workflow, do the three time-lapse proxies change

enough to alter the detectability interpretation, or is the frequency-dependent effect diluted along the way?

4. **Next realism step.** Once the frequency-dependent stage is constrained, which extension is the most defensible next step: more realistic reservoir architecture, more realistic saturation topology, or a depleted-field initial state with residual CH_4?

1.6 Thesis outline

Chapter 2 develops the theoretical framework: the compact low-frequency ingredients, the Gassmann–Wood and Gassmann–Hill limits, and the frequency-dependent patchy-saturation hierarchy. Chapter 3 states the Dutch reservoir context, defines the three reference cases of Table 3.1, and carries the low-frequency reference model forward in compact form. Chapter 4 develops the frequency-dependent modelling: benchmark validation, an idealised parameter scan, the sphere-versus-slab geometry sensitivity, and the embedding into the 1D monitoring workflow. Chapter 5 carries the resulting response into migrated seismic amplitudes through laterally invariant layer-cake scenarios at true reservoir depth, under both a single-block and a layered Dutch overburden. Chapter 6 interprets the outcome against the literature and states the limitations and assumptions, and Chapter 7 draws the conclusions and the future work. The appendices collect supplementary equations, frame-model details, and compact benchmark notes.

Chapter 2

Background and Theory

In this chapter, we introduce the theoretical ingredients that remain active in the current workflow. The low-frequency background stays compact because it already underpins the baseline reference in Chapter 3. We then present the frequency-dependent extension with a strict hierarchy: the Johnson branching-function model is the main model in our work, the White concentric-sphere model is the external benchmark, and we retain exact reference solutions for simple geometries only as benchmark reference. We defer detailed support formulas to the appendices.

2.1 Fluid properties under reservoir conditions

The workflow requires density, bulk modulus, and viscosity for each pore fluid at reservoir pressure and temperature. These quantities enter both the low-frequency reference model and the frequency-dependent patchy-saturation models.

2.1.1 Brine

We compute brine density ρ_b and P-wave velocity V_b from the Batzle–Wang correlations (Batzle & Wang, 1992), and define the brine bulk modulus as

$$K_b = \rho_b V_b^2. \tag{2.1}$$

Brine viscosity η_b does not appear in Gassmann substitution, which assumes that pore pressure has already equilibrated. We retain it because it controls the rate at which that equilibration occurs, and therefore the frequency at which the relaxed description ceases to apply. Section 2.2 makes that dependence explicit.

2.1.2 Gas

We compute gas density and bulk modulus from the Peng–Robinson equation of state (Peng & Robinson, 1976). Solving the Peng–Robinson equation of state for the compressibility factor Z_c gives the molar volume V_m and the gas density ρ_g as

$$V_m = \frac{Z_c RT}{P}, \qquad \rho_g = \frac{M}{V_m}, \tag{2.2}$$

where M is the molar mass, R the universal gas constant, and T and P the reservoir temperature and pressure. We write the thermodynamic compressibility factor as Z_c to distinguish it from the acoustic impedance Z, which is used throughout the rest of this thesis and is unrelated.

The bulk modulus required for wave propagation is the adiabatic modulus K_S, not the isothermal modulus K_T. A seismic wave compresses the pore fluid faster than heat can diffuse out of the compressed region, so the compression is adiabatic rather than isothermal, and the adiabatic modulus is the stiffer of the two (Batzle & Wang, 1992). We obtain K_T from the equation of state and convert using the heat-capacity ratio $\gamma = c_p/c_v$:

$$K_T = -V_m \left(\frac{\partial P}{\partial V_m}\right)_T, \qquad K_S = \gamma K_T. \tag{2.3}$$

The distinction matters for hydrogen, where γ departs appreciably from unity at reservoir conditions, so that using K_T in place of K_S would underestimate the gas stiffness and overstate the modelled fluid contrast.

Our workflow supports H_2, CH_4, and CO_2. Hydrogen is the target of this thesis. Methane is included because a practical Dutch storage site would be a depleted gas field containing residual methane, so the same equation of state must supply methane properties whenever that initial state is modelled. Carbon dioxide is included for comparison with the published CO_2 storage literature and is not a results branch of this work.

2.2 Scale, diffusion length, and the limits of a mixing law

Gassmann's equation, introduced in Section 2.3, is exact within its own assumptions, and those assumptions are not in question here. What limits the low-frequency description of a partially saturated rock is not Gassmann's equation but the step that precedes it: the replacement of two coexisting fluids by a single effective fluid modulus. That replacement presumes that brine and gas share one pore pressure at every instant. Whether they do is a question about length scales.

The diffusion length. When a compressional wave passes through a rock containing both brine-saturated and gas-saturated regions, it compresses the two unequally, because the gas is far more compressible. The resulting pore-pressure difference drives fluid flow from the stiffer regions toward the softer ones. That flow is diffusive, and over one wave period at angular frequency ω it penetrates a characteristic distance

$$L_d \sim \sqrt{\frac{D}{\omega}}, \qquad D = \frac{\kappa}{\eta\phi^2}\,\frac{PR - Q^2}{P + 2Q + R}, \tag{2.4}$$

where κ is the permeability, η the fluid viscosity, ϕ the porosity, and D the Biot slow-wave diffusivity of Johnson (2001, Eq. 8), in which P, Q and R are the Biot coefficients of Johnson (2001, Eq. 5). This is the diffusivity our implementation uses throughout; the coefficients are written out in Appendix C. The diffusion length L_d is the distance over which pore pressure can equalise within a wave period. It falls as the square root of frequency, so the higher the frequency, the smaller the region that can equilibrate.

The patchiness criterion. Comparing L_d with the characteristic size a of the gas-saturated regions divides the behaviour into two regimes (Mavko & Mukerji, 1998; Müller et al., 2010).

When $a \ll L_d$, pressure equilibrates fully across the patches within a wave period. The two fluids behave as a single fluid at a common pressure, the effective modulus is the uniform-mixing average of Equation (2.11), and the rock response is the relaxed Gassmann–Wood limit. This is the regime in which a static mixing law is not merely convenient but correct.

When $a \gg L_d$, no appreciable flow occurs between patches during a wave period. Each patch behaves as an independently saturated region, the responses are averaged at the level of the saturated rock rather than the fluid, and the result is the stiff, unrelaxed Gassmann–Hill limit.

Between these regimes, when a and L_d are comparable, flow occurs but does not complete. Pore pressure partially equalises, out of phase with the passing wave. The rock is then stiffer than the relaxed limit but softer than the unrelaxed one, the velocity depends on frequency, and energy is lost to the incomplete flow. This is mesoscopic wave-induced fluid flow, and it is the mechanism this thesis models.

Why this is not a property of the reservoir. Two consequences follow, and both matter for how our results should be read.

First, patchiness is not a fixed attribute of a rock. The same reservoir, with the same fluid distribution, is uniformly mixed at low enough frequency and patchy at high enough frequency, because the comparison in Equation (2.4) involves ω. Describing a reservoir as patchy is therefore incomplete without a frequency, and any statement about patchy behaviour in this thesis is a statement about a rock, a fluid, a patch size, and a frequency together. The comparison is equally local in space: the diffusion length varies with the local permeability and viscosity, and the patch size with the local saturation distribution, so one region of a plume can lie in the frequency-dependent regime while uniform mixing holds elsewhere in the same reservoir at the same frequency.

Second, because L_d depends on permeability and viscosity while a is a property of the saturation distribution, the frequency at which the transition occurs is set by reservoir properties that vary between formations. This is why the same seismic band can sample the relaxed limit in one reservoir and the transition itself in another, and it is the origin of the case-dependent behaviour reported in Chapters 4 and 5.

Hydrogen as a limiting case. Hydrogen is an extreme fluid on both axes of this comparison. It is far more compressible than brine, which makes the pressure contrast driving the flow unusually large, and it is far less viscous, which makes the flow itself unusually rapid (Gei et al., 2024). The viscosity contrast is the larger of the two, and it can be quantified at the conditions we actually model. The pore-pressure and geothermal-temperature conditions of the three reference cases in Table 3.1 are Control at 9.91 MPa and 40 °C, Bunter-style at 22 MPa and 62.5 °C, and Rotliegend-style at 30 MPa and 94 °C. Evaluated with the correlations implemented in this workflow at these conditions, hydrogen viscosity is 9.2, 9.9 and 10.7 μPa s respectively, against 963, 707 and 529 μPa s for the brine of the corresponding case. Hydrogen is therefore between fifty and one hundred times less viscous than the brine it displaces, and it is the least viscous of the three gases the workflow supports: at the same three conditions methane gives 10.1, 15.7 and 17.5 μPa s and carbon dioxide 16.2, 18.1 and 20.0 μPa s, so the ordering $\eta_{\mathrm{H_2}} < \eta_{\mathrm{CH_4}} < \eta_{\mathrm{CO_2}} \ll \eta_b$ holds at every case condition (Wang et al., 2022). The carbon dioxide values follow a dilute-gas correlation and understate the viscosity of the dense supercritical phase that carbon dioxide forms at the shallower two conditions; they are quoted to fix the ordering, not as an accurate description of a carbon dioxide storage case, which this thesis does not model. Both properties push the response toward the frequency-dependent regime rather than away from it, which is why a mixing law alone is a weaker approximation for hydrogen storage than it would be for a denser, more viscous pore gas.

2.3 Gassmann fluid substitution

Gassmann's equation gives the saturated bulk modulus of a porous rock in the relaxed limit, in which pore pressure is fully equilibrated throughout the representative volume (Gassmann, 1951; Smith et al., 2003):

$$K_{\mathrm{sat}} = K_{\mathrm{dry}} + \frac{\left(1 - \dfrac{K_{\mathrm{dry}}}{K_0}\right)^2}{\dfrac{\phi}{K_f} + \dfrac{1-\phi}{K_0} - \dfrac{K_{\mathrm{dry}}}{K_0^2}}, \tag{2.5}$$

where K_0 is the mineral bulk modulus, K_{dry} the dry-frame bulk modulus, ϕ the porosity, and K_f the effective fluid bulk modulus. The shear modulus is unchanged by the pore fluid:

$$G_{\mathrm{sat}} = G_{\mathrm{dry}}. \tag{2.6}$$

Bulk density and seismic velocities follow from

$$\rho_{\mathrm{bulk}} = (1-\phi)\rho_m + \phi\rho_f, \tag{2.7}$$

$$V_p = \sqrt{\frac{K_{\mathrm{sat}} + \frac{4}{3}G_{\mathrm{sat}}}{\rho_{\mathrm{bulk}}}}, \tag{2.8}$$

$$V_s = \sqrt{\frac{G_{\mathrm{sat}}}{\rho_{\mathrm{bulk}}}}, \tag{2.9}$$

where ρ_m is the mineral density and ρ_f the density of the pore fluid, taken as the saturation-weighted average of the brine and gas densities. Fluid densities mix linearly regardless of how the fluids are distributed, so ρ_{bulk} is unaffected by the choice of mixing law, and every difference between the closures compared in this thesis enters through K_{sat} alone.

Equation (2.5) is a relaxed-limit result. It is the correct description whenever the criterion of Section 2.2 places the rock in the equilibrated regime, and the frequency-dependent formulations introduced in Section 2.6 must reduce to it as $\omega \to 0$.

The Biot coefficient and modulus,

$$\alpha = 1 - \frac{K_{\mathrm{dry}}}{K_0}, \qquad \frac{1}{M} = \frac{\phi}{K_f} + \frac{\alpha - \phi}{K_0}, \tag{2.10}$$

allow Equation (2.5) to be written compactly as $K_{\mathrm{sat}} = K_{\mathrm{dry}} + \alpha^2 M$. We use this form because the same quantities recur in the frequency-dependent formulation and in the supporting expressions collected in Appendix C.

2.4 Partial saturation and fluid mixing laws

When the pore space contains both brine and gas, the relaxed workflow replaces the two-fluid system by a single effective fluid modulus before applying Gassmann substitution. How that replacement is made is not a detail. It encodes an assumption about how the two fluids are distributed in the pore space, and that assumption, rather than Gassmann's equation itself, is what limits the low-frequency description.

2.4.1 The two limiting distributions

Two idealised distributions bracket the possibilities. If the two fluids are mixed finely enough that pore pressure equilibrates between them within a wave period, the fluid phases share a single pressure and the effective modulus is the harmonic, or Reuss, average of the two fluid moduli (Wood, 1955):

$$\frac{1}{K_f^{\mathrm{W}}} = \frac{S_b}{K_b} + \frac{S_g}{K_g}. \tag{2.11}$$

This is Wood's law. It is the soft, lower limit: because it is a harmonic average, a small volume fraction of a very compressible phase dominates the result. For hydrogen and brine, where $K_g \ll K_b$, Wood mixing therefore produces a large modulus reduction at low gas saturation and a shallow response thereafter.

The opposite idealisation is that the two fluids do not communicate at all within a wave period, so that each phase sustains its own pore pressure. The corresponding stiff limit is the arithmetic, or Voigt, average,

$$K_f^{\mathrm{V}} = S_b K_b + S_g K_g, \tag{2.12}$$

which is dominated by the stiffer phase and gives a nearly linear response in saturation.

The two averages are bounds in the sense that any physically admissible distribution of the same two fluids, at the same saturation, produces an effective modulus between them (Mavko & Mukerji, 1998). Which bound is approached depends on whether pressure equilibration is complete within a wave period, and that is a question about length scales and frequency rather than about saturation. Section 2.6 makes that dependence explicit; here it is enough to note that Equations (2.11) and (2.12) are frequency-independent idealisations of the two extremes, and that neither describes the intermediate case.

Equations (2.11) and (2.12) bound the fluid modulus only; the distinct rock-level limits, which average the separately saturated rock responses rather than the fluid moduli, are introduced in Section 2.6.2.

2.4.2 Brie mixing as an interpolation

Real partial saturation is rarely at either extreme, and the Brie relation provides a one-parameter interpolation between them (Brie et al., 1995):

$$K_f^{\mathrm{B}} = (K_b - K_g)\, S_b^e + K_g, \tag{2.13}$$

with e a stiffening exponent. At $e = 1$ Equation (2.13) reduces exactly to the Voigt average of Equation (2.12). As e increases the curve softens toward Wood, which it approaches only asymptotically: for a brine-to-gas modulus contrast as large as that between brine and hydrogen, an exponent of order forty is required before the Brie curve resembles Wood (Papageorgiou et al., 2016). Between those extremes, increasing e moves the predicted modulus from a nearly linear saturation dependence toward one dominated by the compressible phase at low saturation.

The exponent is conventionally described as empirical, and for the original formulation that is accurate: it was obtained by fitting acoustic logs from unconsolidated shaly sands, which yielded values of roughly 2 to 5, and it carries no independent physical determination (Brie et al., 1995). That characterisation is, however, too strong for the later derivation of Papageorgiou et al. (2016), who obtain a Brie-like mixing law from a capillary-pressure argument and express the equivalent exponent in terms of a parameter bounded by the arithmetic and harmonic limits. Their law is better described as phenomenological than empirical: the interpolation parameter acquires a physical interpretation in terms of capillary pressure, even though its saturation dependence still requires a pore-scale model and is left as a fitted quantity in practice.

Choice of exponent. We adopt $e = 5$ throughout, and treat it as a stated modelling assumption rather than a measured property of Dutch sandstone. Two considerations support the choice. It lies at the upper edge of the range originally fitted by Brie et al. (1995). And it has peer-reviewed precedent for gas and brine in North Sea sandstone, having been adopted at Sleipner as an effective-fluid description intermediate between homogeneous and patchy responses (Queißer & Singh, 2013; Dupuy et al., 2017). That precedent, rather than any hydrogen-specific calibration, is what the choice rests on: we are not aware of a published Brie exponent fitted to hydrogen and brine, and we do not claim one.

Two weaknesses of that choice should be stated plainly. No fitted Brie exponent has been published for Rotliegend or Bunter Sandstone, or for the Dutch sector, so any value is transferred rather than calibrated. And the transfer crosses more than lithology: the Sleipner precedent is a soft, unconsolidated, high-porosity sand at roughly a kilometre depth, whereas our cases run to 2800 m in consolidated sandstone, so a single exponent is being carried across a range of depth and consolidation state over which there is no reason to expect it to be constant. We therefore report the sensitivity of the comparison to the exponent rather than relying on a single value alone.

Admissibility over the modelled range. Because Equation (2.13) interpolates at the fluid level while the relevant limits act at the rock level, we verify directly that Gassmann substitution with a Brie-mixed fluid remains bracketed by the Gassmann–Wood and Gassmann–Hill limits over the saturation range we model. It does, for all three reference cases and for every exponent between 3 and 6, with the tightest margin approaching the lower limit at the highest saturation and exponent considered. Outside our modelled range, at hydrogen saturations above roughly 0.6, the Brie prediction falls marginally below the Gassmann–Wood limit, by at most about one part in a thousand. This is a known consequence of applying a fixed exponent to a fluid pair with a very large modulus contrast, in which the S_b^e term collapses toward the gas endpoint faster than the harmonic average does, and it has been documented as a failure of large-exponent Brie curves to remain consistent with the Reuss lower bound (Rozhko, 2020). It does not affect any result reported here.

2.4.3 Role in this thesis

Wood and Brie enter our workflow in two distinct roles, and separating them matters for what follows.

Wood mixing is structural. It defines the relaxed limit of the frequency-dependent models, and the low-frequency reference model of Chapter 3 is built on it.

Brie mixing is comparative. It is not required by any of the frequency-dependent formulations and is not the closure used in Chapter 4. We carry it because it approximates the effect of patchiness with a single fitted parameter, without modelling the underlying pressure diffusion; comparing the uniform-mixing limit, the Brie interpolation, and the frequency-dependent formulation separates how much of the response is due to saturation distribution alone from how much requires frequency dependence.

2.5 Dry rock-frame models

The dry-frame moduli control both the low-frequency response and the width of the modulus bracket within which the frequency-dependent response must lie. Our workflow retains the three frame options established in Stekelenburg (2026): (i) a constant frame, (ii) a MacBeth-style pressure-sensitive frame (MacBeth, 2004), and (iii) a Gei-style Krief–Athy compaction-linked frame (Gei et al., 2024).

Option (ii) is the default reservoir frame for the Rotliegend-style and Bunter-style cases reported in Chapter 4. Option (i) is retained only for the Control case, fixed at $K_{\mathrm{dry}} = 12$ GPa and $G_{\mathrm{dry}} = 10$ GPa throughout this thesis, so that one case carries a frame that is deliberately independent of depth and pressure.

Two considerations motivate the MacBeth choice. First, Table 2 of MacBeth (2004) supplies sandstone-specific parameterisations that map onto the Dutch candidate formations: a Rotliegend Facies B re-fit row for the Rotliegend-style case, and a Berea Sandstone row as the closest available clean continental fluvial analogue for the Bunter-style case. Second, the MacBeth form retains an explicit differential-pressure dependence in both K_{dry} and G_{dry}, which is appropriate over the abandonment-pressure range characteristic of storage in a depleted field.

Option (iii) remains available as a comparison but is not used for the results reported here, because it is calibrated to soft, high-porosity Utsira-style sand (Gei et al., 2024) and does not map onto the consolidated Dutch sandstone targets covered by the parameterisations of MacBeth (2004). This difference in frame stiffness is one of the two reasons our response differs quantitatively from that study, and we return to it in Chapter 6.

We do not repeat the full parameterisations here; they are collected in Appendix B.3, and the case-specific values adopted in this thesis are listed in Table 3.1. The point to carry forward is that the frame choice remains active throughout: the frequency-dependent stage does not replace the dry-frame description, it reuses it, and the bracket between the relaxed and unrelaxed limits derived in Section 2.6.2 is set by the frame as much as by the fluids.

2.6 Frequency-dependent patchy saturation

Section 2.2 established that a static mixing law describes a partially saturated rock correctly only when the patch size is small compared with the diffusion length, and that between the two limiting regimes the response depends on frequency. This section sets out the models we use to describe that intermediate regime. A passing compressional wave generates pressure gradients between gas-rich and brine-rich regions, drives local fluid flow, and produces velocity dispersion and attenuation

(White, 1975; Johnson, 2001; Pride et al., 2004; Müller et al., 2010). This mesoscopic wave-induced fluid flow is the rock-physics mechanism at the centre of this thesis.

We use three geometrical idealisations of the patch distribution, shown in Figure 2.1: a gas-rich sphere embedded in brine, the same geometry with the two fluids exchanged, and alternating parallel slabs. These are not competing descriptions of a real reservoir, which contains no such regular structures. They are tractable geometries for which the diffusion problem can be solved, and comparing them establishes how sensitive the response is to the assumed shape of the patches rather than to their size alone.

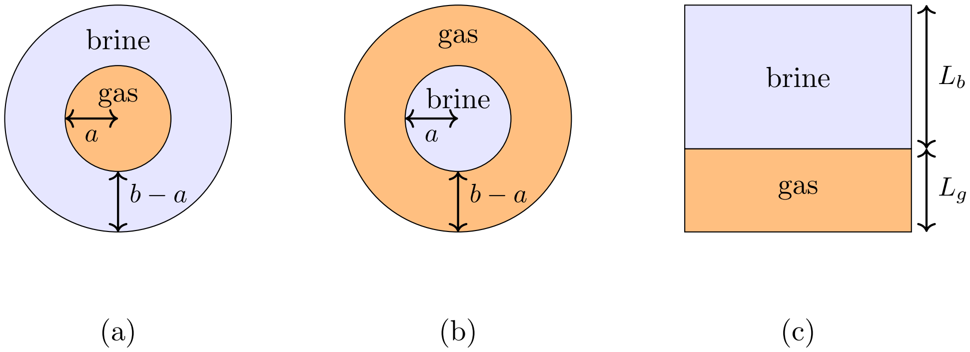

Figure 2.1: Gas- and brine-rich patch geometries and their dimensions are shown for the three Johnson-model benchmarks. (a) A gas-rich inner sphere of radius a within a brine-rich sphere of radius b. (b) A brine-rich inner sphere within a gas-rich outer shell. (c) Alternating gas-rich and brine-rich slabs of thicknesses L_g and L_b. The White model is restricted to geometry (a).

2.6.1 Methods we use

Our workflow uses several fluid-substitution and partial-saturation methods at different points in the modelling chain, and their roles are distinct rather than competing. Table 2.1 summarises them.

Relaxed Gassmann substitution with the Wood average defines the low-frequency reference model of Chapter 3, and the Hill average defines its high-frequency counterpart. Together they bracket the frequency-dependent response, as set out in Section 2.6.2. The Brie relation is carried as a comparison closure, positioned between those two limits, and is used in Chapter 3 rather than in the frequency-dependent modelling.

The frequency-dependent core is the Johnson branching-function model. The White concentric-sphere model serves as an external benchmark, and the exact solutions of Johnson (2001, Appendix B) as internal benchmark support. The distinction matters for how the results should be read: the branching-function model is the only path used to generate the results reported in Chapters 4 and 5, and the other two exist to establish that it is trustworthy within the geometry where all three apply.

Table 2.1: Fluid-substitution and partial-saturation methods used in this thesis. Gassmann substitution with the Wood average is the relaxed limit and the basis of the low-frequency reference model of Chapter 3; the Hill average is the unrelaxed limit; the Johnson branching-function model spans the two as a function of frequency. Brie mixing is a comparison closure between the two limits, used in the low-frequency comparison of Chapter 3. The White concentric-sphere model and the exact reference solutions are benchmarks for the branching-function model within their respective shared domains.

Method	Regime	Geometry	Citation
Gassmann substitution	low-frequency limit	single fluid	(Gassmann, 1951; Smith et al., 2003)
Gassmann + Wood's law	low-frequency limit	uniform mixture	(White, 1975)
Gassmann + Brie	low-frequency, empirical	patchy mixing	(Brie et al., 1995; Papageorgiou et al., 2016)
Gassmann–Hill	high-frequency limit	separated patches	(Johnson, 2001)
Johnson branching function	frequency-dependent	geometry-parametric	(Johnson, 2001)
White concentric-sphere	frequency-dependent	concentric sphere	(White, 1975; Carcione et al., 2003)
Exact reference solutions	frequency-dependent	slab and sphere	(Johnson, 2001, App. B)

2.6.2 The relaxed and unrelaxed limits

The regimes identified in Section 2.2 correspond to two limiting moduli, and every frequency-dependent model used here is constructed to interpolate between them.

Relaxed limit. When the patch size is small compared with the diffusion length, pore pressure equilibrates fully within a wave period, the two fluids act as one at a common pressure, and the saturated modulus follows from Gassmann substitution with the Wood effective fluid of Equation (2.11):

$$K_{\mathrm{GW}} = K_{\mathrm{dry}} + \alpha^2 M\left(K_f^{\mathrm{W}}\right). \tag{2.14}$$

This is the soft limit, and it is the modulus that the low-frequency reference model of Chapter 3 predicts.

Unrelaxed limit. When the patch size is large compared with the diffusion length, no appreciable flow occurs between patches within a wave period. Each region responds as though it were separately and fully saturated, and the composite modulus is the saturation-weighted harmonic average of the separately saturated P-wave moduli, the Hill average:

$$\frac{1}{K_{\mathrm{GH}} + \frac{4}{3}G} = \frac{S_1}{K_{\mathrm{sat},1} + \frac{4}{3}G} + \frac{S_2}{K_{\mathrm{sat},2} + \frac{4}{3}G}, \tag{2.15}$$

with $G = G_{\mathrm{dry}}$ and $K_{\mathrm{sat},j}$ the single-fluid Gassmann moduli of the two regions. This is the stiff limit.

What the limits do and do not bound. The frequency-dependent response of a given patch distribution is bounded by these two limits,

$$K_{\mathrm{GW}} \le K(\omega) \le K_{\mathrm{GH}}, \tag{2.16}$$

and both the branching-function model and the exact solutions recover them as $\omega \to 0$ and $\omega \to \infty$ respectively. Recovering both limits is the primary internal check on every frequency-dependent computation reported in this thesis.

Equation (2.16) constrains the frequency-dependent response of a patchy medium, not an arbitrary static closure applied at the fluid level: averaging fluid moduli and averaging saturated rock moduli are different constructions. The bracketing of the Brie closure of Equation (2.13) was therefore verified explicitly over the modelled saturation range (Section 2.4.2) rather than assumed.

Two further consequences of Equation (2.16) shape how the results should be read. The width of the bracket is set by the frame as much as by the fluids: a stiffer dry frame narrows the separation between K_{GW} and K_{GH}, so the same patch distribution produces a smaller frequency-dependent excursion in a consolidated sandstone than in an unconsolidated one. And the bracket says nothing about where within it the response lies at a given frequency. That position is set by the relaxation frequency, which depends on the patch size, the permeability, and the fluid viscosity, and it is the quantity that determines whether the frequency dependence is visible in the seismic band at all.

2.6.3 The Johnson branching-function model

The main model we use here is the Johnson branching-function model (Johnson, 2001). In the form used in the active implementation, the dynamic bulk modulus is written directly in modulus space as

$$K(\omega) = K_{\mathrm{GH}} - \frac{K_{\mathrm{GH}} - K_{\mathrm{GW}}}{1 - \zeta + \zeta\sqrt{1 + i\omega\tau/\zeta^2}}, \tag{2.17}$$

which recovers the Gassmann–Wood and Gassmann–Hill bounds in the low- and high-frequency limits.

The two parameters in Eq. 2.17 have a direct physical interpretation. The relaxation time τ sets the frequency scale of the transition. The parameter ζ controls the shape of that transition. In the current implementation, we do not treat these two parameters as fitting terms. Instead, we derive them from the low-frequency parameter T and the high-frequency coefficient G following Johnson's asymptotic construction. Appendix C gives the supporting formulas.

We foreground this model in our work because it is the only implemented path that is both compact and scalable beyond a single geometry-specific boundary-value problem.

2.6.4 The White concentric-sphere model

White's model is the classical analytical treatment of patchy saturation in concentric spheres (White, 1975). In the corrected White family used here (Dutta & Odé, 1979; Dutta & Seriff, 1979; Carcione et al., 2003), the complex bulk modulus can be written compactly as

$$K(\omega) = \frac{K_\infty}{1 - K_\infty W(\omega)}, \tag{2.18}$$

where K_∞ is the high-frequency limit and $W(\omega)$ collects the diffusive coupling terms.

Here, we retain the White concentric-sphere model only as a benchmark. Its classical scope is the gas-inner concentric-sphere geometry. That makes it valuable in the shared benchmark domain, but too restrictive to function as the main scalable model. We collect the supporting formulas used in the implementation in Appendix E.

2.6.5 Exact reference solutions for simple geometries

Johnson (2001, Appendix B) also provides exact solutions for two simple geometries only: periodic slabs and concentric spheres. We retain these exact reference solutions in the codebase because they provide internal benchmark checks on the branching-function model. Their role in our work is narrow. They support validation in simple geometries, but they are not the main production path. We collect their notation mapping and implementation-support formulas in Appendix D.

2.7 Seismic observables from elastic and complex moduli

The thesis uses two observable families. The baseline chapter (Chapter 3) works in the relaxed Gassmann limit and reports intrinsic elastic-property observables that follow directly from the saturated moduli of Section 2.3: the bulk density ρ_{bulk} of Eq. (2.7), the P-wave velocity V_p of Eq. (2.8), the shear-wave velocity V_s of Eq. (2.9), and the derived ratio V_p/V_s together with the P-impedance $Z_p = \rho_{\mathrm{bulk}} V_p$. These quantities are frequency-independent in the relaxed limit because $G_{\mathrm{sat}} = G_{\mathrm{dry}}$ by Gassmann and ρ_{bulk} depends only on saturation.

The frequency-dependent chapter (Chapter 4) replaces K_{sat} with a complex, frequency-dependent bulk modulus $K(\omega)$ and works with the complex P-wave modulus

$$\mathcal{M}(\omega) = K(\omega) + \tfrac{4}{3}G_{\mathrm{dry}}. \tag{2.19}$$

The phase velocity and inverse quality factor follow as

$$V_p(\omega) = \left[\mathrm{Re}\left(\sqrt{\frac{\rho_{\mathrm{bulk}}}{\mathcal{M}(\omega)}}\right)\right]^{-1}, \tag{2.20}$$

$$Q^{-1}(\omega) = \frac{\mathrm{Im}[\mathcal{M}(\omega)]}{\mathrm{Re}[\mathcal{M}(\omega)]}. \tag{2.21}$$

For the 1D monitoring-workflow embedding in Chapter 4, the additional time-lapse proxies are

$$Z(\omega) = \rho_{\mathrm{bulk}} V_p(\omega), \tag{2.22}$$

$$R(\omega) = \frac{Z_{\mathrm{res}}(\omega) - Z_{\mathrm{cap}}}{Z_{\mathrm{res}}(\omega) + Z_{\mathrm{cap}}}, \tag{2.23}$$

$$\Delta t_{\mathrm{2way}}(\omega) = 2h\left[\frac{1}{V_p^{\mathrm{mon}}(\omega)} - \frac{1}{V_p^{\mathrm{base}}(\omega)}\right]. \tag{2.24}$$

In Eq. (2.23), the subscripts *res* and *cap* denote reservoir and caprock impedance, respectively. In Eq. (2.24), h is the reservoir thickness, and the superscripts *mon* and *base* denote monitor and baseline P-wave velocity, respectively. These three quantities require either a caprock impedance contrast or a reservoir thickness, and they therefore enter only at the Chapter 4 embedding stage, not in the baseline of Chapter 3.

Chapter 3

Reservoir Context and the Low-Frequency Reference Model

This chapter retains the work of Stekelenburg (2026) in compact form only. Its role is not to re-argue low-frequency detectability in full detail. Its role is to define the elastic-property reference state against which we judge the frequency-dependent results of Chapter 4.

The baseline assumes the relaxed Gassmann limit throughout: pore pressure equilibrates fully within one wave period, and each elastic-property observable is therefore frequency-independent. The full static sweep and extended figure set remain documented in Stekelenburg (2026).

3.1 Reservoir context

Chapter 1 introduced two Dutch sandstone reservoir groups as the leading candidates for porous-media UHS. The Permian Rotliegend sandstones provide many of the onshore Noord-Holland candidates and combine moderate porosity (roughly 0.18 to 0.22) with permeabilities in the 50 to 500 mD range. The Triassic Bunter sandstones in Zuid-Holland and the wider southern North Sea bring stronger layering and depositional heterogeneity (Reijnen-Mooij et al., 2025; Loyola et al., 2026). The national screening reports identify 27 candidate fields in Noord-Holland and Zuid-Holland alone; across the screened set, abandonment pressures span 10 to 200 bar and reservoir depths typically exceed 1000 m. Figure 3.1 shows the schematic distribution of the two formation groups across the Netherlands.

These two formations are chosen because they are the UHS-suitable candidates identified by the national screening (Reijnen-Mooij et al., 2025; Loyola et al., 2026), not because of any incidental match to specific MacBeth Table 2 rows. The two-group scoping is introduced by the present thesis in line with the Netherlands-specific framing of Section 1.2; the Research Module of Stekelenburg (2026) ran on a single generic constant-frame parameter point and did not commit to formation-specific cases. The baseline workflow remains at the rock-physics and 1D monitoring scale, so the two formations are best read as parameter envelopes within which the workflow is exercised, rather than as commitments to specific fields.

3.1.1 Our reference cases

Throughout Chapters 3 and 4 we work with three reference cases drawn from this envelope. Table 3.1 collects their parameter values in one place; all later figures

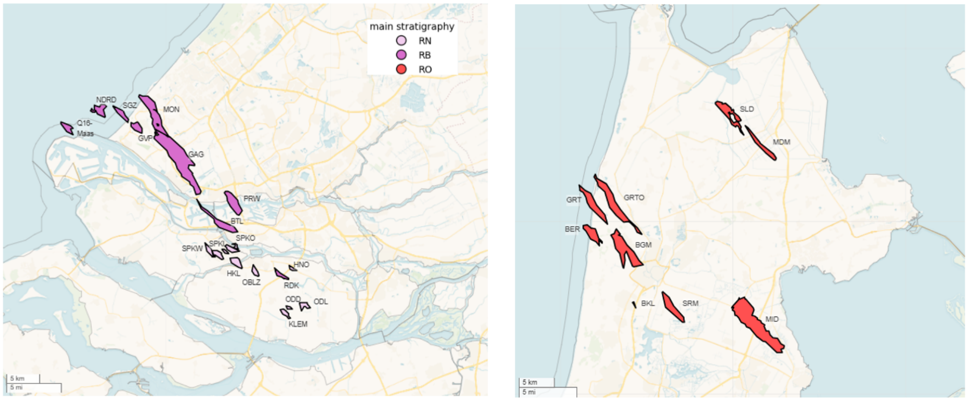

Figure 3.1: Locations and main reservoir stratigraphy of candidate underground hydrogen storage fields: (left) Upper- and Middle-Bunter candidates in Zuid-Holland; (right) Rotliegend candidates in Noord-Holland. RN, RB, and RO denote Upper Bunter, Middle Bunter, and Rotliegend, respectively. Adapted from Reijnen-Mooij et al. (2025, Fig. 21).

and sweeps in this thesis derive their case parameters from this table. The Control case is a generic constant-frame point at a representative mid-depth, mid-porosity location and anchors the baseline workflow. The Rotliegend-style and Bunter-style cases adopt porosity and pore pressure from cluster centroids in Reijnen-Mooij et al. (2025, Fig. 23) and use MacBeth (2004) Table 2 rows for the dry-frame parameterisation: Rotliegend Facies B (italic re-fit row) for the Rotliegend-style case and Berea Sandstone as the closest clean continental fluvial analogue for the Bunter-style case. Berea is adopted because the MacBeth (2004) database does not contain Triassic Buntsandstein samples; the rationale for choosing Berea over the available Forties facies rows is given in Appendix B. The reservoir-style labelling preserves the parameter-envelope reading and avoids any claim that these cases represent specific named fields.

Table 3.1: Thesis reference cases. MacBeth parameters are from MacBeth (2004, Table 2): the Rotliegend-style case uses the Rotliegend Facies B italic re-fit row; the Bunter-style case uses the Berea Sandstone row. Porosity, depth, pore-pressure, and permeability ranges quoted in the "HyTROS range" column are read from the HyTROS Figures 22, 23, 24, 26 and 28 stratigraphic-group distributions (Reijnen-Mooij et al., 2025); the reference-case values are cluster-centroid reads from those distributions. RO denotes Rotliegend, UB Upper Bunter, MB Middle Bunter. Differential pressure and derived dry-frame moduli are computed from Eq. (B.2) and Eqs. (B.6)–(B.7), respectively, using $\rho_{\text{ov}} = 2300\ \text{kg/m}^3$.

Parameter	Symbol	Control	Rotliegend-style	Bunter-style	HyTROS range	Source
Depth	z	1000 m	2800 m	1750 m	RO: 2000–2300 m; UB: 1700–2750 m; MB: 2600–3500 m	HyTROS Fig. 22
Porosity	ϕ	0.15	0.18	0.22	RO: 0.16–0.23; UB: 0.10–0.23; MB: 0.05–0.22	HyTROS Fig. 22
Pore pressure	p_p	9.91 MPa	30 MPa	22 MPa	10–200 bar across all groups; field-dependent	HyTROS Fig. 26, 28
Salinity	S_b	0.15	0.20	0.15	data sparse; representative envelope only	HyTROS §6.2 salinity
Frame model	n/a	constant	MacBeth	MacBeth	n/a	Sec. 2.5
MacBeth K_∞	K_∞	n/a	22.39 GPa	18.27 GPa	n/a	MacBeth Tab. 2
MacBeth G_∞	G_∞	n/a	17.48 GPa	17.43 GPa	n/a	MacBeth Tab. 2
MacBeth S_K	S_K	n/a	0.62	0.38	n/a	MacBeth Tab. 2
MacBeth S_G	S_G	n/a	0.42	0.38	n/a	MacBeth Tab. 2
MacBeth P_K	P_K	n/a	15.75 MPa	11.91 MPa	n/a	MacBeth Tab. 2
MacBeth P_G	P_G	n/a	56.47 MPa	8.16 MPa	n/a	MacBeth Tab. 2
Differential pressure	p_d	12.8 MPa	33.3 MPa	17.6 MPa	n/a	derived
Dry-frame K at p_d	K_{dry}	12.0 GPa	18.70 GPa	16.03 GPa	n/a	derived
Dry-frame G at p_d	G_{dry}	10.0 GPa	12.47 GPa	16.27 GPa	n/a	derived
Reservoir thickness	h	100 m	80 m	100 m	sand thickness varies by group, see HyTROS Fig. 24	this work
Permeability*	κ	100 mD	50 mD	50 mD	RO: 50–500 mD; UB: 50–1000 mD; MB: 10–500 mD	HyTROS Fig. 24
Patch radius*	a	5 cm	5 cm	15 cm	n/a	this work

*Permeability and patch radius enter only the frequency-dependent stage in Chapter 4 and are listed here for completeness; the baseline of Chapter 3 does not use them. The within-cluster uncertainty is quantified explicitly in Section 3.2.7.

3.1.2 Previous findings carried forward

The Research Module (Stekelenburg, 2026) ran the same Gassmann plus Wood plus MacBeth plus Batzle–Wang plus Peng–Robinson EOS workflow across the full parameter envelope at a single generic constant-frame point. It reported the response in time-lapse difference proxies ($\Delta Z/Z$, ΔR, $\Delta t_{2\text{way}}$) rather than the intrinsic elastic properties used in the present chapter. Three findings from that work are carried forward and shape the present baseline rather than being re-derived here.

First, the brine-to-H_2 response is dominated by a steep drop at low gas saturation: roughly half of the full-saturation response is reached by $S_{H_2} \approx 0.10$, and a substantial fraction is reached by $S_{H_2} \approx 0.05$. This is what the Wood-mixing column of Appendix A.3 shows in the difference-proxy language; the same shape is recovered in the intrinsic elastic properties of Figure 3.2 in Section 3.2.3. This shape is the reason the present thesis cares about the frequency-dependent response in the low-saturation regime specifically: that is where the detectability question is decided, and that is where the relaxed Gassmann assumption is most stressed.

Second, the Wood-versus-Brie sensitivity sweep (Appendix A.3, reproduced from the Research Module) showed that Brie mixing softens and delays the low-saturation response by an amount that scales with the stiffening exponent, while preserving its sign and qualitative shape. Wood remains the active closure because it defines the relaxed Gassmann–Wood limit and produces the strongest low-saturation response, the regime where the detectability question lives; the selection rationale is given in Section 2.4.

Third, of the three dry-frame parameterisations compared there, the MacBeth frame is the one that maps directly onto the Dutch candidate formations, and it is adopted as the default reservoir frame with the constant frame retained for the Control case; the selection rationale is given in Section 2.5, and the resulting frame stiffness across differential pressure is shown in Section 3.2.5.

The remaining ingredients, the Batzle–Wang brine correlations and the Peng–Robinson gas-property path, are inherited unchanged (Section 2.1, Appendix A).

What the present chapter adds, on top of these inherited findings, is the translation from the Research Module's difference-proxy language into the intrinsic elastic-property observables that Chapter 4 extends. The elastic-property baseline curves, the Control-point modulus path, and the HyTROS-justified sub-case sweeps are the contributions beyond the Research Module.

3.2 Workflow and elastic-property response

The baseline workflow models a single uniform sandstone reservoir layer of thickness h and evaluates brine-to-H_2 substitution under the low-frequency Gassmann assumption. The subsections below state the workflow inputs, review the substitution mechanism, present the three-case elastic-property baseline, and quantify its salinity and porosity sensitivities, closing with the aquifer-initial reference state.

3.2.1 Workflow inputs and parameter envelope

The workflow choices identified in Section 3.1.2 (Wood mixing, MacBeth reservoir frame, constant Control frame, Batzle–Wang brine, Peng–Robinson gas) are applied here without further discussion. Their literature defense is given in Chapter 2 (Section 2.4 and Section 2.5).

Pressure and temperature are mapped from depth by linear gradients (Table 3.2). The hydrostatic gradient sets the pore pressure $p_p(z)$; an overburden gradient sets the confining pressure $p_c(z)$ and hence the differential pressure $p_d = p_c - p_p$. The geothermal gradient sets the in-situ temperature $T(z)$. Salinity S_b is specified directly per case. The depth-to-pressure mapping is repeated for reference in Appendix B.3.

Table 3.2: Linear gradients used in the baseline workflow to map depth to pore pressure, confining pressure, and temperature. The same gradients are used in the reference-case sweeps of Chapter 4.

Quantity	Symbol	Value	Notes
Surface temperature	T_0	10 °C	this work
Geothermal gradient	dT/dz	30 °C/km	this work
Hydrostatic gradient	dp_p/dz	9.81 MPa/km	freshwater g
Overburden gradient	dp_c/dz	22.56 MPa/km	$\rho_{\mathrm{ov}} = 2300$ kg/m^3

The Control reference parameter point (depth 1000 m, salinity 0.15) sits inside the envelope of effective pressures 10 to 200 MPa and reservoir temperatures 30 to 70 °C that bracket the candidate Dutch UHS reservoirs reported by Reijnen-Mooij et al. (2025) and Loyola et al. (2026). Above 70 °C hydrogen enters a thermodynamic regime where Peng–Robinson predictions become less reliable; below 30 °C microbial reactions on H_2 become a separate risk that is not addressed here. Within this envelope, the workflow has been exercised in the Research Module across the full ranges (Stekelenburg, 2026); the qualitative trends in the elastic-property observables are robust to changes in pressure, temperature, and salinity, with the modulus contrast driven mainly by the gas-versus-brine compressibility ratio, which is largely insensitive to small variations in those three parameters at the depths of interest.

3.2.2 Gassmann substitution mechanism

Gassmann substitution (Gassmann, 1951; Smith et al., 2003), Eq. (2.5), computes the saturated bulk modulus of a porous rock from its dry-frame bulk modulus, the mineral bulk modulus, the porosity, and an effective fluid bulk modulus. The shear modulus is unchanged by the fluid, Eq. (2.6); only K_{sat} depends on saturation. The effective fluid modulus in the relaxed limit is the Wood harmonic average of brine and gas moduli, weighted by saturation (White, 1975), Eq. (2.11). Replacing brine with H_2 at modest gas saturation drives a steep drop in the Wood-averaged fluid modulus, because gas is nearly three orders of magnitude more compressible than brine, and that drop propagates through Eq. (2.5) into a reduced K_{sat} and a reduced P-wave velocity (Eq. (2.8)). The same fluid-substitution step also reduces bulk density (Eq. (2.7)), because gas is much less dense than brine.

We report this prediction through three intrinsic elastic-property observables that are functions of gas saturation alone in the relaxed limit: the bulk density $\rho_{\mathrm{bulk}}(S_g)$ of Eq. (2.7), the velocity ratio $V_p(S_g)/V_s(S_g)$ formed from Eqs. (2.8) and (2.9), and the P-impedance $Z_p(S_g) = \rho_{\mathrm{bulk}}(S_g)\,V_p(S_g)$. Time-lapse difference quantities such as $\Delta Z/Z$, ΔR, and Δt_{2way} are observables of the Chapter 4 embedding stage, where the caprock impedance and reservoir thickness enter; they are reported there, not here.

3.2.3 Three-case low-frequency baseline

Figure 3.2 shows the low-frequency Gassmann response of the three reference cases of Table 3.1 under Wood fluid mixing. The Control case recovers the representative constant-frame point used to anchor the baseline workflow; the Rotliegend-style and Bunter-style cases combine the reservoir parameter picks of Table 3.1 with their MacBeth dry frames. All three cases retain the same first-order behaviour: V_p/V_s falls sharply at low gas saturation and then flattens; bulk density falls roughly linearly with saturation; P-impedance falls steeply at low saturation and then changes more gradually, broadly like V_p/V_s, but retains a clearer downward inclination above $S_{H_2} \approx 0.10$. The reservoir-style cases sit above the Control case in V_p/V_s and Z_p at all saturations because their pressure-stiffened MacBeth frames carry larger dry-frame moduli (see Section 3.2.5).

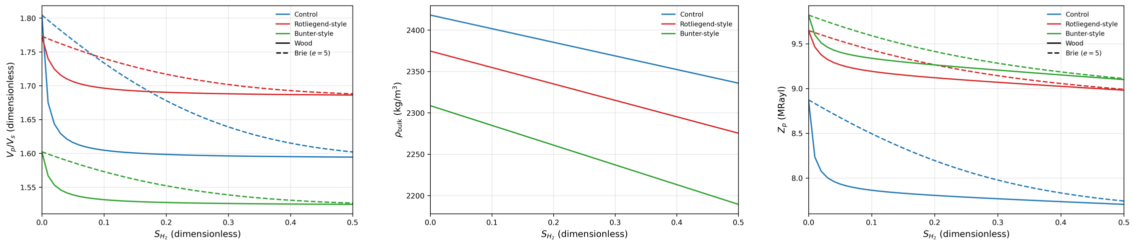

(a) (b) (c)

Figure 3.2: V_p/V_s, bulk density ρ_{bulk}, and P-wave impedance Z_p are shown against hydrogen saturation for the three reference cases. The V_p/V_s and Z_p panels show the low-frequency Gassmann–Wood limit and the Brie curve with stiffening exponent $e = 5$; the density panel carries one curve per case because fluid densities mix linearly and the two closures coincide there. (a) V_p/V_s. (b) ρ_{bulk}. (c) Z_p. The Control case uses a constant dry frame; the other cases use pressure-sensitive MacBeth frames.

3.2.4 Modulus path at the Control reference point

The mechanism behind the saturation curves of Figure 3.2 is most direct at the Control reference point. Figures 3.3 and 3.4 show the two steps that connect fluid substitution to the elastic-property response at this point. Figure 3.3 shows the Wood-mixed effective fluid bulk modulus $K_f^{\text{W}}(S_{H_2})$ together with the brine and hydrogen endmember bulk moduli (K_b and K_g) at the Control pressure-temperature state. Figure 3.4 shows the resulting Gassmann saturated bulk modulus $K_{\text{sat}}(S_{H_2})$ together with the dry-frame floor $K_{\text{dry}} = 12$ GPa.

The Wood average drops steeply at low gas saturation because $K_g \ll K_b$ at the Control state, so a small gas volume fraction already pulls K_f^{W} down by nearly an order of magnitude. That drop is propagated through Eq. (2.5) into K_{sat}, which approaches the dry-frame floor as S_{H_2} grows. The same coupling between fluid mixing and saturated modulus underlies the Control curves in Figure 3.2; it is shown explicitly here because the saturation curves of the derived observables compress the underlying modulus story into ratios and impedance products.

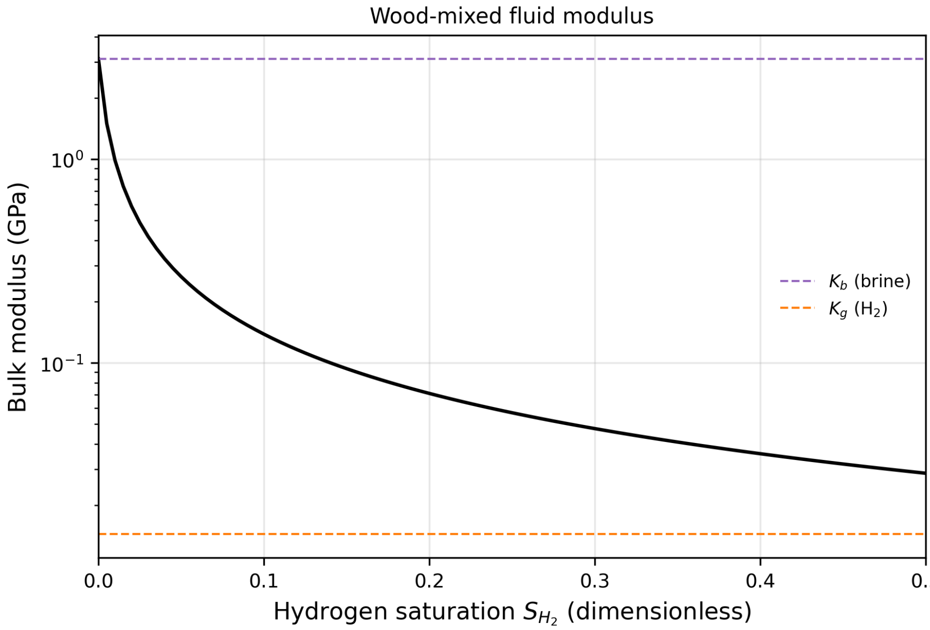

Figure 3.3: Wood-mixed effective fluid bulk modulus K_f^{W} is shown against hydrogen saturation at the Control reference point. The brine and hydrogen endmember bulk moduli are included for comparison.

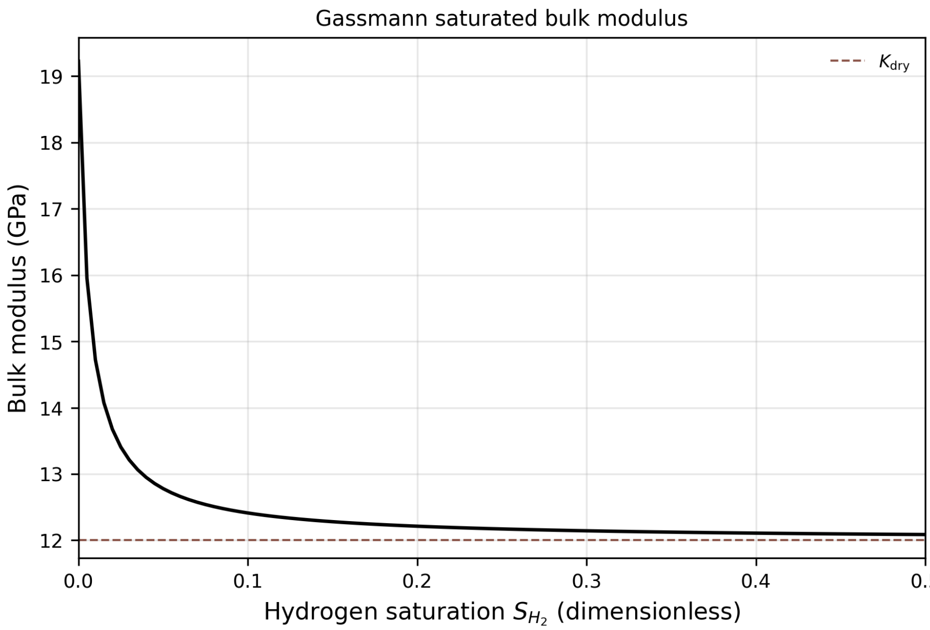

Figure 3.4: Gassmann saturated bulk modulus K_{sat} is shown against hydrogen saturation at the Control reference point. The dry-frame modulus is included as its lower limit.

3.2.5 Frame-stiffness diagnostic

The frame stiffness that drives the magnitude differences between the three cases in Figure 3.2 is shown directly in Figure 3.5. The Control frame is constant by construction. The MacBeth Rotliegend Facies B and Berea parameterisations both stiffen with differential pressure, and at the operating differential pressures of the two reservoir cases (vertical markers) the resulting K_{dry} values are several times the Control floor. This frame-stiffness picture carries forward unchanged into the frequency-dependent stage of Chapter 4.

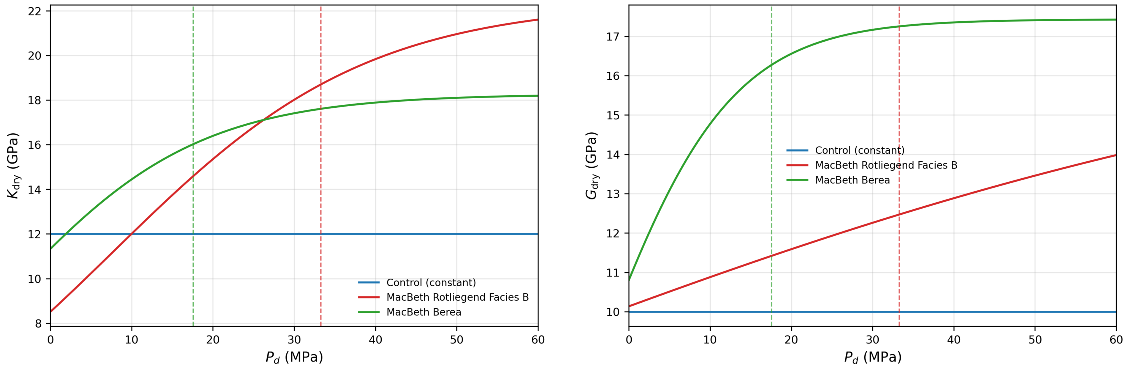

(a) (b)

Figure 3.5: Dry-frame bulk and shear moduli are shown against differential pressure for the three reference cases. (a) K_{dry}. (b) G_{dry}. Vertical markers indicate the reservoir differential pressure of each case; the Control moduli are constant and the other cases follow MacBeth pressure-sensitive frames.

3.2.6 Salinity sensitivity at the Control reference point

Salinity affects the brine bulk modulus and density through the Batzle–Wang correlations (Batzle & Wang, 1992) and therefore the brine endmember of the Wood mixing. Reijnen-Mooij et al. (2025) noted that publicly available salinity data for Dutch formation waters is limited, so the salinity values used here should be read as a representative range rather than a field-specific value. Figure 3.6 overlays the saturation response at the Control parameter point (depth 1000 m, $\phi = 0.15$, $K_{\text{dry}} = 12$ GPa, $G_{\text{dry}} = 10$ GPa) for brine salinities $S_b \in \{0.05, 0.10, 0.15, 0.20\}$ weight fraction. The elastic-property response is mildly sensitive to salinity across this range: higher salinity stiffens the brine endmember and shifts V_p/V_s and Z_p upward across all saturations, while bulk density increases roughly linearly with salinity. The brine-to-H_2 separation remains the dominant feature in all three panels. Because the salinity effect is small compared to the porosity and frame effects discussed below, salinity sensitivity is not pursued further at the reservoir sub-cases.

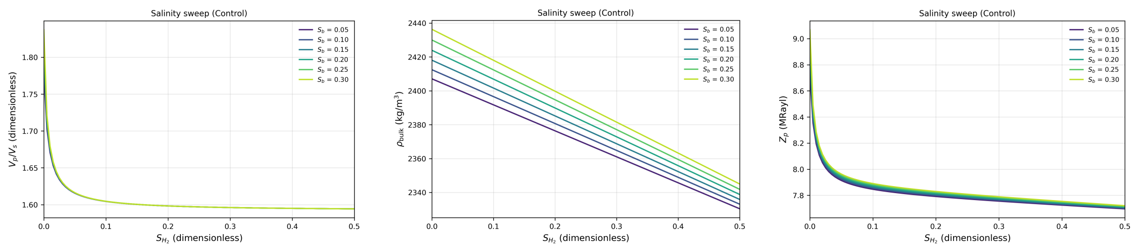

(a) (b) (c)

Figure 3.6: V_p/V_s, bulk density ρ_{bulk}, and P-wave impedance Z_p are shown against hydrogen saturation for salinity weight fractions $S_b = 0.05$ to 0.30 at the Control point. (a) V_p/V_s. (b) ρ_{bulk}. (c) Z_p.

3.2.7 Porosity sensitivity across HyTROS reservoir sub-cases

The Rotliegend-style and Bunter-style reference cases of Table 3.1 are cluster-centroid reads from the HyTROS field distribution (Reijnen-Mooij et al., 2025, Fig. 22 and Fig. 23), not field-specific values. To make the within-cluster uncertainty explicit, we evaluate the low-frequency baseline response at three reservoir sub-cases drawn directly from the HyTROS stratigraphic split, with porosity ranges defended against the HyTROS Figure 22 distributions.

The three sub-cases are: Rotliegend, Upper Bunter, and Middle Bunter. Rotliegend is treated as a single depth-porosity cluster because the HyTROS Figure 22 Rotliegend top-reservoir distribution is narrow ($\sim$ 2000 to 2300 m TVDSS (metres true vertical depth subsea)) and the porosity distribution is tightly clustered around 0.20 ($\sim$ 0.16 to 0.23). Bunter is split into Upper Bunter (shallower, higher porosity, located south and south-east of Rotterdam) and Middle Bunter (deeper, lower porosity, with the deepest fields trailing toward $\phi \approx 0.04$) following the HyTROS stratigraphic separation, because the porosity-depth trend within Bunter is too wide to be captured by a single sub-case. Salinity is held at the formation-appropriate value from Table 3.1 (0.20 for Rotliegend, 0.15 for both Bunter sub-cases); the salinity dependence within each sub-case is expected to be small based on Figure 3.6.

For each sub-case the dry-frame moduli are evaluated at the corresponding differential pressure using the MacBeth (2004) parameter row of Table 3.1 (Rotliegend Facies B for the Rotliegend sub-case, Berea Sandstone for both Bunter sub-cases). The computed $(K_{\text{dry}}, G_{\text{dry}})$ values appear in the figure annotations.

The Rotliegend-style and Bunter-style cases in Table 3.1 lie within their respective sub-case envelopes but not at the cluster median: the Rotliegend-style case at $z = 2800$ m, $\phi = 0.18$ lies on the deep, low-porosity edge of its distribution, whereas the Bunter-style case at $z = 1750$ m, $\phi = 0.22$ lies on the shallow, high-porosity edge of the Upper Bunter distribution. Figure 3.7 provides the exemplar, and the same porosity sensitivity was evaluated for the Upper and Middle Bunter intervals in Section A.4; both follow the Rotliegend pattern, although the Middle Bunter sweep extends to lower porosity and a deeper reference depth and therefore spans a broader elastic-property envelope. These sweeps intentionally bracket rather than reproduce the downstream reference cases, which remain at the Table 3.1 values in Chapter 4.

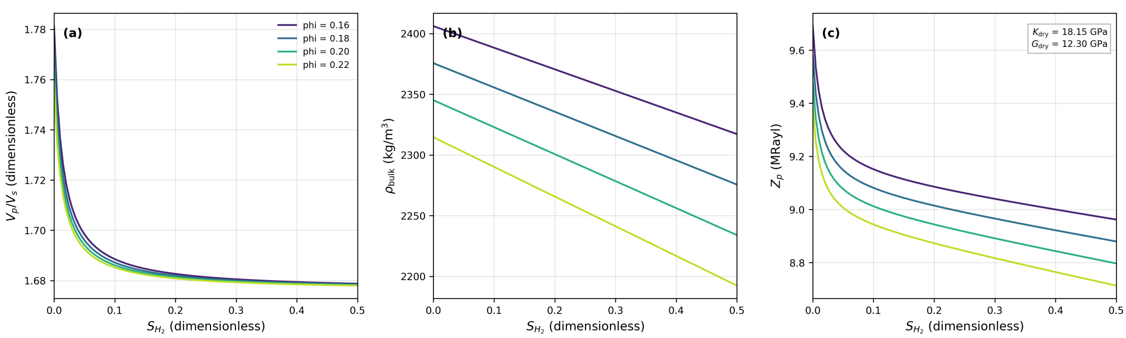

Figure 3.7: V_p/V_s, bulk density ρ_{bulk}, and P-wave impedance Z_p are shown against hydrogen saturation for Rotliegend porosities $\phi \in \{0.16, 0.18, 0.20, 0.22\}$ at $z = 2400$ m and $S_b = 0.20$. (a) V_p/V_s. (b) ρ_{bulk}. (c) Z_p.

3.2.8 Aquifer-initial reference state

The baseline workflow as described above is, by construction, an aquifer-initial reference state: the pre-injection rock is fully brine-saturated, with no residual gas phase. This is the appropriate reference for a saline-aquifer storage scenario, but it is not the realistic initial state of a depleted gas-field UHS target, where a residual methane cushion remains after production. The depleted-gas-field case has been treated by Bagheri et al. (2026) for a heterogeneous Dutch reservoir with a methane cushion gas, using the same Wood plus Gassmann closure that underpins the present baseline. Their analysis confirms that hydrogen enrichment produces measurable time-lapse signatures in that setting under the low-frequency assumption. We therefore retain the aquifer-initial baseline as the reference state for the rest of this work, and treat the extension to a depleted-gas initial condition with residual CH_4 as a planned realism step rather than as a current result. The frequency-dependent question addressed in Chapter 4 is independent of this distinction: it asks how the relaxed Gassmann assumption breaks down once mesoscopic patchiness is present, and that question is sharpened, not softened, by working from the cleaner aquifer-initial reference.

3.3 Why the low-frequency reference model is retained

The low-frequency baseline remains useful for three reasons.

First, it fixes the reference state. Chapter 4 does not revisit whether H_2 substitution should produce any seismic response at all. That question is already answered positively in the low-frequency framework. The Chapter 4 question is narrower: does frequency dependence change the interpretation enough to matter?

Second, the baseline provides trend expectations for the elastic-property observables. If the active frequency-dependent workflow produces V_p/V_s, ρ_{bulk}, or Z_p behaviours that reverse their established low-frequency trends, that would be scientifically important. If it preserves the same signs but changes the magnitude, the interpretation is different.

Third, the baseline defines the relaxed limit that the frequency-dependent models recover as $\omega \to 0$; the Gassmann–Wood state is a physical anchor in the active workflow, not only a narrative reference.

3.4 Why the low-frequency reference model is not sufficient

The baseline is intentionally incomplete for the question we address in our work.

The main limitation is that the baseline contains no frequency dependence: every elastic-property observable is a single number, independent of seismic frequency. That is justified only when pore pressure equilibrates between gas-rich and brine-rich regions within a wave period. For Dutch depleted-gas-field conditions and for plausible patch sizes generated by H_2 injection, the relaxation frequency set by patch size and permeability (Section 2.2) can fall within or near the seismic band, where the relaxed assumption fails and the response becomes dispersive and attenuating. The extension developed in Chapter 4 addresses exactly this regime.

Two further limitations remain important but are not the immediate focus at this stage. The initial state in the baseline is fully brine- saturated, whereas practical UHS targets are more likely to be depleted gas fields with residual CH_4. The reservoir geometry is a single uniform layer, so layering, transition zones, and more realistic saturation topology are absent. These realism steps remain open, but they are only worth prioritising after we isolate the frequency-dependent effect itself and embed it back into the monitoring workflow.

Chapter 4

Frequency-Dependent Modelling

This chapter presents the frequency-dependent patchy-saturation modelling that extends the low-frequency baseline of Chapter 3: the physical basis for patchy saturation, the model hierarchy, the benchmark validation, an idealised parameter scan, the sphere-versus-slab geometry sensitivity, and the embedding of the frequency-dependent modulus into the three reference cases of Table 3.1.

4.1 Physical basis for patchy saturation

The baseline of Chapter 3 treats the pore fluid as a single effective phase obtained by Wood mixing of brine and hydrogen, which is then substituted into the dry frame through the Gassmann equation. That construction assumes the two fluids are mixed uniformly at the pore scale, so that wave-induced pore-pressure increments equilibrate instantaneously everywhere. Laboratory and field evidence indicates that pore-scale uniform mixing is the exception rather than the rule in reservoirs that have undergone drainage, injection, or depletion. Hydrogen and brine instead segregate into clusters, or patches, at scales much larger than a pore but smaller than a seismic wavelength.

Several mechanisms drive this segregation. Capillary entry pressure differs between pore sizes, so during drainage the non-wetting gas invades the larger pores first and is stranded as disconnected clusters in the finer pore network, whereas imbibition produces more regular distributions. Direct X-ray and ultrasonic imaging of partially saturated sandstone confirms that drainage leaves nonuniform gas clusters and that the resulting fluid distribution, not saturation alone, controls the measured velocities (Lebedev et al., 2009; Müller et al., 2010). The low viscosity of hydrogen further promotes preferential flow into high-permeability bands rather than uniform frontal advance, and cyclic injection and production traps hydrogen within those permeable zones, giving a plume geometry more complex than a uniform cloud (Bagheri et al., 2026). Geological layering, permeability streaks, and gravity segregation reinforce the same outcome.

Uniform saturation is therefore not a separate physical state from patchy saturation but its low-frequency limit. Whether a given patch responds as uniformly mixed or patchy is set by the comparison of Section 2.2: at frequencies low enough that the diffusion length $L_d \sim \sqrt{D/\omega}$ of Equation (2.4) exceeds the patch size, pore pressure equilibrates across the patch within a wave period and the relaxed description holds; at higher frequencies it does not. Because L_d depends on the local permeability and viscosity, and the patch size on the local saturation distribution, this comparison is made cell by cell rather than once per reservoir: parts of the same plume can respond as uniformly mixed while other parts are patchy at the same frequency. The question this chapter addresses is whether, at seismic frequencies and for hydrogen-brine patches in Dutch sandstone, the equilibrated regime still applies, and how large the discrepancy is where it does not.

The patchy-saturation models used here apply on a representative elementary volume that is large compared with a single patch but small compared with a seismic wavelength. This scale ordering, $\lambda_{\text{seismic}} \gg L_{\text{REV}} \gg a$, defines the mesoscopic regime of Section 2.2 in which wave-induced fluid flow operates (Müller et al., 2010). Below the lower bound, of order a pore diameter, the relevant dissipation mechanism is microscopic squirt flow rather than mesoscopic patch-scale flow; above the upper bound, of order a wavelength, patches become quasi-resolvable and scattering becomes significant. For seismic frequencies this places the operational patch-size envelope at roughly 1 mm to 1 m. The patch radii considered in this chapter (5 to 20 cm) fall within this envelope.

4.2 Model hierarchy and governing formulas

The frequency-dependent workflow replaces the static partial-saturation closure of the baseline (Wood or Brie mixing followed by Gassmann substitution) with a complex, frequency-dependent bulk modulus $K(\omega)$. All other workflow components, including fluid properties, dry-frame models, and downstream observables, remain as defined in Chapter 2.

The Johnson branching-function model (Johnson, 2001) is the principal frequency-dependent patchy-saturation method used in this thesis. It parametrises the transition between the relaxed (Gassmann–Wood) and unrelaxed (Gassmann–Hill) limits via two analytically computed parameters, without requiring a new boundary-value solution for each patch geometry. This makes it the more applicable modelling path at this stage of our work (Müller et al., 2010).

The dynamic bulk modulus is that of Equation (2.17), which interpolates between the Gassmann–Wood and Gassmann–Hill limits of Eqs. (2.14)–(2.15) through the analytically computed pair (τ, ζ) of Johnson (2001, Eqs. (44)–(45)): the relaxation time τ sets the frequency scale of the transition, the shape parameter ζ its width and asymmetry, and neither is a free fitting parameter. For the benchmark geometries tested in this chapter, the values of ζ that Johnson (2001) reports on Fig. 1 span from 0.43 to 2.35. Appendix C gives the full derivation chain from poroelastic inputs to (τ, ζ).

Two structural properties of Eq. (2.17) should be made explicit before it is applied. First, it is an analytic interpolation between the two asymptotic limits rather than a numerical solution to a boundary-value problem. Johnson (2001) shows that this branching function satisfies the causality requirement that all singularities lie on the negative imaginary ω-axis, and that it reproduces both the low-frequency and high-frequency asymptotic series of the exact concentric-sphere and periodic-slab solutions by construction. The benchmark in Section 4.3 quantifies the residual deviation from the exact reference solutions on the geometries where those solutions exist. Second, the closed-form structure means that geometry enters $K(\omega)$ only through the two scalars (τ, ζ), and that evaluating $K(\omega)$ at one parameter point is cheap. This is what makes the same modelling stack reusable cell-by-cell when the workflow is later extended to heterogeneous layered reservoirs and to two-dimensional distributions of patch parameters; an exact boundary-value solver would have to be re-derived for each new geometry.

Two limitations follow from the same closed-form structure. First, Eq. (2.17) produces a single relaxation peak per parameter point, set by a single τ. A distribution of patch sizes in a real reservoir would superpose multiple such peaks and yield a broader, smoother $Q^{-1}(\omega)$ response. The standard route to capture this is a composite-frequency-function construction over the patch-size distribution (Pride et al., 2004; Müller et al., 2010), which is identified as a planned extension. Second, the analytically derived (τ, ζ) used here are specific to a single concentric-sphere gas-inner geometry. Other patch geometries have different (τ, ζ) (Johnson, 2001, Fig. 1), and arbitrarily shaped patches require either a closed-form T and S/V when one exists or a numerical evaluation otherwise (Müller et al., 2010). The benchmark validation in Section 4.3 covers the concentric-sphere case used throughout this chapter.

The White concentric-sphere model (White, 1975; Dutta & Odé, 1979; Dutta & Seriff, 1979; Carcione et al., 2003) provides an independent benchmark in the concentric-sphere domain, using $K(\omega) = K_\infty/(1 - K_\infty W(\omega))$ as stated in Eq. (2.18). It is restricted to the concentric-sphere geometry but valuable as an external cross-check. Implementation details are given in Appendix E.

Exact reference solutions for periodic slabs and concentric spheres, based on the boundary-value construction in Johnson (2001, Appendix B), serve as internal benchmark support. They provide geometry-specific solutions against which both the branching-function model and its asymptotic parameters can be tested. We collect the relevant formulas in Appendix D.

A final modelling choice concerns the relationship between patch geometry and reservoir depth. The concentric-sphere gas-inner geometry used throughout this chapter is most defensible where hydrogen is present as discrete, disconnected clusters in a brine continuum, which corresponds to residual or capillary-trapped gas after drainage (Lebedev et al., 2009; Müller et al., 2010). It is a stronger idealisation in two other settings that occur within a depleted-field storage column. Directly beneath the caprock a hydrogen gas cap is effectively single-phase and near-uniform in saturation, so its modulus is given by Gassmann substitution with pure hydrogen and no patchy correction is required; the relevant boundary there, the gas-water contact, is a wavelength-scale interface that enters the model as an impedance contrast rather than as a patch. In the transition zone around that contact, and along high-permeability streaks through which hydrogen has preferentially migrated, the gas-brine alternation is set by bedding and by capillary entry pressure against near-horizontal layering, so a periodic-slab geometry is geometrically more faithful than a sphere. The slab case is the second geometry for which Johnson (2001, Appendix B) provides an exact construction of (τ, ζ), and it enters Eq. (2.17) at the same closed-form cost as the sphere. This chapter retains the single concentric-sphere geometry for all reservoir cells so that the frequency-dependent mechanism is isolated from geometric variation. The depth-dependent geometry refinement, in which the patch model is selected per cell according to the local saturation state, is identified as the natural next step on the modelling ladder and is carried into the scenario modelling that follows.

4.3 Benchmark validation

4.3.1 Johnson Fig. 1 reproduction

The validation reproduces the benchmark behaviour shown in Johnson (2001, Fig. 1) for three geometries: concentric spheres with a gas-inner sphere, concentric spheres with a gas-outer shell, and periodic slabs. The parameter set follows Johnson (2001, Table 1): $\phi = 0.284$, $K_0 = 35$ GPa, $K_{\mathrm{dry}} = 2.637$ GPa, $G_{\mathrm{dry}} = 1.740$ GPa, $\kappa = 10^{-13}$ m^2 (≈ 0.1 mD), and $S_g = 0.10$.

Figure 4.1 shows, for each geometry, the exact reference solution (solid lines), the branching-function model (dashed), and the low-frequency and high-frequency asymptotic expressions used to compute τ and ζ. The Gassmann–Wood and Gassmann–Hill limits appear as horizontal bounds.

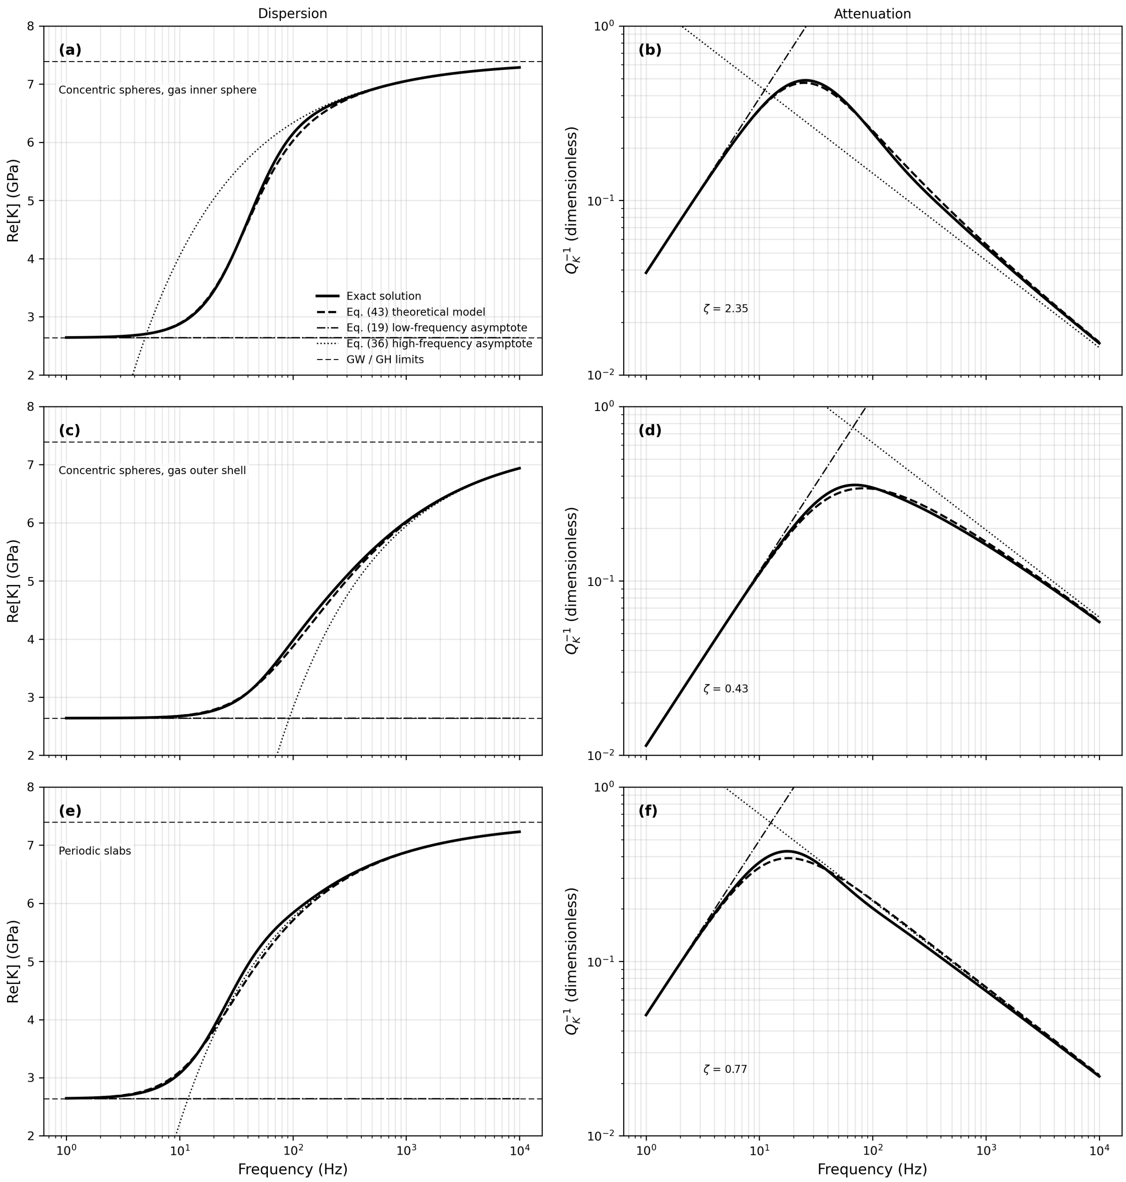

Figure 4.1: Bulk-modulus dispersion Re$[K]$ and attenuation Q_K^{-1} are shown against frequency for three benchmark geometries. (a) and (b) Gas-inner sphere; (c) and (d) gas-outer sphere; (e) and (f) periodic slabs, with dispersion followed by attenuation in each pair. Solid curves are exact solutions, dashed curves are the branching-function model, and the other curves mark the asymptotes and Gassmann bounds.

The branching-function model tracks the exact reference solutions closely across the full transition for all three geometries. The computed ζ values reproduce those that Johnson (2001) annotates on Fig. 1: $\zeta \approx 2.35$ for the gas-inner sphere, $\zeta \approx 0.43$ for the gas-outer shell, and $\zeta \approx 0.77$ for the periodic slab. The gas-inner sphere value exceeds unity, which is the regime in which the branching function develops a simple pole and the attenuation Q^{-1} overshoots its high-frequency asymptote; this overshoot is visible in the corresponding panel and is reproduced by the implementation. Johnson (2001) notes that the exact solution does not share this pole structure (it has instead a countable sequence of poles on the negative imaginary axis), so the branching function slightly misrepresents the singularity structure for $\zeta > 1$ while still tracking the real-frequency curve to within a few percent. Automated checks confirm that $\mathrm{Re}[K] > 0$ everywhere, $\mathrm{Re}[K]$ at low frequency recovers K_{GW} to within 1%, $K_{\mathrm{GW}} \leq \mathrm{Re}[K] \leq K_{\mathrm{GH}}$, and $\mathrm{Im}[K] \geq 0$ (passive dissipation). All checks pass for all three geometries.

4.3.2 White cross-check

The White model applies only in the gas-inner concentric-sphere domain. There, the Johnson branching-function model, the exact concentric-sphere reference solution, and the White concentric-sphere model can be compared directly. Figure 4.2 shows this comparison in both dispersion and attenuation.

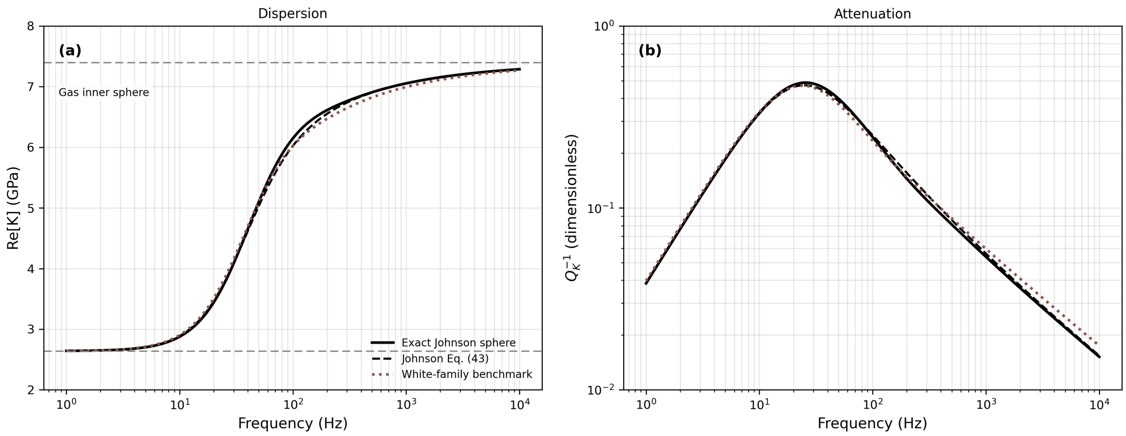

Figure 4.2: Bulk-modulus dispersion and attenuation are shown against frequency for the gas-inner concentric-sphere benchmark. (a) $\mathrm{Re}[K]$. (b) Q_K^{-1}. Curves compare the Johnson branching function, the exact sphere solution, and the White concentric-sphere model.

The three model families agree closely in the shared benchmark domain, confirming that the Johnson implementation is consistent with the independent White model in the concentric-sphere case.

4.3.3 Benchmark scope

The benchmark evidence supports three statements: (1) the exact reference solutions reproduce Johnson's published simple-geometry benchmark behaviour; (2) the branching-function model tracks those reference solutions closely; and (3) the White concentric-sphere model agrees with the Johnson model in the shared domain. The benchmark does not validate the branching function for arbitrary geometries beyond the tested cases. It confirms the implementation for periodic slabs and concentric spheres, which is sufficient to proceed to parameter scans in the concentric-sphere setting used below.

4.4 Idealised parameter scan

We perform parameter scans over saturation, permeability, and patch size using the Johnson branching-function model in the constant-frame concentric-sphere gas-inner geometry, in order to identify which parameter combinations push the relaxation transition into the seismic band. The scans use the Control case of Table 3.1 ($z = 1000$ m, $\phi = 0.15$, $S_b = 0.15$, $K_{\text{dry}} = 12.0$ GPa, $G_{\text{dry}} = 10.0$ GPa) and the same fluid-property workflow as the baseline. The mineral phase is taken as quartz-dominated sandstone with $K_0 = 36.6$ GPa and $\rho_m = 2650$ kg/m^3 throughout this thesis. At the Control-case reservoir conditions the Peng–Robinson equation of state gives, for the gas, $K_{H_2} = 14.4$ MPa and $\eta_{H_2} = 9.2$ μPa·s; Batzle–Wang gives, for the brine, $K_b = 3.11$ GPa and $\eta_b = 0.96$ mPa·s. The frequency axis spans 10^{-1}–10^4 Hz on a logarithmic grid (500 points).

The purpose of these scans is to map how permeability and patch size control whether the relaxation falls in the seismic band (1–100 Hz) and how large the resulting velocity dispersion and attenuation are.

4.4.1 Reference-case frequency scan

For the reference case, we fix permeability at $\kappa = 100$ mD and patch radius at $a = 5$ cm, and scan gas saturation from 0 to 50%. Figure 4.3 shows the resulting $V_P(f)$ and $Q^{-1}(f)$ curves.

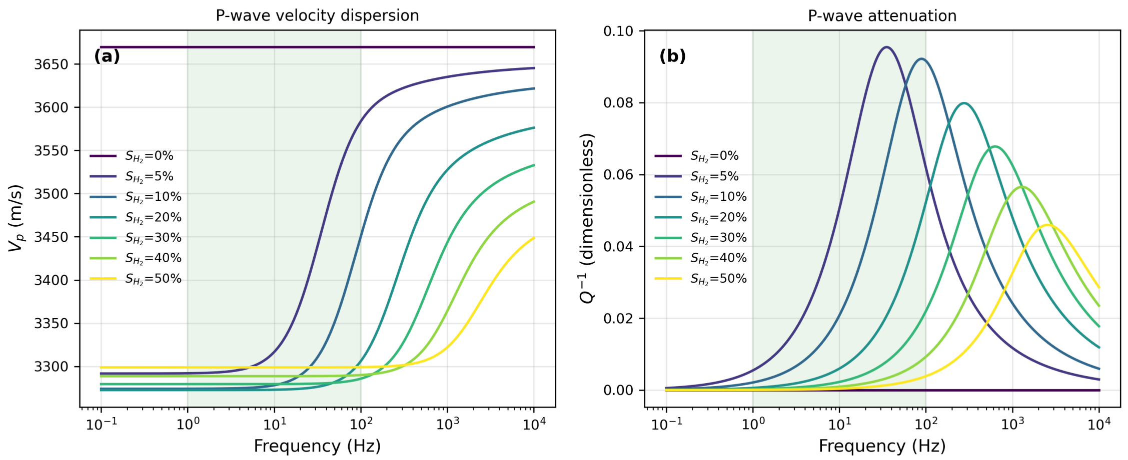

Figure 4.3: V_P and Q^{-1} are shown against frequency for hydrogen saturations from 0% to 50% at $\kappa = 100$ mD and $a = 5$ cm. (a) V_P. (b) Q^{-1}. The green band marks 1–100 Hz.

Figure 4.4 shows that the saturation dependence of V_s is mild and monotonically increasing, because reducing brine and adding hydrogen reduces the bulk density ρ_{bulk} while G_{sat} stays constant. The V_p/V_s ratio drops sharply at low gas saturation, mirroring the V_p behaviour discussed above; the small numerator change in V_p relative to the small denominator change in V_s amplifies into a larger ratio change. The

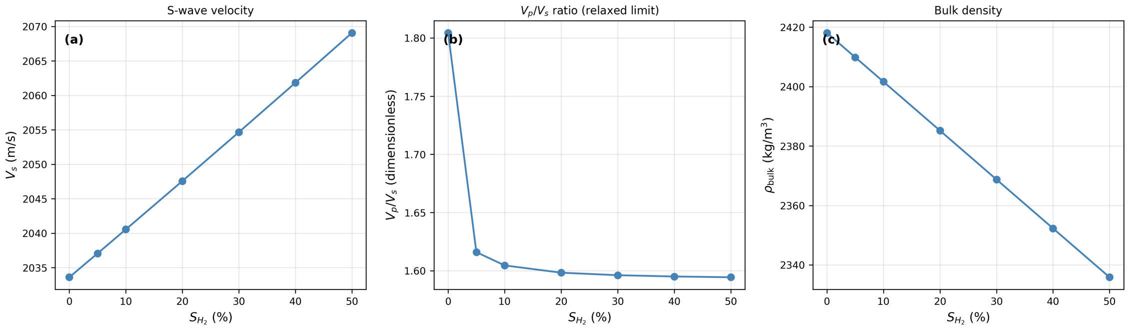

Figure 4.4: Shear-wave velocity V_s, the V_p/V_s ratio, and bulk density ρ_{bulk} are shown against hydrogen saturation at the relaxed limit for $\kappa = 100$ mD and $a = 5$ cm. (a) V_s. (b) V_p/V_s. (c) ρ_{bulk}.

bulk density falls roughly linearly with gas saturation, as expected from a simple volume-fraction argument.

The dispersion window widens and the attenuation peak shifts to higher frequencies as gas saturation increases. At low saturations ($S_{H_2} = 5\%$), the peak attenuation frequency is $f_{\mathrm{peak}} \approx 36$ Hz and falls squarely in the seismic band; the velocity dispersion between 30 and 100 Hz is $\sim$153 m/s. At $S_{H_2} = 20\%$, f_{peak} has moved to approximately 280 Hz, which is above the seismic band, and the 30–100 Hz velocity dispersion has dropped to $\sim$29 m/s. At high saturations ($S_{H_2} \geq 30\%$), the relaxation occurs above 600 Hz and the seismic-band velocity is nearly frequency-independent.

The maximum Q^{-1} ranges from $\sim$0.046 at $S_{H_2} = 50\%$ to $\sim$0.095 at $S_{H_2} = 5\%$. These are substantial attenuation values, but they occur at the respective peak frequencies; the attenuation at seismic frequencies is lower for cases where the peak has moved out of band.

4.4.2 Permeability scan

For the permeability scan, we fix $S_{H_2} = 20\%$ and $a = 5$ cm and vary κ across 1, 10, 100, 1000, and 10 000 mD. Figure 4.5 shows the result.

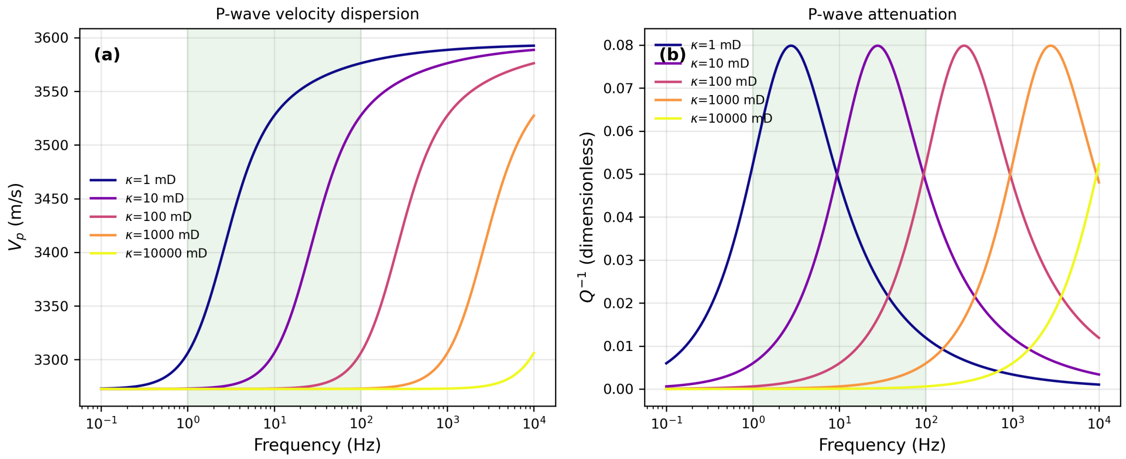

Figure 4.5: V_P and Q^{-1} are shown against frequency for permeability from 1 mD to 10 000 mD at $S_{H_2} = 20\%$ and $a = 5$ cm. (a) V_P. (b) Q^{-1}.

At this saturation, the relaxed-limit values are $V_s = 2048$ m/s, $V_p/V_s = 1.598$, and $\rho_{\mathrm{bulk}} = 2385$ kg/m^3, independent of permeability in this framework. The peak of the $\kappa = 10\,000$ mD curve lies above 10^4 Hz and is clipped by the frequency axis.

The peak attenuation frequency scales approximately linearly with permeability, shifting from $f_{\text{peak}} \approx 2.8$ Hz at $\kappa = 1$ mD to $f_{\text{peak}} \approx 2800$ Hz at $\kappa = 1000$ mD. The maximum Q^{-1} is $\sim$0.080 and is essentially independent of permeability (it depends on the saturation and the modulus contrast, not the diffusivity). The permeability controls where the attenuation peak sits, not how large it is.

For permeabilities in the range 1–100 mD, the relaxation frequency falls within or near the seismic band. The 30–100 Hz velocity dispersion is largest at $\kappa = 10$ mD ($\sim$107 m/s), where $f_{\text{peak}} \approx 29$ Hz sits directly in the seismic band. At higher permeabilities ($\kappa \geq 1000$ mD), the transition has moved above the seismic band, and the rock is effectively in the relaxed (Gassmann–Wood) regime at seismic frequencies.

The same permeability scan was evaluated at $S_{H_2} = 10\%$ and $S_{H_2} = 30\%$ in Section C.10; both variants follow the $S_{H_2} = 20\%$ pattern in Figure 4.5, with permeability setting the peak frequency while saturation controls the attenuation magnitude. At $\kappa = 100$ mD, f_{peak} rises from $\sim$90 Hz at 10% through $\sim$280 Hz at 20% to $\sim$640 Hz at 30%, while Q^{-1}_{max} falls from approximately 0.092 through 0.080 to 0.068. Consequently, some combinations without seismic-band sensitivity at 20% do show measurable sensitivity at 10%, where the peak is both stronger and lower in frequency at fixed permeability.

4.4.3 Patch-radius scan

The patch radius is the least directly constrained parameter in the model, so the values swept here are anchored to the ranges used in the patchy-saturation literature. He et al. (2020) swept patch radii of 0.5, 2, 5, and 10 cm for a gas-brine White–Dutta–Odé model in sandstone of comparable permeability to the Rotliegend case; larger patches lower the attenuation peak and shift it to lower frequencies, with $R_W = 10$ cm giving the strongest time-lapse response. Picotti and Carcione (2017) used a 50 cm outer-sphere radius for a CO_2 leakage scenario, and Gei et al. (2024), the only study consulted here that applies the Johnson model specifically to hydrogen, used a 30 cm radius at storage depth, although without an accompanying sensitivity study. Independent laboratory and stochastic-modelling work places the dominant heterogeneity scale at the centimetre level: centimetre-scale velocity heterogeneity in sandstone (Müller et al., 2010) and a von-Karman correlation length of order 10 cm in patchy gas-water distributions (Rubino & Holliger, 2012). The 5 to 20 cm range used here therefore brackets the values these studies adopt while remaining inside the mesoscopic envelope of Section 4.1. The reference value $a = 5$ cm used for the Control and Rotliegend-style cases sits at the centre of the He et al. (2020) scan; the larger $a = 15$ cm used for the Bunter-style case is chosen to place the relaxation peak in the lower seismic band, as discussed in Section 4.5.3.

The patch-radius scan fixes $S_{H_2} = 20\%$ and $\kappa = 100$ mD and varies the inner-sphere radius a from 0.5 cm to 50 cm. Figure 4.6 shows the dispersion and attenuation curves; Figure 4.7 shows f_{peak} as a function of patch radius.

The relaxed-limit values $V_s = 2048$ m/s, $V_p/V_s = 1.598$, and $\rho_{\text{bulk}} = 2385$ kg/m^3 are independent of patch radius and match the permeability scan in Figure 4.5.

The relaxation frequency scales approximately as $f_{\text{peak}} \propto a^{-2}$, as expected from the diffusive relaxation time $\tau \propto a^2/D$. Small patches ($a \leq 2$ cm) place the relaxation well above the seismic band; large patches ($a = 50$ cm) push it below 3 Hz. For this permeability, patch radii in the range $\sim$5–20 cm bracket the relevant mesoscopic range, with the in-band portion roughly 10–20 cm because $a = 5$ cm gives $f_{\text{peak}} \approx 280$ Hz

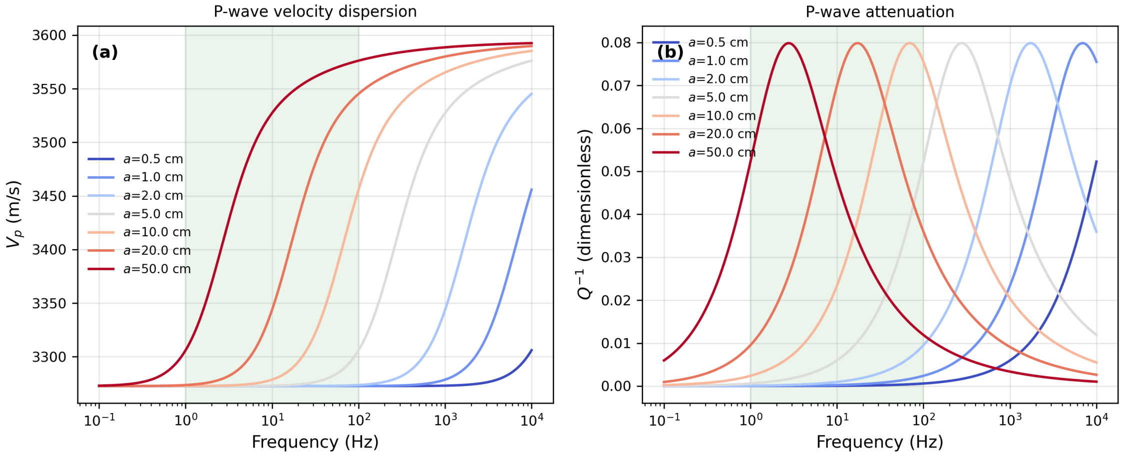

Figure 4.6: V_P and Q^{-1} are shown against frequency for patch radii from 0.5 cm to 50 cm at $S_{H_2} = 20\%$ and $\kappa = 100$ mD. (a) V_P. (b) Q^{-1}.

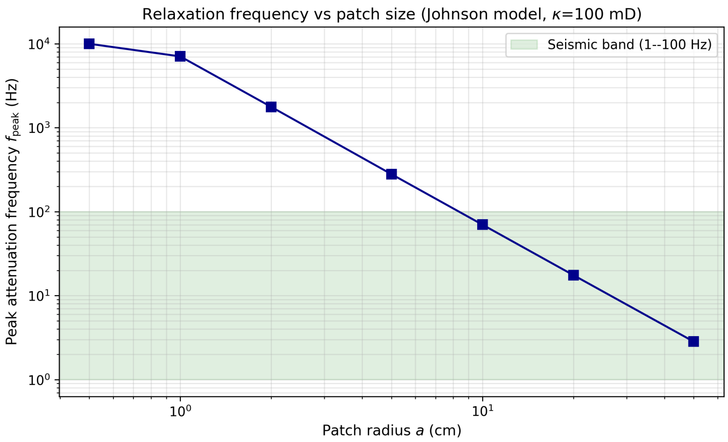

Figure 4.7: Peak attenuation frequency is shown against patch radius at $S_{H_2} = 20\%$ and $\kappa = 100$ mD. The green band marks 1–100 Hz.

above the band. At $a = 10$ cm, $f_{\text{peak}} \approx 70$ Hz and the 30–100 Hz velocity dispersion is ∼137 m/s. At $a = 20$ cm, $f_{\text{peak}} \approx 18$ Hz and the dispersion is ∼70 m/s.

The maximum Q^{-1} is again ∼0.080 and independent of patch size, confirming that the patch radius controls the frequency placement of the attenuation peak, not its amplitude.

4.4.4 Geometry sensitivity: sphere versus slab

The scans presented so far use the gas-inner concentric-sphere geometry throughout, and Figure 4.8 compares it with the periodic-slab geometry. That geometry is an idealisation: as set out in Section 4.1, the gas–brine alternation in a depleted-field storage column is controlled by bedding rather than by isolated spherical patches. The periodic-slab geometry of Johnson (2001) is the bedding-faithful counterpart, alternating gas-saturated and brine-saturated layers, and it is available at the same closed-form cost through the second exact (τ, ζ) construction. This subsection quantifies how much the choice between the two geometries changes the seismic-band response at fixed saturation and fixed characteristic scale.

For the slab the controlling length is the bedding period $L = L_1 + L_2$, with gas fraction $S_g = L_1/L$ and surface-to-volume ratio $S/V = 2/L$. For the sphere the controlling length is the gas-patch radius a, with the outer radius b fixed by S_g. To compare the two on a common footing we hold the rock and fluid inputs at the Bunter-style case of Table 3.1, fix $S_{H_2} = 0.20$, and scan the characteristic gas-region dimension across the same physical range used in the patch-radius scan of Section 4.4. The Bunter-style case is chosen because it is the only reference case whose relaxation frequency falls inside the seismic band (Section 4.5.3), so geometry choice has the largest effect on the imaged response there.

Both geometries are evaluated from the Johnson branching function (Eq. 2.17) with the (τ, ζ) parameters computed for each geometry, so that the comparison is consistent with the production workflow used in Chapter 5 rather than relying on the exact reference solutions. The exact periodic-slab and concentric-sphere solutions of Johnson (2001, Appendix B) are retained as a validation gate: at the matched scale they confirm that the branching curves track the exact solutions to within the intrinsic accuracy of the two-parameter interpolation. That accuracy is a few percent in $\text{Re}[K]$ and somewhat larger in Q^{-1} (sphere 3.2 % and 4.8 %; slab 4.5 % and 7.7 % at the Bunter point), which is the characteristic asymptotic residual of the branching function in a broad-crossover regime and not a transcription artefact. The branching function is used throughout because it is the geometry-general path that carries forward into the embedding and the scenarios; the exact solvers confirm it is faithful to within that stated residual.

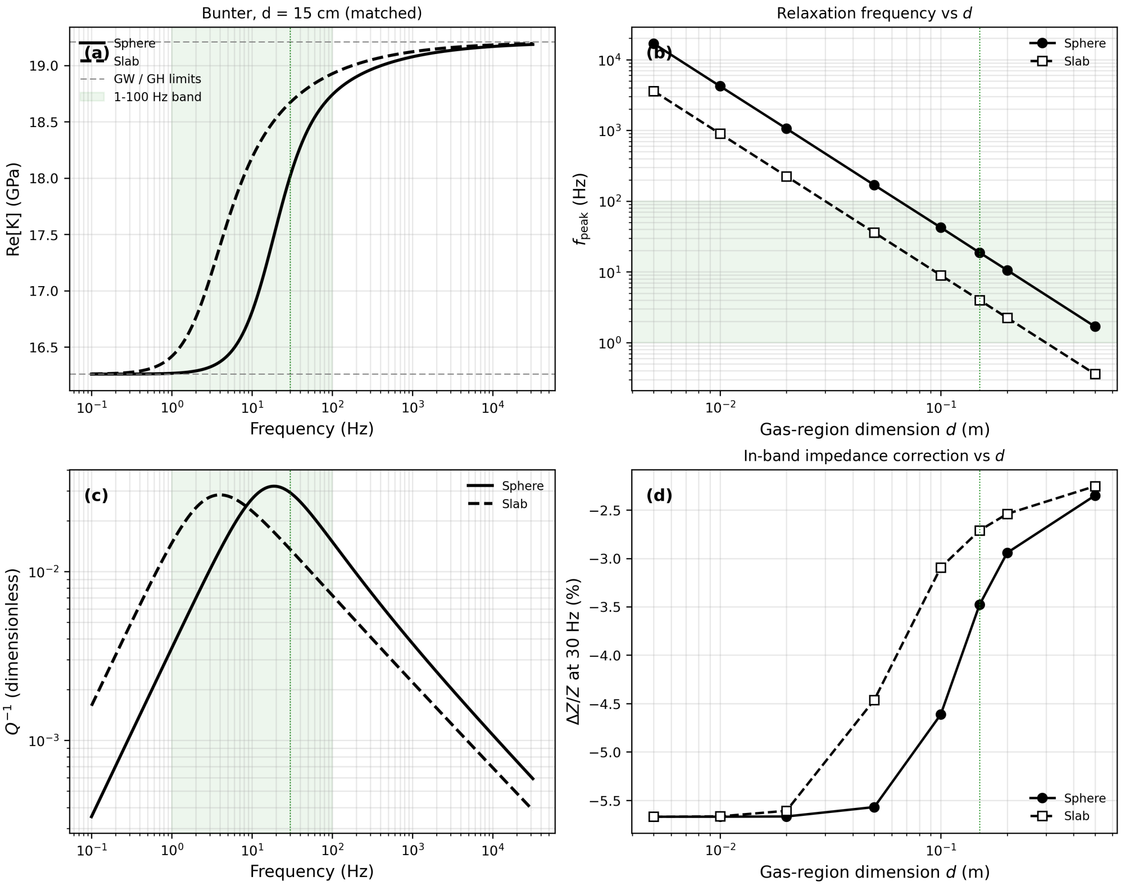

Figure 4.8: Bulk modulus, attenuation, relaxation frequency, and impedance correction compare sphere and slab geometries for the Bunter-style case at $S_{H_2} = 0.20$. (a) $\mathrm{Re}[K]$ against frequency. (b) f_{peak} against gas-region dimension. (c) Q^{-1} against frequency. (d) $\Delta Z/Z$ at 30 Hz against gas-region dimension.

At the matched characteristic scale $d = 0.15$ m and $S_{H_2} = 0.20$, switching from the concentric sphere to the periodic slab reduces the seismic-band $\Delta Z/Z$ correction at 30 Hz from $-3.5\,\%$ to $-2.7\,\%$, a difference of about 0.8 percentage points, or roughly $22\,\%$ in relative terms. The relaxation frequency differs markedly between the two geometries at the same d: $f_{\mathrm{peak}} \approx 19$ Hz for the sphere and ≈ 4 Hz for the slab. This difference follows from matching the gas-region linear dimension rather than the relaxation behaviour; at equal d the slab bedding period $L = d/S_g$ is several times the sphere outer diameter, giving the slab a longer characteristic diffusion length and a lower relaxation frequency. Because the slab peak lies below the seismic band while the sphere peak lies within it, at 30 Hz the slab has relaxed further toward the Gassmann–Hill limit and shows the smaller impedance contrast. The geometry idealisation is therefore a second-order but non-negligible modelling choice in the in-band Bunter regime: it does not alter the sign or order of magnitude of the monitoring signal, but it scales the in-band correction by about a fifth. The sphere is retained as the primary geometry for the embedding of Section 4.5, with the slab result bounding the geometry sensitivity. Per-cell geometry selection across a layered column is a later step on the modelling ladder and is not pursued here.

4.4.5 Summary of scan results

The three scans show that the frequency-dependent effect is controlled jointly by gas saturation, permeability, and patch size. Permeability and patch radius together set the diffusive relaxation time, and therefore the frequency at which the attenuation peak sits; gas saturation sets the modulus contrast between gas-rich and brine-rich regions, and therefore the magnitude of the velocity dispersion and attenuation. Certain range-combinations of all three parameters result in measurable frequency-dependent effects in the seismic band; outside those combinations, the rock behaves as the relaxed Gassmann–Wood limit predicts at seismic frequencies. The main findings of the three scans are:

- The peak attenuation magnitude ($Q^{-1}_{\max} \approx 0.05$–$0.10$) is set by the gas saturation and the modulus contrast between gas-rich and brine-rich regions; saturation therefore controls how large the dispersion and attenuation can become, while permeability and patch size do not.
- The peak attenuation frequency scales linearly with permeability and inversely with the square of the patch radius. For the tested parameter ranges, f_{peak} spans from $\sim$3 Hz to above 10^4 Hz.
- For κ in the range 1–100 mD and a in the range 5–20 cm, the scans bracket the relevant mesoscopic range; at $\kappa = 100$ mD, the in-band patch-radius portion is roughly 10–20 cm, while $a = 5$ cm gives $f_{\text{peak}} \approx 280$ Hz above the band.
- At higher permeabilities or smaller patch sizes, the rock is effectively in the relaxed Gassmann–Wood regime at seismic frequencies, and the frequency-dependent effect is negligible regardless of saturation.

Three caveats apply to the scan results when interpreted toward realistic monitoring scenarios. First, the patch radius a is treated here as a representative single length scale. In a real reservoir, the saturation distribution after H_2 injection has multiple length scales set by the depositional facies, capillary entry pressures, and gravity segregation; the dominant patch size is itself uncertain (Johnson, 2001; Müller et al., 2010). The 5–20 cm range used here brackets the mesoscopic patch sizes typically considered in the literature, but specific Dutch reservoir patch sizes are not identified or tested by the present modelling. Second, the shape parameter ζ in the Johnson branching function is derived from the patch surface-to-volume ratio and tortuosity, and is fixed by the chosen reference geometry (concentric sphere with gas-inner). A more general treatment would scan ζ explicitly to capture distributions of patch sizes and non-spherical patch geometries; this is identified as a planned extension of the present work. Third, and related, each (κ, a, S_g) combination in the scan produces a single relaxation peak at a single f_{peak}, because Eq. (2.17) carries one τ per evaluation. A reservoir with a distribution of patch sizes would superpose multiple such peaks and yield a broader, smoother $Q^{-1}(\omega)$ response than any single scan curve shown here (Pride et al., 2004; Müller et al., 2010). The sharp f_{peak} values reported below should therefore be read as the centre of mass of the relaxation distribution for a dominant patch size, not as the monochromatic response expected from a real reservoir.

A separate caveat applies to the practical seismic bandwidth at depth. The two evaluation frequencies used in Section 4.5 (30 Hz and 100 Hz) bracket typical surface-seismic monitoring bands, but high-frequency content (above $\sim$50 Hz) attenuates strongly in the overburden and is not realistically recoverable from deep reservoirs at 2–3 km depth in a real survey. The 100 Hz evaluation is therefore a methodological reference, not a prediction of practical detectability for the deeper cases. The 30 Hz evaluation is the more practically realistic one for the Rotliegend-style case at 2800 m. The shallower Bunter-style case at 1750 m could support a somewhat higher usable frequency, of order 40 Hz, so the 30 Hz evaluation is conservative rather than optimistic for that case.

A clarification on how the relaxation peak depends on saturation is warranted, because the dependence is stronger than the closed-form time scale alone suggests. The diffusive time scale that sets the nominal frequency placement, $\tau \propto a^2/D_0$ with $D_0 \approx \kappa K_b/\eta_b$ a brine-phase diffusivity, is independent of gas saturation. The peak attenuation frequency reported throughout this chapter, however, is not read from this τ. It is computed as the numerical maximum of $Q^{-1}(\omega)$ on the logarithmic frequency grid, and that maximum shifts with saturation because the Gassmann–Wood and Gassmann–Hill bounds that the branching function interpolates between are themselves saturation-dependent. This is the origin of the upward migration of f_{peak} with increasing S_{H_2} documented above. The same behaviour appears in the White concentric-sphere model, whose peak frequency is strongly saturation-dependent even though its geometric time scale is not (White, 1975; Dutta & Seriff, 1979). Any direct comparison between the Johnson branching function and the White model at matched conditions must therefore either calibrate the branching-function time scale to the White peak or state the saturation at which the comparison is made; the f_{peak} values reported here are the numerically located maxima of $Q^{-1}(\omega)$, not the analytic τ^{-1}.

Whether this frequency-dependent effect materially changes the detectability interpretation when embedded into the broader 1D monitoring workflow is the subject of the next section.

4.5 Embedding into the 1D monitoring workflow

The idealised scan of Section 4.4 mapped how gas saturation, permeability, and patch size control the relaxation frequency, but it did not yet answer the detectability question: does the frequency-dependent effect survive once placed back into the monitoring proxies inherited from the low-frequency baseline? This section evaluates this question by comparing the Johnson branching-function result at 30 Hz and 100 Hz against the low-frequency (Gassmann–Wood) baseline for the same three monitoring proxies used in Chapter 3: $\Delta Z/Z$, ΔR, and Δt_{2way}.

We use the three thesis reference cases of Table 3.1: the constant-frame Control case, the Rotliegend-style case, and the Bunter-style case. The Control case isolates the frequency-dependent mechanism from frame-model realism. The two reservoir-style cases inherit their porosity, pore pressure, depth, and dry-frame parameterisation from Section 3.1.1 and Section 2.5; the case-specific dry-frame moduli evaluated at each case's reservoir differential pressure are shown directly in Figure 3.5. The low-frequency baseline responses against which the embedding stage compares are shown case-by-case in Figure 3.2.

For the Rotliegend-style case, the MacBeth parameters of Table 3.1 (Rotliegend Facies B re-fit) at the case differential pressure $p_d \approx 33$ MPa give $K_{\text{dry}} \approx 18.7$ GPa and $G_{\text{dry}} \approx 12.5$ GPa. The permeability is set at $\kappa = 50$ mD, at the low edge of the HyTROS Rotliegend range (50–500 mD), to give the strongest candidate frequency-dependent response inside that envelope. The patch radius $a = 5$ cm matches the reference scan.

For the Bunter-style case, the MacBeth parameters of Table 3.1 (Berea Sandstone row) at the case differential pressure $p_d \approx 17$ MPa give $K_{\text{dry}} \approx 16.0$ GPa and $G_{\text{dry}} \approx 16.3$ GPa. The combination of low permeability ($\kappa = 50$ mD, at the low edge of the HyTROS Upper-Bunter range) and a moderately large patch radius ($a = 15$ cm) is chosen specifically to push the relaxation peak into the lower seismic band at $S_{H_2} = 0.20$, which is the regime in which any frequency-dependent correction is largest. Appendix B gives the Berea-over-Forties rationale.

For each case we compute all proxies relative to the $S_{H_2} = 0$ baseline at the same frequency: the low-frequency baseline uses the Gassmann–Wood saturated modulus, while the Johnson results use $\text{Re}[K(\omega)]$ from the branching-function model evaluated at 30 Hz and 100 Hz.

4.5.1 Control case

Figure 4.9 shows the monitoring proxies for the Control case. Because $f_{\text{peak}} \approx 280$ Hz at $S_{H_2} = 0.20$ lies well above the seismic band, the Johnson result at both evaluation frequencies stays close to the Gassmann–Wood baseline across the saturation range, with deviations below 10% (Table 4.1) and no sign changes at any saturation.

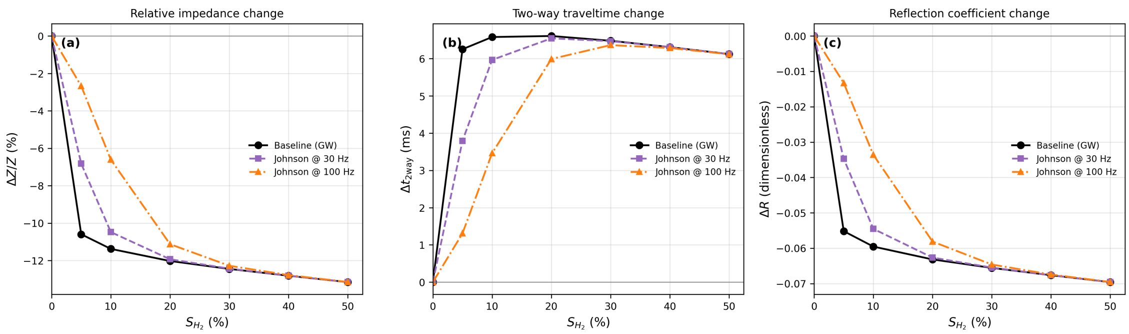

Figure 4.9: Relative impedance change, two-way traveltime change, and reflection-coefficient change are shown against hydrogen saturation for the Control case. (a) $\Delta Z/Z$. (b) Δt_{2way}. (c) ΔR. Curves show the low-frequency baseline and Johnson results at 30 and 100 Hz for $\kappa = 100$ mD and $a = 5$ cm.

4.5.2 Rotliegend-style case

Figure 4.10 shows the Rotliegend-style case. The MacBeth frame at $p_d \approx 33$ MPa is substantially stiffer than the Control frame, which compresses the Gassmann–Wood to Gassmann–Hill spread and reduces the absolute baseline response, as already visible in Figure 3.2. The relaxation peak again sits well above the seismic band ($f_{\text{peak}} \approx 286$ Hz at $S_{H_2} = 0.20$), so the frequency-dependent deviations remain below 10% at both evaluation frequencies (Table 4.1), with no sign changes: this case is firmly in the regime where the low-frequency baseline interpretation transfers without material correction.

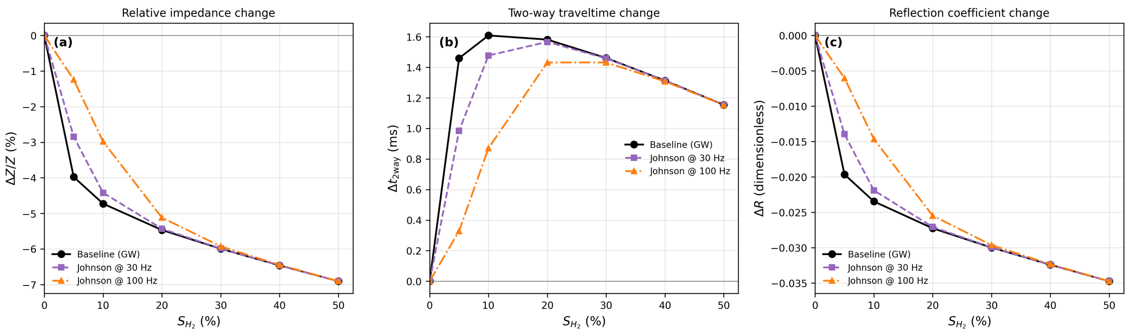

Figure 4.10: Relative impedance change, two-way traveltime change, and reflection-coefficient change are shown against hydrogen saturation for the Rotliegend-style case. (a) $\Delta Z/Z$. (b) $\Delta t_{2\text{way}}$. (c) ΔR. Curves show the low-frequency baseline and Johnson results at 30 and 100 Hz for the MacBeth frame, $\kappa = 50$ mD, and $a = 5$ cm.

4.5.3 Bunter-style case

Figure 4.11 shows the Bunter-style case. The MacBeth frame at $p_d = 17.6$ MPa is moderately stiffer than the Control frame, so the baseline response is again reduced relative to the Control. What distinguishes this case is the relaxation frequency: the combination of low permeability ($\kappa = 50$ mD) and a moderately large patch radius ($a = 15$ cm) places $f_{\text{peak}} \approx 19$ Hz in the lower seismic band at $S_{H_2} = 0.20$, so at the 30 and 100 Hz evaluation frequencies the rock has moved past its relaxation peak toward the unrelaxed Gassmann–Hill regime. The result is a large frequency-dependent reduction of all three proxies, from 39–62% at 30 Hz to 54–85% at 100 Hz (Table 4.1), with no sign changes at any saturation. This is the regime in which a frequency-blind monitoring workflow would most strongly overestimate the seismic signature of hydrogen injection. Because the absolute Gassmann–Wood baseline is itself moderate, the absolute proxy values under the Johnson model are smaller still, raising a separate detectability concern that is independent of the frequency-dependent question studied here.

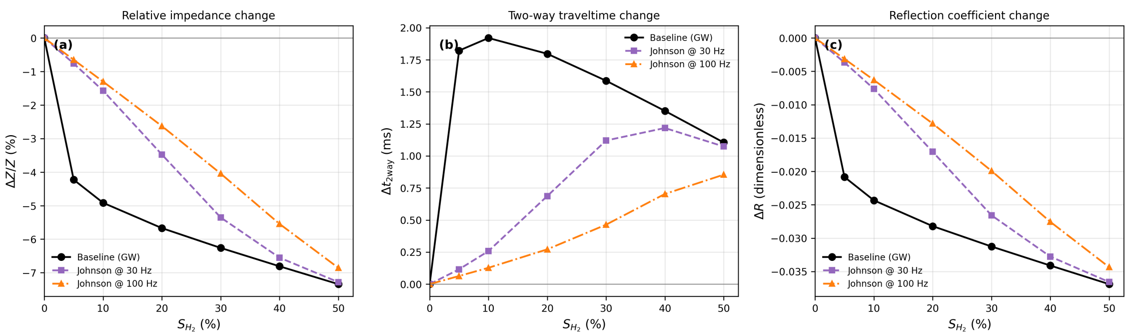

Figure 4.11: Relative impedance change, two-way traveltime change, and reflection-coefficient change are shown against hydrogen saturation for the Bunter-style case. (a) $\Delta Z/Z$. (b) $\Delta t_{2\text{way}}$. (c) ΔR. Curves show the low-frequency baseline and Johnson results at 30 and 100 Hz for the MacBeth frame, $\kappa = 50$ mD, and $a = 15$ cm.

4.5.4 Summary of embedding results

The pattern in Table 4.1 is clear: the two cases whose relaxation peak lies well above the seismic band show corrections below 10% at both evaluation frequencies, while

Table 4.1: Monitoring proxies at $S_{H_2} = 0.20$ for the three reservoir cases: low-frequency baseline (GW) and Johnson branching-function model at 30 Hz and 100 Hz. Deviations are quoted relative to the baseline magnitude. The Control case uses a constant frame; the Rotliegend-style and Bunter-style cases use the MacBeth-style frame of Appendix B.3.

Case	Proxy	Baseline	J @ 30 Hz	J @ 100 Hz	dev. 30 Hz	dev. 100 Hz
Control	$\Delta Z/Z$ (%)	−12.03	−11.94	−11.13	0.7%	7.5%
	Δt_{2way} (ms)	6.61	6.55	5.99	1.0%	9.4%
	ΔR	−0.0632	−0.0627	−0.0581	0.8%	8.1%
Rotliegend	$\Delta Z/Z$ (%)	−5.47	−5.43	−5.12	0.7%	6.4%
	Δt_{2way} (ms)	1.58	1.57	1.43	1.0%	9.5%
	ΔR	−0.0273	−0.0271	−0.0255	0.7%	6.6%
Bunter	$\Delta Z/Z$ (%)	−5.67	−3.48	−2.63	39%	54%
	Δt_{2way} (ms)	1.80	0.69	0.27	62%	85%
	ΔR	−0.0282	−0.0171	−0.0128	39%	55%

the Bunter-style case, with its peak inside the band, is reduced by 39–62% at 30 Hz and 54–85% at 100 Hz across the three proxies.

Two consequences follow. First, the frequency-dependent effect acts in the same direction in all three cases: it reduces proxy magnitudes relative to the baseline, because the rock is partially stiffened toward the Gassmann–Hill limit at seismic frequencies. No proxy sign reversals occur at any saturation level. The frequency-dependent effect therefore does not create false negatives where the baseline predicts a detectable response; it does, however, reduce the margin available for detection in cases where the relaxation falls in the seismic band. Second, frame stiffness controls the absolute baseline magnitude, not the size of the frequency-dependent correction relative to that baseline: the stiff MacBeth frames in both reservoir cases compress the Gassmann–Wood to Gassmann–Hill spread, but the relative correction is still controlled by where f_{peak} falls relative to the seismic band, which in turn depends on permeability, patch radius, and (through the diffusivity) porosity and viscosity.

4.6 Interim conclusion

The frequency-dependent patchy-saturation effect, as modelled by the Johnson branching-function model in the concentric-sphere gas-inner geometry, does not change the first-order detectability interpretation: the signs of $\Delta Z/Z$, ΔR, and Δt_{2way} match the low-frequency baseline in every case, at every saturation. The quantitative correction, however, is case dependent. Where the reservoir parameters place the mesoscopic relaxation inside the seismic band, as the combination of low permeability and a moderately large patch radius does in the Bunter-style case, the proxies are reduced by tens of percent (Table 4.1), and a workflow that predicts magnitudes from the Gassmann–Wood closure and compares them to field observations at 30–100 Hz would overestimate the expected response. Where the relaxation sits well above the band, as in the Control and Rotliegend-style cases, the baseline is an adequate description at seismic frequencies and the correction stays below 10%.

This conclusion is methodologically aligned with Gei et al. (2024), who characterised the frequency-dependent acoustic properties of a soft, high-porosity, high-permeability Utsira-like sand with the same modelling stack. The present results extend that exploration to the stiffer, lower-permeability conditions representative of Dutch reservoirs and ask the monitoring consequence rather than the rock-physics characterisation; Chapter 6 develops this relation and the comparison with the wider literature.

The branching-function structure supports the next rung of the modelling ladder directly: each layer of a column can carry its own (ϕ, κ, S_g, a) tuple, and therefore its own (τ, ζ) and $K(\omega)$, at the same closed-form cost. A 1.5D column with an overburden, a caprock interface, and a layered reservoir saturation is therefore the natural next step, taken in Chapter 5, with the single-relaxation-spectrum and single-geometry caveats of Section 4.4.5 carried forward.

Chapter 5

Scenario-Based Seismic Detectability

The previous chapter established the frequency-dependent patchy-saturation response of a hydrogen-bearing reservoir and reduced it to a set of monitoring proxies at a single reservoir interface: the fractional acoustic impedance change $\Delta Z/Z$, the two-way travel-time shift $\Delta t_{2\text{way}}$, and the reflection-coefficient change ΔR. These proxies are interface quantities. They assume that whatever happens at the top of the reservoir is what a survey records. A monitoring survey does not record an interface number; it records seismic reflection data, from which a migrated image is produced. Between the proxy of Chapter 4 and the recording surface sits a stratigraphic column rather than a single interface, finite-bandwidth wave propagation, and a migration operator. This chapter is the final rung of the modelling ladder: it embeds the frequency-dependent response in a migrated section under Dutch reservoir conditions and asks whether the effect survives the synthetic and imaging workflow or is diluted by it.

5.1 From interface proxy to imaged section

The monitoring proxies described in Chapter 4 are evaluated at one interface and carry no information about the rest of the column. Translating them into an imaged reflector requires three steps, each of which can in principle attenuate or distort the signal. First, the single interface is replaced by a depth column: an overburden, a caprock, a reservoir carrying a prescribed hydrogen distribution, and an underburden. Second, a finite-bandwidth wavefield is propagated through that column, so the reservoir response is convolved with a source wavelet and combined with reflections from every other contrast in the section. Third, the prestack response is migrated back to depth, so the recovered reflector amplitude depends on the migration velocity model as well as on the true contrast.

The geometry used here is laterally invariant. The medium is a one-dimensional depth column; the Kennett reflectivity method (Kennett, 1983) turns it into a two-dimensional prestack shot record in offset and time; and a slant-stack phase-shift loop images that record back to depth. This is a one-and-a-half-dimensional treatment, not a true two-dimensional one. Lateral homogeneity is an assumption, made so that the reservoir-fluid response can be isolated from structural and lateral-heterogeneity effects, and heterogeneous two-dimensional modelling is left to future work. This boundary is stated here because it conditions every result in the chapter: the migrated reflector amplitudes report what survives propagation and imaging through a layered but laterally uniform medium, and nothing about lateral variation.

A second assumption, the single real scalar velocity per layer that makes each synthetic a frequency-domain proxy evaluated band by band rather than a time-domain waveform, follows from the propagator and is stated with the workflow in Section 5.4.

When the reservoir is given a single uniform hydrogen saturation, the migrated top-reservoir reflector reproduces the interface proxies of Chapter 4, recovering $\Delta Z/Z$ to within 0.05 percentage points, Δt_{2way} to within 0.05 ms, and ΔR to within 5×10^{-4} at the reference frequency. This agreement validates the synthetic-to-migration chain end to end before any non-uniform topology or layered overburden is introduced, and it is the regression reference against which the rest of the chapter is built.

5.2 The laterally invariant layer-cake model

The reservoir is embedded in a flat-layered, laterally invariant column, referred to here as the layer-cake model. From top to bottom the column consists of an overburden terminating in a sealing caprock, the reservoir interval carrying a prescribed hydrogen distribution, and an underburden that provides a base-reservoir contrast. The reservoir frame is one of the two Dutch reservoir analogues carried through the thesis, the Rotliegend-style and Bunter-style cases, whose porosities, pressures, permeabilities, dry-frame parameters, and patch radii are listed in Table 3.1 and are not restated here. The Control case is retained as a synthetic anchor with a constant dry frame. A construction schematic for the layered treatment of this column is provided as Figure F.1 in Appendix F.1.

Each reservoir is embedded at its true depth: the reservoir top is placed at the case depth of Table 3.1, so the Bunter-style reservoir sits at 1750 m, the Rotliegend-style reservoir at 2800 m, and the Control anchor at 1000 m. The full overburden column spans from the surface to that depth. This is a change from the embedding used in the initial development of the workflow, in which the reservoir was placed beneath a fixed reference overburden of fixed thickness regardless of the case depth; that earlier choice decoupled the reservoir elastic properties, which are evaluated at the true case depth, from the seismic geometry, which was shallow. Embedding at true depth removes that inconsistency, so the migrated section now corresponds to a reservoir at its actual Dutch depth, with the corresponding overburden thickness, travel time, and reflector geometry.

In this chapter the overburden is treated in two ways. The first is a single homogeneous block spanning the full column to the reservoir top, which keeps the reservoir-fluid response cleanly separated from overburden interference while the saturation scenarios are introduced. The second is a realistic layered overburden, built from the Dutch stratigraphy, added as a controlled robustness step in Section 5.5.4. Both treatments reach the same true reservoir depth and terminate in an identical caprock immediately above the reservoir.

5.3 Saturation scenarios

Hydrogen injected into a depleted gas reservoir does not distribute itself uniformly. Buoyancy, capillarity, and reservoir layering all shape where the gas accumulates, and the resulting saturation topology controls both the average reservoir saturation seen at the top reflector and the patch-scale geometry that drives the frequency-dependent response. Five idealised topologies were constructed to span this range. Two are carried in the body of this chapter; the remaining three are documented in Appendix F.

The two main scenarios are chosen to bracket the realistic cases. The *uniform* scenario assigns a single hydrogen saturation across the whole reservoir. It is the cleanest possible topology and serves as the regression reference that anchors the synthetic-to-migration chain to Chapter 4. The *gravity-segregated* scenario places a hydrogen-rich zone at the top of the reservoir with saturation decreasing downward, representing the buoyant gas cap expected when injected hydrogen rises against the caprock. This configuration is the one most consistent with the geological controls on hydrogen storage in depleted Dutch sandstone fields identified by Loyola et al. (2026) and with the storage behaviour described in the study of Bagheri et al. (2026), which makes it the most thesis-relevant of the realistic topologies.

The three additional scenarios extend the range without adding to the body argument. A *capillary transition* topology imposes a smooth downward saturation decay representing a capillary-controlled transition zone; a *baffled* topology introduces a low-permeability interval that discontinuously separates storage zones; and a *Rotliegend-style continuous* topology represents a laterally continuous reservoir body with moderate vertical variation. Each is described, with its migrated response, in Appendix F.

Across all scenarios the hydrogen cap is given a soft target saturation band of 0.15 to 0.60, consistent with the saturations reached in cyclic storage operation; brine-dominated and baseline intervals below this band are permitted. This band is a design guide rather than a constraint enforced in the modelling: each scenario run reports the per-layer minimum and maximum hydrogen saturation as a visibility diagnostic, so the realised saturations can be read against the intended band without the topology being clamped.

5.4 Synthetic and imaging workflow

Each scenario is converted into the property logs that drive the synthetic data modeling. The layer-cake column is sampled at a fixed depth increment into compressional velocity, shear velocity, and bulk density profiles. The reservoir layers carry the frame moduli, density, and hydrogen distribution of the scenario; the overburden, caprock, and underburden carry their assigned elastic properties. A zero-phase Ricker wavelet at the chosen centre frequency is used as the source. The prestack response is computed with the Kennett reflectivity solver (Kennett, 1983), and the resulting shot record is migrated to depth with a slant-stack phase-shift-plus-interpolation loop that uses a smoothed background velocity. The migrated top-reservoir reflector amplitude is the quantity carried into the results.

Figure 5.1 shows this chain at its midpoint and end for the Bunter-style uniform scenario at 30 Hz: the prestack record that the reflectivity solver returns, and the migrated depth section from which the top-reservoir amplitude is read.

Three acoustic states are propagated for each scenario, differing only in how the reservoir compressional velocity is set. The *all-brine* state fills the reservoir with brine through Gassmann substitution and represents the pre-injection baseline. The *relaxed* state imposes the scenario hydrogen distribution and computes the reservoir velocity from the Gassmann–Wood low-frequency limit, which carries the saturation change but no dispersion. The *tuned* state imposes the same hydrogen distribution and takes the reservoir velocity from the real part of the Johnson complex velocity

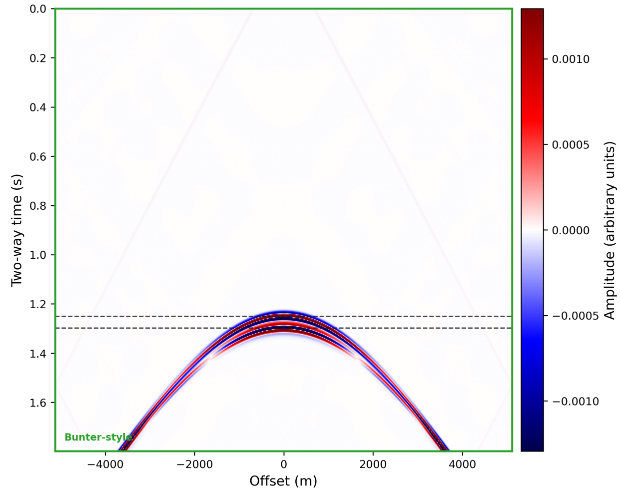

(a)

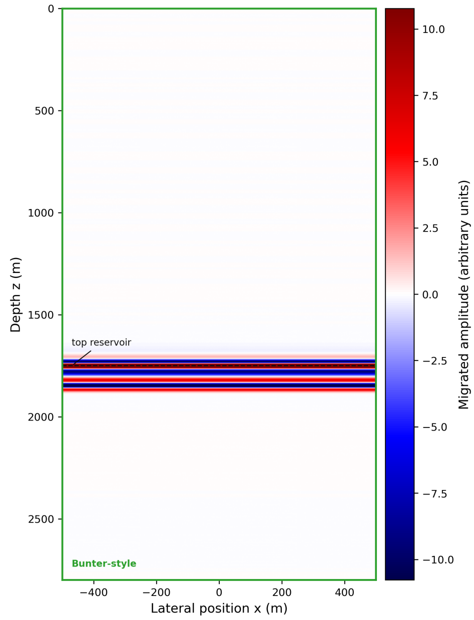

(b)

Figure 5.1: In the migrated section the lateral axis carries no physical information: the medium is laterally invariant, so a single migrated trace is repeated across a nominal lateral extent for display. Prestack shot record and migrated depth section for the Bunter-style uniform scenario under the single-block overburden, tuned state, at a 30 Hz Ricker centre frequency. (a) Common-source gather in offset and two-way time. (b) Migrated depth section; the top-reservoir reflector at 1750 m is the surface on which the amplitude differences of this chapter are read.

evaluated at the Ricker centre frequency, which carries the frequency-dependent patchy response. Two differences between these states are carried through the chapter. The difference between the all-brine and tuned states is the monitoring signal (TmB): the full seismic change produced by replacing reservoir brine with patchy hydrogen. The difference between the relaxed and tuned states is the frequency-dependent survival (TmR): the part of the monitoring signal that exists only because the patchy response is dispersive rather than relaxed. Both differences are read on the migrated top-reservoir reflector.

What each monitoring proxy integrates over. The three monitoring proxies carried through this chapter do not sample the reservoir in the same way, and the distinction matters once the saturation profile is no longer uniform. The impedance change $\Delta Z/Z$ and the reflection-coefficient change ΔR are properties of the top-reservoir interface, and are evaluated from the uppermost reservoir sublayer alone. The two-way travel-time shift $\Delta t_{2\mathrm{way}}$ is not an interface quantity at all: it is a thickness-weighted sum of the slowness change over every reservoir sublayer, and therefore integrates the whole saturation profile. For the uniform topology the distinction is empty, since the reservoir is a single layer. For the gravity-segregated and capillary topologies it is real, and it is the reason the impedance proxies and the travel-time proxy can move in opposite directions: the former respond only to the gas-rich cap, while the latter also carries the more weakly saturated interval beneath it. Where this chapter refers to the three quantities collectively as monitoring proxies, that grouping is inherited from the homogeneous single-layer reservoir of Chapter 4 and should be read with this difference in mind.

The migrated synthetics are generated at an applied frequency set of 15, 30, 50, and 100 Hz, which spans the seismic monitoring band and resolves the frequency dependence of the response. The cheaper interface-proxy evaluations and the in-panel display markers use a set of 10, 30, 50, and 100 Hz.

Scope of the modelled physics. The single-real-scalar-velocity assumption noted in Section 5.1 is a property of the reflectivity method used here, which accepts one real velocity per layer and cannot be modified to carry a complex velocity or an intrinsic quality factor through propagation. The tuned state therefore injects the dispersive response through the real part of the Johnson velocity at the imaging frequency: it captures the velocity dispersion that shifts the reflector, but not the intrinsic attenuation that the same model predicts during propagation. Modelling intrinsic attenuation in propagation with a viscoacoustic or viscoelastic treatment admitting a complex, frequency-dependent modulus would require a different propagator and is outside the scope of this chapter.

Discretisation of layered scenarios. The continuous saturation profiles of the gravity-segregated and capillary scenarios are represented by a finite number of constant-saturation layers. The sensitivity of the migrated reflector to this layer count was tested for the capillary transition scenario, which has the smoothest profile and is therefore the most demanding case; the test and its convergence are reported in Appendix F.

5.5 Results

The results are reported as the two migrated top-reservoir reflector amplitude differences defined in Section 5.4: the monitoring signal TmB and its frequency-dependent part TmR. Figure 5.2 summarises both across scenarios, cases, and applied frequencies.

What the migrated amplitudes are. The amplitude carried through this chapter requires a word of definition, because it is neither a reflection coefficient nor a calibrated physical quantity. The reflectivity method returns dimensionless reflection and transmission responses, and the source wavelet is applied at unit peak amplitude, so a single ray-parameter trace of the migrated image is in reflection-coefficient units. We verified this directly: at the top-reservoir reflector the single-ray-parameter migrated amplitude reproduces the analytic normal-incidence reflection coefficient to three or four significant figures for every case tested. The quantity reported here, however, is the full stack, formed as an unnormalised sum over the ray-parameter axis. That sum multiplies the reflectivity-scaled trace by an effective count of coherently summing traces.

That multiplier is not a constant. It depends on the case, the frequency, and the overburden treatment, varying by a factor of about eight across the results reported below (Table F.1), because the evanescent limit removes different portions of the ray-parameter aperture at different frequencies and depths. Absolute stack amplitudes are therefore not comparable across frequencies or across cases, and no single scaling converts them to reflectivity.

For this reason every amplitude difference in this chapter is reported both as the raw migrated value and as a percentage of the all-brine reference amplitude at the same case, frequency, and overburden treatment, measured at the same depth sample. The scaling is common to the three coupling states, so it cancels exactly from that ratio, and the percentages are directly comparable across the frequency axis in a way the raw values are not. The percentage should be read as the fractional change in the imaged reflector relative to its pre-injection amplitude. Where this chapter compares magnitudes across frequency or between cases, it is the percentages that carry the comparison.

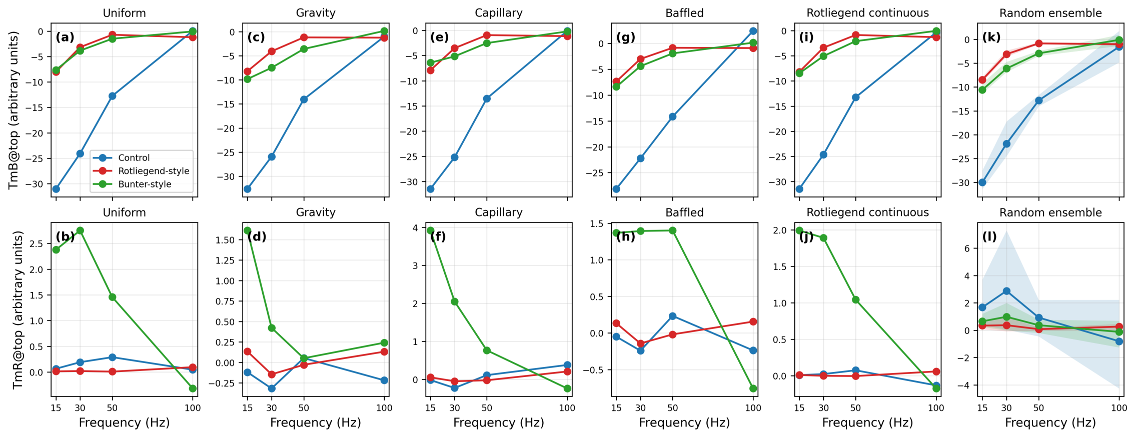

Figure 5.2: Migrated top-reservoir reflector amplitude differences are shown against applied frequency across six saturation scenarios under the single-block overburden. (a), (c), (e), (g), (i), and (k) show TmB (tuned minus all-brine); (b), (d), (f), (h), (j), and (l) show TmR (tuned minus relaxed). The panel pairs are uniform, gravity-segregated, capillary, baffled, Rotliegend continuous, and random ensemble, respectively. Colours identify the three cases; shaded bands show the random-ensemble spread.

5.5.1 Uniform reference scenario

The full migrated-amplitude values for the uniform scenario are given in Table F.2 in Appendix F.1. The migrated images behind these values, for both overburden treatments, are reproduced in Figures F.4 and F.6 of the same appendix. The monitoring signal is large for every case: replacing reservoir brine with patchy hydrogen produces a strong impedance contrast regardless of frequency, and this is a Gassmann-level fluid-substitution effect that does not by itself test the thesis question. The monitoring signal is largest for the shallow Control anchor and smallest for the deep Rotliegend-style case, consistent with the reduced reflector amplitude expected at greater depth. The frequency-dependent part tells a different and case-dependent story. For the Control and Rotliegend-style cases it is a small fraction of the monitoring signal, remaining at or below a few percent across the band. For the Bunter-style case it is of the same order as the entire monitoring signal.

Figure 5.3 shows the frequency-dependent part of the monitoring signal as a fraction of the reference reflector amplitude, which is the form in which magnitudes are comparable across the frequency axis. The separation between the cases is stark. Under the uniform topology the Bunter-style correction reaches 8.1, 10.3 and 7.4 percent of the reference amplitude at 15, 30 and 50 Hz, while the Control and Rotliegend-style corrections remain at or below 1.5 percent everywhere in the band. Under the gravity-segregated topology the Bunter-style correction falls to 5.5 percent at 15 Hz and to below 2 percent above it, because the impedance response is generated at the gas-rich cap while the reservoir beneath it is only weakly saturated. The ordering between cases is unchanged by the topology; only the magnitude moves.

Expressing both differences as percentages of the all-brine reference amplitude makes the size of the frequency-dependent correction directly readable. In the Bunter-style case the monitoring signal at 30 Hz is -14.3 percent of the reference reflector amplitude while the dispersive part is $+10.3$ percent, and the two carry opposite signs: the dispersive correction opposes the bulk fluid-substitution response rather than adding to it. The relaxed state alone would therefore predict a monitoring signal

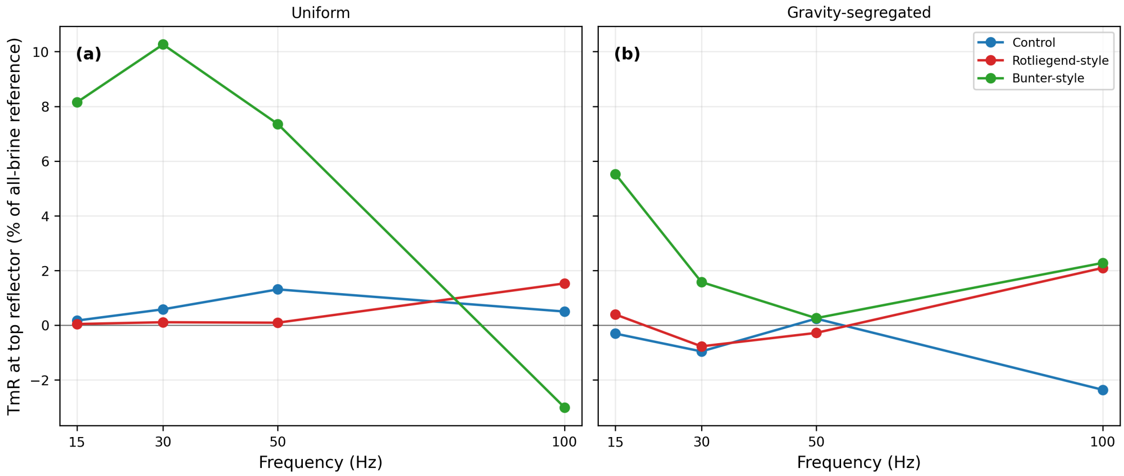

Figure 5.3: The frequency-dependent part of the migrated monitoring signal, tuned minus relaxed at the top reservoir reflector, is shown against applied frequency as a percentage of the all-brine reference amplitude at the same sample. (a) Uniform topology. (b) Gravity-segregated topology. Both under the single-block overburden. Colours identify the three cases.

of −24.6 percent, so a low-frequency description overpredicts the imaged change by a factor of 1.72 at that frequency. The same comparison gives factors of 1.31 at 15 Hz and 2.00 at 50 Hz for the Bunter-style case, against 1.00 to 1.02 for the Control and Rotliegend-style cases across the same band. The frequency-dependent physics therefore changes the predicted monitoring signal by one to two percent where the relaxation peak lies above the seismic band, and by a factor of two where it lies inside it. This overprediction is insensitive to the overburden treatment: the same comparison under the layered Dutch overburden gives factors of 1.32, 1.71 and 1.79 at the same three frequencies. It is, however, specific to the uniform topology; under the gravity-segregated topology the same factors fall to between 1.01 and 1.16, because the impedance proxies respond only to the gas-rich cap while the dispersive correction is generated throughout the reservoir.

At the top of the applied band the comparison becomes unstable rather than merely small. At 100 Hz the monitoring signal has decayed to a fraction of a percent of the reference amplitude in several case and scenario combinations, and in four of the forty-eight combinations examined the relaxed and tuned states predict monitoring amplitudes of opposite sign. Three of those four are the Bunter-style case, which is where the dispersive correction is largest. A sign reversal of this kind is a genuine consequence of dispersion rather than a numerical artefact, but it occurs where the underlying signal is too weak to support a quantitative statement, and we therefore restrict the overprediction comparison above to 15 to 50 Hz.

This case separation follows directly from Chapter 4: the in-band correction is controlled by where the relaxation peak falls relative to the seismic band, and only the larger-patch Bunter-style case places that peak inside it, so the whole band samples the dispersive part of the response there. Where the monitoring signal itself approaches zero, at the highest frequency for the shallower cases, the ratio $|TmR/TmB|$ inflates without the dispersive correction strengthening, and should be read against the small absolute TmR value rather than as a large effect. Figure 5.5 shows the three migrated reflectors for the Bunter-style case, where the separation between the relaxed and tuned traces is the dispersive part.

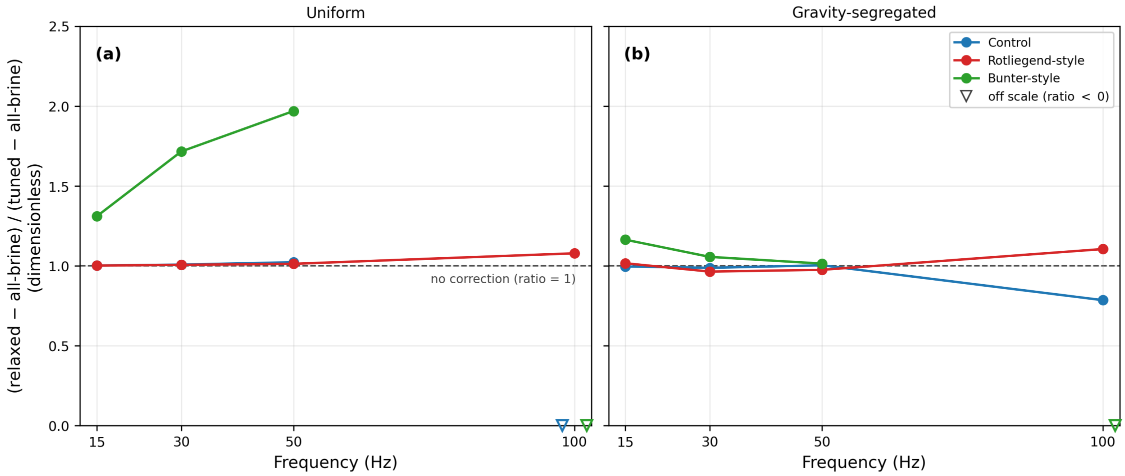

Figure 5.4: The ratio of the monitoring signal predicted by the relaxed state to that predicted by the tuned state is shown against applied frequency, giving the factor by which a low-frequency description overpredicts the imaged change. (a) Uniform topology. (b) Gravity-segregated topology. Both under the single-block overburden. The dashed line at unity marks agreement between the two descriptions. Points marked at the lower axis lie off scale because the ratio is negative there, indicating that the two descriptions predict monitoring amplitudes of opposite sign.

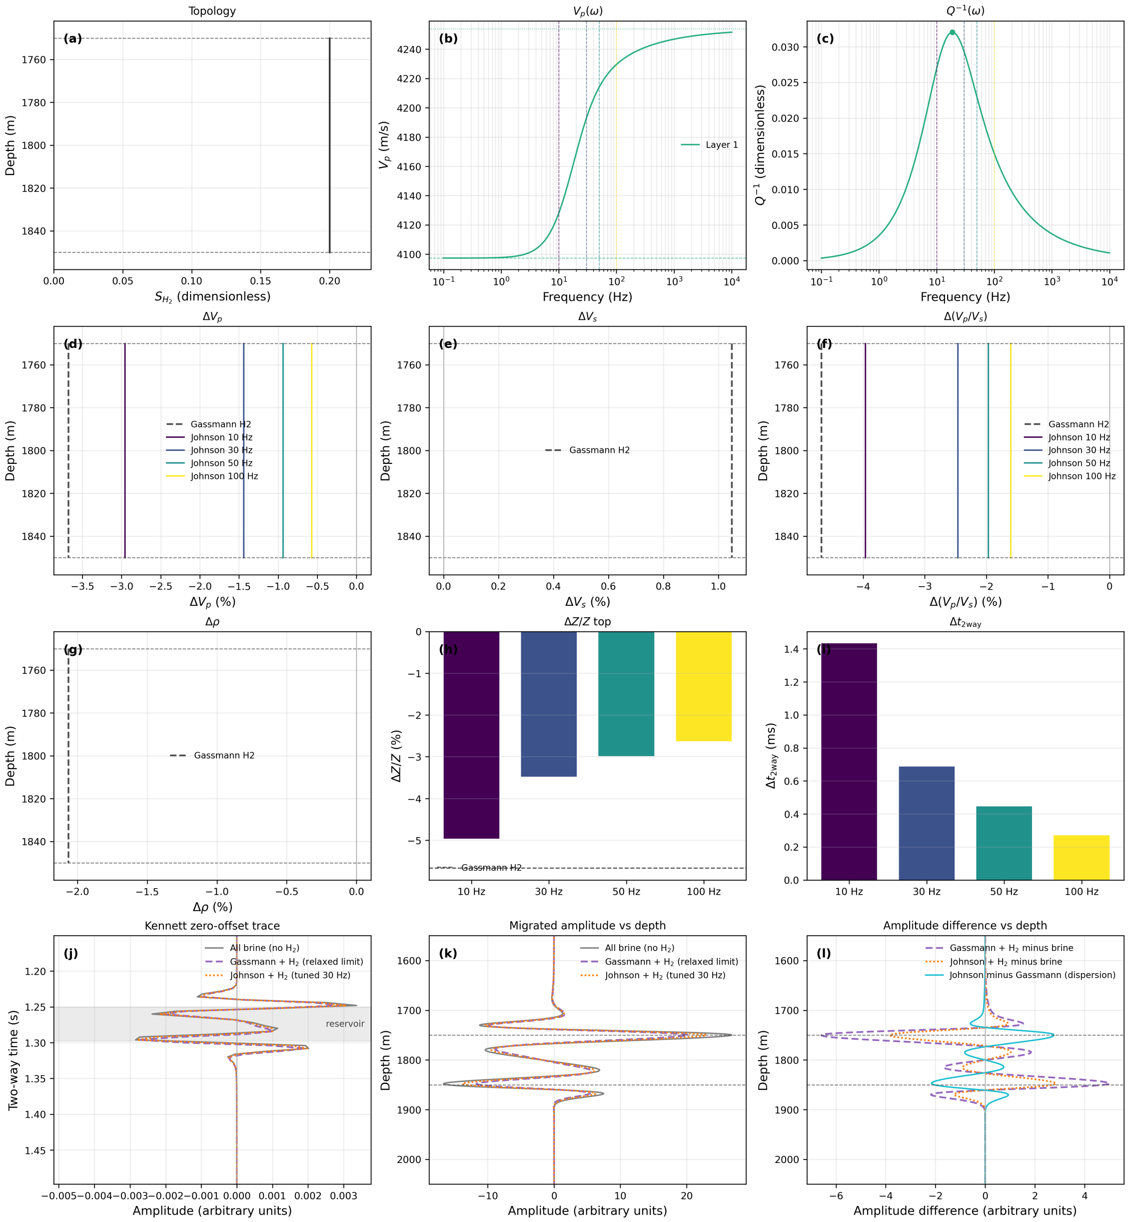

Figure 5.5: Hydrogen saturation, elastic-property changes, monitoring proxies, and synthetic seismic amplitudes are shown for the Bunter-style uniform scenario. (a) Saturation topology; (b) $V_p(\omega)$; (c) $Q^{-1}(\omega)$. (d) ΔV_p; (e) ΔV_s; (f) $\Delta(V_p/V_s)$; (g) $\Delta\rho$. (h) Top-reservoir $\Delta Z/Z$; (i) $\Delta t_{2\text{way}}$. (j) Zero-offset trace; (k) migrated amplitude; (l) migrated amplitude difference.

5.5.2 Gravity-segregated gas cap

The full migrated-amplitude values for the gravity-segregated scenario are given in Table F.4 in Appendix F.1. In this topology, the hydrogen-rich zone sits at the top of the reservoir against the caprock. The qualitative pattern of the uniform scenario is preserved but weaker: a large monitoring signal for every case, a frequency-dependent part of a few percent for Control and Rotliegend-style, and a frequency-dependent part for Bunter-style that is reduced relative to the uniform case, reaching roughly a sixth of the monitoring signal at the low end of the band and falling to a few percent by mid-band. The magnitudes shift because the top-reservoir reflector now reads the high-saturation cap rather than a uniform interior, but the case ordering is unchanged: Bunter-style still carries the largest dispersive correction and the other two cases remain at the percent level. The saturation topology modulates the size of the effect; it does not change which case carries it.

The full two-way travel-time shifts for the uniform and gravity-segregated topologies are given in Table F.3 in Appendix F.1.

5.5.3 Patch-radius sensitivity

The case-dependence above attributes the survival of the dispersive correction to the placement of the peak relaxation frequency, which is set primarily by the patch radius. To confirm that this is the operative control rather than a coincidence of the chosen case parameters, the patch radius was swept for the gravity-segregated scenario while the other reservoir properties were held fixed. Figure 5.6 shows the result. As the patch radius increases, the peak relaxation frequency moves downward into the seismic band, and the frequency-dependent part of the migrated signal grows accordingly; small patch radii push the peak above the band and the correction collapses. This reproduces, in the migrated section, the mechanism identified at the interface level in Chapter 4, and it confirms that the Bunter-style survival is a consequence of its larger patch radius placing the relaxation peak in band, not of its frame or its specific reservoir conditions.

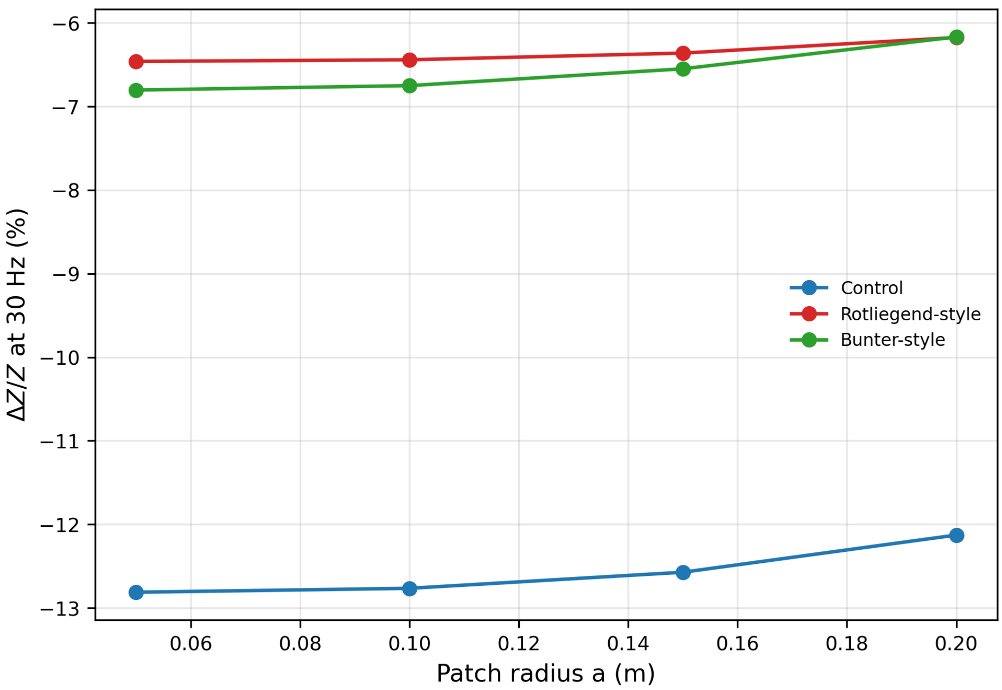

Figure 5.6: The frequency-dependent part of the migrated top-reservoir signal is shown against patch radius for the gravity-segregated scenario.

5.5.4 A layered overburden: does the surviving signal persist?

The scenarios to this point place the reservoir beneath a single homogeneous overburden block. This isolates the reservoir-fluid response but understates the overburden, which in the Dutch setting is layered and seismically active. The HyTROS field review (Reijnen-Mooij et al., 2025) distinguishes two overburden styles that map onto the two reservoir cases. Above the Rotliegend reservoirs of Noord-Holland sits a Zechstein salt seal, a high-velocity, low-density unit that forms strong impedance contrasts at its boundaries. Above the Bunter reservoirs of Zuid-Holland the seal is shale-dominated and layered: a Röt claystone, then several hundred metres of Jurassic Altena shale, with a Cretaceous chalk and a thick Cenozoic clastic section higher in the column. A single block captures neither character.

To test whether the frequency-dependent signal established above survives a realistic overburden, each primary scenario is rerun beneath a layered cake built from these two stratigraphies, spanning the full column to the true reservoir depth, and the result is compared against the single-block case. The comparison is controlled by construction. The layer immediately above the reservoir is held identical between the two treatments, so the top-reservoir reflection coefficient is unchanged; the realistic units, the Zechstein salt for the Rotliegend-style case and the Altena shale, chalk, and clastics for the Bunter-style case, are stacked above that matched layer. Any change in the migrated reflector therefore comes from overburden interference and from the more structured migration-velocity background, not from a change in the reflector itself. In the migrated sections the strong contrasts, the Zechstein salt boundaries for the Rotliegend-style case and the reservoir reflector for all cases, image clearly; the weaker clastic contrasts are gentler and are partly smoothed by the migration background velocity, as expected.

The same three coupling states are used, and the overburden is identical across all three. It is therefore a common factor, but it is a common *factor* rather than a common offset, and this distinction governs how the layered treatment should be read. To a good approximation the migrated reflector amplitude is the two-way transmission response of the overburden multiplying the reservoir reflection response. A common multiplicative factor divides out of a ratio but does not cancel from a difference: it scales it. The relaxed-to-tuned difference is therefore not independent of the overburden, even though the reservoir reflection coefficient that generates it is. What is strictly overburden-independent is the interface quantity ΔR itself, and the travel-time shift, which is computed over the reservoir interval alone and never encounters the overburden.

This matters because the transmission response is frequency dependent, and a single homogeneous block places its transmission nulls at different frequencies than a realistic layered column does. Where the simplified overburden happens to place a null, the dispersive difference at that frequency is suppressed by the overburden rather than by the reservoir physics. Comparing the two treatments therefore tests something stronger than robustness: it separates the part of the frequency dependence that belongs to the reservoir from the part contributed by the propagation path. The medium remains laterally invariant and the propagation uses one real scalar velocity per layer, so internal and surface multiples are not modelled in this pass. The step tests transmission through the layered section and the effect of a structured migration background; multiple contamination is a more stringent test left to future work, alongside the salt-directly-on-reservoir geometry that a fully faithful Rotliegend column would impose.

The full frequency-dependent values under the two treatments are tabulated side by side in Table F.5 in Appendix F.1, and the outcome separates sharply by case. For the Bunter-style case, the case in which the dispersive correction survives migration, the layered values track the single-block values closely: the same sign, the same order of magnitude, and the same shape across the band, with the largest deviation well under twenty percent. For the Control and Rotliegend-style cases the values differ by much larger factors, reaching an order of magnitude at Control 100 Hz and a factor of nearly four at Rotliegend-style 50 Hz. Those are not competing measurements of a robust quantity. They are two measurements of a signal small enough that the transmission path, rather than the reservoir, dominates its value, and they confirm that the dispersive correction is not recoverable in those cases under either treatment. Figure 5.7 shows the same comparison graphically.

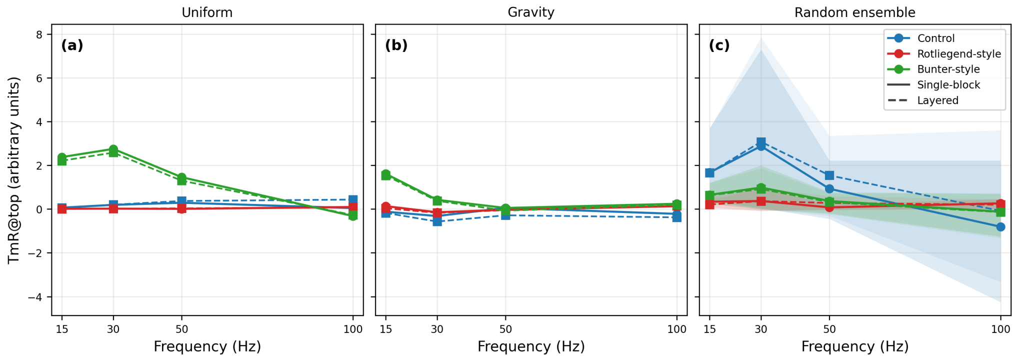

Figure 5.7: The frequency-dependent migrated top-reservoir amplitude difference TmR (tuned minus relaxed) is shown against applied frequency for two overburden treatments. (a) Uniform scenario. (b) Gravity-segregated scenario. (c) Random ensemble. Solid curves show the single-block overburden, dashed curves the layered Dutch overburden, and colours the three reservoir cases.

Two conclusions follow, and they should be kept apart.

The first is empirical and concerns the case that carries the thesis result. For the Bunter-style case the dispersive correction is essentially unchanged by the realistic overburden. This is not guaranteed by the construction of the comparison, for the reason given above, but it is what the modelling returns, and it holds across the whole applied band. The case in which the frequency-dependent signature is recoverable is therefore also the case in which that signature is insensitive to how the overburden is represented, which is the more useful of the two properties for a monitoring feasibility assessment.

The second concerns the two cases in which the correction is small. There, the choice of overburden changes the value by factors of up to nine. The direction of that change is not systematic: at some frequencies the layered treatment yields a larger dispersive difference than the single block, at others a smaller one. This is the signature of transmission nulls moving in frequency as the overburden structure changes, and it means that any statement about the magnitude of the dispersive signal in the Control and Rotliegend-style cases is a statement about the propagation path as much as about the reservoir. The conclusion that survives is the ordering, not the values: under both treatments those cases remain one to two orders of magnitude below the Bunter-style case, and neither treatment brings them within reach.

The absolute reflector amplitude is modulated by the realistic overburden in both cases, most strongly where a high-impedance unit such as the Zechstein salt is present, so the full monitoring signal must always be read against the overburden in place. The travel-time shift is unaffected for a different and stronger reason: it is computed over the reservoir interval alone, so the overburden never enters the sum, and its invariance is true by construction rather than a result.

5.5.5 Recoverability under acquisition noise

The results above establish how large the frequency-dependent correction is. Whether it can be recovered is a separate question, and it depends on the noise level of the survey that would have to measure it. We add band-limited noise to the migrated sections at prescribed signal-to-noise ratios, apply an independent noise realisation to each of the two states being differenced, since a time-lapse difference does not cancel non-repeatable noise, and ask at what level the relaxed-to-tuned difference ceases to be distinguishable from zero across repeated realisations. The signal-to-noise ratio is defined as the root-mean-square amplitude of the all-brine migrated section over the reservoir interval, divided by the standard deviation of the added noise.

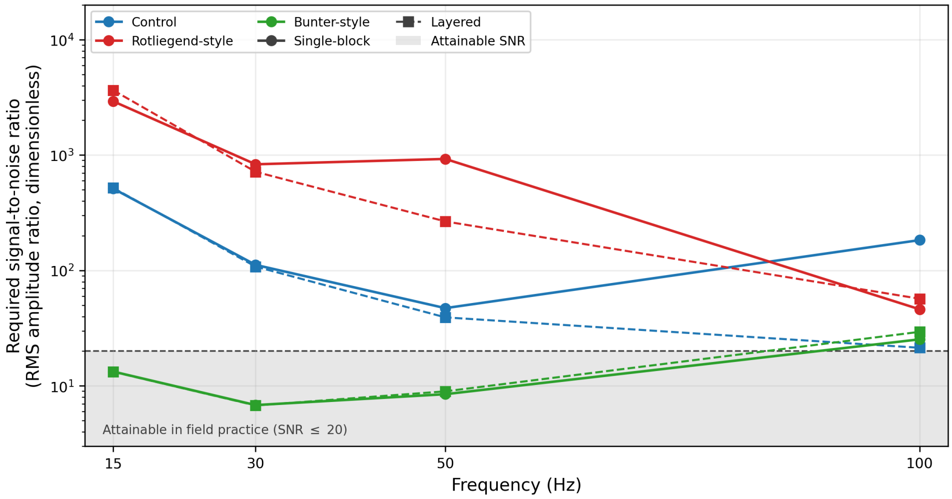

Figure 5.8: The signal-to-noise ratio required for the frequency-dependent part of the monitoring signal to be distinguishable from zero is shown against applied frequency, on a logarithmic scale, for the uniform topology. Solid curves give the single-block overburden and dashed curves the layered Dutch overburden. Colours identify the three cases. The shaded region marks the range attainable in field practice.

Figure 5.8 separates the cases by two orders of magnitude. The Bunter-style dispersive signature requires signal-to-noise ratios of roughly 7 to 13 between 15 and 50 Hz, which is within reach of practical acquisition, and the requirement is almost identical under the two overburden treatments. The Control and Rotliegend-style signatures require ratios from about 46 to nearly 3000, which is not. The conclusion drawn from the noise-free amplitudes therefore survives the addition of noise, and in a stronger form: the dispersive correction is not merely diluted in two of the three cases, it is unrecoverable in them under any realistic survey. With ten noise realisations the observed spread understates the tails of the distribution, so this distinguishability test is mildly optimistic; the Control and Rotliegend-style verdicts fail even under the optimistic test, and the marginal case is the Bunter-style recovery at the top of the band.

The more informative result is what noise does not affect. The full monitoring signal remains recoverable at a signal-to-noise ratio of five in the very cases where the dispersive part is unrecoverable at twenty. The fragility belongs to the frequency-dependent refinement and not to the detectability of hydrogen injection itself, which is the sharpest available statement of the two-level separation that runs through these results.

5.5.6 Ensemble random patchiness

The five scenarios to this point impose deterministic, idealised saturation topologies. Each fixes the patch radius at the case value and prescribes the saturation profile by a single rule. Injected hydrogen does neither: the patch scale and the local saturation both vary from place to place within the reservoir, in ways that are not known in advance. This scenario tests whether the case-dependent survival established above is a property of the rock-physics mechanism or an artefact of the particular idealised topologies, by replacing the prescribed topology with an ensemble of random realisations.

In each realisation the reservoir is divided into ten equal-thickness layers, and both the patch radius and the hydrogen saturation are drawn at random and independently for every layer: the patch radius log-uniformly over the mesoscopic range used in the sweep of Chapter 4, and the saturation uniformly over the soft target band. Ten independent realisations are drawn per case, each with its own seed, and the migrated top-reservoir reflector amplitude differences are computed for every realisation. The full ensemble summary is given in Table F.6 in Appendix F.1. For each case and applied frequency, it reports the across-realisation mean of the monitoring (TmB) and frequency-dependent (TmR) differences, with the range across realisations reporting the spread. With ten realisations per case, the reported tenth and ninetieth percentiles approximate the observed extremes rather than percentiles of an underlying distribution, and are read here as the spread of the ensemble rather than as a confidence interval. The ensemble is run under both overburden treatments so that the overburden robustness of the deterministic scenarios can be checked against randomised patchiness as well.

Under the single-block overburden the ensemble is complete, and its result qualifies the deterministic picture rather than simply reproducing it. The random per-layer mix of patch radius and saturation partially smears the in-band dispersive signal of the Bunter-style case. The ensemble-mean TmR sits well below the deterministic uniform value at the frequencies where the deterministic effect is largest (0.65 and 0.98 at 15 and 30 Hz, against deterministic uniform values of 2.38 and 2.75), and the 10-to-90 band reaches down toward zero. The coherent dispersive correction therefore survives in the mean but is markedly weaker and more uncertain than any single idealised topology suggests. The monitoring signal TmB, in contrast, is larger in magnitude than in the deterministic uniform case, consistent with the higher mean saturation produced by independent per-layer draws from the target band.

Layered overburden. The ensemble has since been completed under the layered Dutch overburden as well, over the same ten seeds per case, and it confirms the single-block result rather than qualifying it further. The layered-overburden Bunter-style ensemble mean TmR is 0.63 at 15 Hz and 0.92 at 30 Hz, against 0.65 and 0.98 under the single block, and the percentile bands are likewise close. The same smearing of the deterministic value is therefore present under both treatments, and the ensemble reproduces under randomised patchiness the pattern already established deterministically: the Bunter-style dispersive signal is insensitive to how the overburden is represented, even though the diagnostic itself is not overburden-independent in general.

5.6 Chapter synthesis

Three statements summarise what the scenarios show. First, replacing reservoir brine with patchy hydrogen produces a large monitoring signal for every case and every topology. This is a fluid-substitution effect at the Gassmann level; it confirms that hydrogen injection is seismically visible, but it is not the novel question of this thesis. Second, the frequency-dependent part of that signal, the part that exists only because the patchy response is dispersive, survives migration where the rock physics places the peak relaxation frequency inside the seismic band, and is diluted to a few percent where that peak lies above the band. Among the Dutch cases tested this means the dispersive correction is recovered in the larger-patch Bunter-style case, where it reaches the same order as the entire monitoring signal, and is effectively lost in the Control and Rotliegend-style cases. Third, the saturation topology modulates the magnitude of the dispersive correction, as the difference between the uniform and gravity-segregated scenarios shows, but it does not change which case carries the effect; the control is the patch-radius-set relaxation frequency, confirmed by the patch-radius sensitivity, and inherited directly from Chapter 4.

This case-dependent survival is also stable under the realistic overburden, within the limits established in Section 5.5.4: because the overburden scales rather than cancels the dispersive difference, the case ordering is robust under both treatments, while the magnitude of the correction must be quoted against a realistic overburden.

The conditional nature of this result, surviving where the relaxation peak is in band and diluted where it is not, is the substantive outcome of the chapter. What it implies for the detectability of UHS in practice, how it sits alongside the rock-physics characterisations in the literature, and what the dispersion-only propagation closure and the laterally invariant geometry leave unaddressed are drawn together in Chapter 6.

Chapter 6

Discussion

The preceding chapters answered our four research questions along a modelling ladder. The low-frequency reference model of Chapter 3 established the reference response and the base scenario (Stekelenburg, 2026). The frequency-dependent modelling of Chapter 4 isolated the patchy-saturation effect at the reservoir interface and mapped its dependence on permeability and patch size. The scenario modelling of Chapter 5 embedded that effect in a migrated section under Dutch reservoir conditions and the modelling asked whether it survives the synthetic and imaging workflow. Here, we interpret those three stages together, place the outcome against the literature, and state the boundaries within which it holds. The conclusions themselves are drawn in Chapter 7.

6.1 Interpretation of the detectability result

What controls the frequency-dependent effect. Patchy hydrogen saturation produces velocity dispersion and attenuation through mesoscopic wave-induced fluid flow (White, 1975; Johnson, 2001; Müller et al., 2010). The magnitude of the in-band correction is not controlled by frame stiffness. It is controlled by where the peak relaxation frequency falls relative to the seismic band, which is set primarily by the patch radius and the permeability. Among our three Dutch cases, only the larger-patch Bunter-style combination places that peak inside the seismic band. The Control and Rotliegend-style combinations place it above the band, so across 15 to 100 Hz they sample essentially the relaxed limit. This is the mechanism established at the interface in Chapter 4 and confirmed in the migrated section by the patch-radius sensitivity of Chapter 5.

Where the model hierarchy holds. In the concentric-sphere gas-inner domain shared by the three models, the Johnson branching-function model reproduces the exact reference solution and tracks the White concentric-sphere benchmark (White, 1975; Johnson, 2001). This establishes that the scalable branching-function path, which parametrises the transition between the relaxed and unrelaxed limits through analytically derived relaxation-time and shape parameters without solving a geometry-specific boundary-value problem for each case, is an adequate stand-in for the exact solution within the tested geometry. It is therefore the model we carry into the parameter scans and the migrated scenarios. The agreement is a benchmark consistency within a shared domain, not a transfer of validation across geometries. The relaxation-time and shape parameters differ by geometry, so an extension beyond the concentric-sphere case would require its own validation rather than inheriting this one.

Two levels of detectability. Our answer separates into two levels, and keeping them apart is the central result of this work. At the first level, the full monitoring signal (the change produced by replacing reservoir brine with patchy hydrogen) is large for every case and every saturation topology. This is a fluid-substitution effect at the Gassmann level. It confirms that hydrogen injection is seismically visible in all three Dutch settings, but it is not the novel element of this work, because it is already predicted by the low-frequency reference model. At the second level, the frequency-dependent part of that signal (the part that exists only because the patchy response is dispersive rather than relaxed) survives migration only where the rock physics places the relaxation peak inside the seismic band. For the Bunter-style case the dispersive correction reaches ten percent of the reference reflector amplitude, against a monitoring signal of fourteen percent, and it opposes rather than reinforces that signal: a low-frequency description overpredicts the imaged change by a factor of between 1.3 and 2.0 across the band. For the Control and Rotliegend-style cases, the same comparison gives factors between 1.00 and 1.02: the dispersive part is diluted to a few percent of the monitoring signal and is, for practical purposes, lost in the migrated image (Tables F.2 and F.4). Our answer to the detectability question is therefore conditional rather than absolute. Frequency-dependent patchy saturation changes the quantitative monitoring interpretation where the relaxation peak is in band, and is otherwise diluted to negligibility by the imaging workflow. The saturation topology modulates the size of the correction, as the difference between the uniform and gravity-segregated scenarios shows, but it does not change which case carries the effect.

Dependence on the overburden. The case-dependent recoverability established above is unchanged under a realistic overburden, but the reason is empirical rather than structural, and the distinction matters. As established in Section 5.5.4, the overburden enters the migrated reflector amplitude multiplicatively, as a two-way transmission response scaling the reservoir reflection, so it scales the relaxed-to-tuned difference rather than cancelling it; only the interface reflection coefficient itself and the reservoir-interval travel-time shift are strictly overburden-independent.

What the modelling returns is a separation by case (Table F.5). For the Bunter-style case, in which the dispersive correction is large, replacing the single-block overburden with a realistic layered Dutch column at true reservoir depth leaves that correction essentially unchanged across the applied band. For the Control and Rotliegend-style cases, in which the dispersive correction is small, the two treatments differ by factors of up to nine, and not systematically in one direction: where the difference is small it is the propagation path rather than the reservoir that controls the value.

Two consequences follow. The magnitude of the dispersive correction cannot be quoted from a simplified overburden: at some frequencies a simplified column happens to transmit very little of the wavefield to the reservoir and back, so the difference is suppressed by the propagation path rather than by the reservoir physics; the realistic overburden is therefore a correction and not merely a robustness check. The case ordering on which our answer rests, however, is stable under both treatments, and it is stable for the case in which the correction is recoverable at all. That is the weaker of the two claims we might have made, and it is the one the modelling supports.

Survival under acquisition noise. The prescribed signal-to-noise-ratio study of Section 5.5.5 sharpens this separation into a statement about recoverability: the dispersive signature is recoverable at signal-to-noise ratios attainable in field practice only for the Bunter-style case, while for the other two cases it would require ratios one to two orders of magnitude beyond acquisition practice. The full monitoring signal, by contrast, remains recoverable at low signal-to-noise ratios in every case. The fragility is therefore specific to the dispersive signature and not to the detectability of hydrogen itself, which is the sharpest form of the two-level separation that runs through our results.

Qualification from the random-patchiness ensemble. The random-patchiness ensemble of Chapter 5, drawing an independent patch radius and saturation for every reservoir layer across repeated realisations, tests whether the case-dependent recoverability is a property of the mechanism or of the particular idealised topologies. It qualifies rather than overturns our interpretation. In the Bunter-style case, the dispersive correction survives in the ensemble mean at the low end of the band, but at roughly a third of the deterministic uniform value, with a spread across realisations that reaches zero by 30 Hz (Table F.6): realistic per-layer heterogeneity smears the in-band signal without erasing it in the mean. The monitoring signal itself grows relative to the uniform case, because the randomised draw has a higher mean saturation. The two overburden treatments agree closely throughout, which extends the case-specific overburden insensitivity from the deterministic topologies to randomised ones, the more demanding test of the two.

6.2 Relation of the migrated results to the interface-level modelling of Chapter 4

The interim comparison drawn at the end of Chapter 4 already placed our interface-level response against the published rock-physics literature. Two points made there carry directly into the migrated result and are worth restating in that light.

First, the interface-level comparison with Gei et al. (2024) showed agreement in the shape of the dispersion curve but a systematic offset in its magnitude, attributable to the stiffer frame and lower permeability of our Dutch cases relative to their Utsira-like sand. The migrated results of Chapter 5 show that this offset is not a cosmetic difference. Because the peak relaxation frequency scales with permeability and inversely with the square of the patch radius, the same offset that shifts the curve at the interface is what moves the peak out of the seismic band for two of our three cases, and it is therefore the direct cause of the case-dependent survival we observe after imaging.

Second, the interface-level agreement with the permeability dependence described by Müller et al. (2010) was noted there as a consistency check. After migration, it becomes a predictive statement: because the relaxation peak position, and not the frame, decides survival, permeability and patch radius are the two reservoir properties that a monitoring feasibility study would need to constrain before predicting whether the dispersive correction is recoverable in a given Dutch field.

6.3 Relation to the wider literature

Complementary scope relative to the closest methodological reference. The closest methodological reference point is Gei et al. (2024), and the relationship is complementary rather than validating. The two studies are methodologically aligned: both use Peng–Robinson gas properties, Gassmann substitution, and the Johnson branching-function model with analytically derived relaxation-time and shape parameters applied to a concentric-sphere gas-inner geometry, and they occupy complementary regions of the parameter space. Gei et al. (2024) characterised the rock-physics response of an Utsira-like North Sea sand (a soft, high-porosity, high-permeability setting in which mesoscopic attenuation is pronounced). We move to stiffer, lower-permeability conditions representative of Dutch depleted gas fields, replace the Krief–Athy frame with a MacBeth-style calibration to reservoir-rock measurements (MacBeth, 2004), and ask a different question: not what the acoustic properties look like, but whether the effect changes the monitoring interpretation once it is imaged. Gei et al. (2024) established the rock-physics characterisation; we test the monitoring consequence. The two are complementary precisely because they probe different rocks and ask different but related questions.

Consistency with the mesoscopic-flow literature. The dependence we find is consistent with the broader understanding of mesoscopic wave-induced fluid flow. Müller et al. (2010) noted that the mechanism is permeability-dependent and most relevant when the characteristic diffusion length matches the seismic wavelength. Our results quantify that dependence, and carry it through migration, for hydrogen and brine patchy saturation in Dutch-style reservoir conditions. Our choice of the gravity-segregated topology as the most relevant realistic case is grounded in the Dutch storage context described by Loyola et al. (2026) and in the study of Bagheri et al. (2026), and the overburden stratigraphies used in the layered treatment follow the HyTROS field review (Reijnen-Mooij et al., 2025).

Comparison with field time shifts from CO_2 monitoring. No hydrogen storage field dataset exists yet against which our modelled time shifts can be checked, so the nearest available comparison is with CO_2-storage monitoring in depleted North Sea gas fields, a setting that shares the depth range, the reservoir lithologies, and the acquisition practice of the Dutch fields considered here. Toh et al. (2025) modelled reservoir-interval time shifts of 0.2 to 3.1 ms across thirty-seven UK Continental Shelf depleted gas fields and 0.1 to 2.6 ms across twenty-seven Norwegian Continental Shelf fields, and assessed them against a minimum detectability threshold of 0.5 ms for towed-streamer acquisition. They further reported that time shifts proved the more reliable observable of the two, since the corresponding amplitude changes rarely reached the normalised root-mean-square threshold required to distinguish them from background noise. Their UK cases were grouped by reservoir and included a Bunter Sandstone group, which makes the comparison a lithological as well as a numerical one.

Our modelled shifts span that published range rather than sitting outside it. At 30 Hz under the uniform topology, the Control case gives $\Delta t_{2\mathrm{way}} = 6.5$ ms, the Rotliegend-style case 1.6 ms, and the Bunter-style case 0.69 ms (Table F.3). The Rotliegend-style value falls comfortably inside the published CO_2 range and three times above the towed-streamer threshold. The Control value exceeds the published range, which is expected rather than encouraging: that case is a deliberately soft, shallow reference construction and not a Dutch storage target. Under the gravity-segregated topology, all three shifts fall, most sharply for the Control case, because the reduced saturation over the lower eighty percent of the reservoir outweighs the higher saturation in the gas-rich cap when the shift is integrated over the full interval.

The comparison also exposes a tension that the amplitude analysis alone does not show. The Bunter-style case, which is the one case whose relaxation peak lies inside the seismic band and therefore the only one where the dispersive correction survives migration, carries the *smallest* time shift of the three. At 30 Hz, it sits just above the 0.5 ms threshold, and by 50 Hz it has fallen to 0.45 ms, below it. The two detectability levels identified in Section 6.1 therefore do not reinforce one another in this case: the frequency band that makes the dispersive signature visible is also the band in which the travel-time signature weakens toward the threshold. A monitoring design for a Bunter-type Dutch reservoir would have to choose which of the two observables to optimise for, and our results indicate it cannot optimise for both.

Three qualifications bound this comparison. Our shifts are produced by fluid substitution alone, whereas the published CO_2 values combine saturation and pressure effects, and Toh et al. (2025) noted that in depleted gas fields the pressure response dominates the subtler saturation response. A like-for-like comparison would therefore require the pressure contribution to be modelled alongside the saturation contribution, which we have not done. Separately, the threshold itself is an acquisition-specific figure for towed-streamer data. Under the definition of detectability adopted here, set out in Section 6.4, it enters as an external reference point for scale, not as a criterion our modelling is equipped to test. Third, the pore fluids differ acoustically: the speed of sound in CO_2 is much lower than in hydrogen, so the velocity contrast produced per unit saturation change differs between the two gases, and the published CO_2 magnitudes are not directly transferable to hydrogen.

6.4 Limitations and Assumptions

Our interpretation holds within a defined scope, and the boundaries of that scope condition how far our interpretation is carried.

What we mean by detectability. Throughout this work, detectability is assessed relative to the fluid effect alone. We ask whether a fluid-substitution change, and specifically its frequency-dependent part, survives forward modelling and migration as a resolvable change in the imaged section. We do not assess whether that change would be recoverable from a real time-lapse survey. Field detectability depends, in addition, on the acquisition geometry and repeatability, the ambient and source-generated noise level, the quality and repeatability of the processing sequence, and the accumulated non-repeatability between survey vintages. All of these are excluded here by construction. The prescribed signal-to-noise-ratio study of Section 5.5.5 is a sensitivity of this idealised result to additive band-limited noise, not a model of field acquisition noise, and it does not change this scope. Our result should therefore be read as an upper bound on what is available to a field survey, not as a prediction of what a field survey would return.

Dispersion-only propagation closure. The reflectivity method we use accepts one real scalar velocity per layer, and so it carries velocity dispersion but not intrinsic attenuation. It cannot propagate a complex velocity or a finite quality factor through the section. The tuned state therefore injects the dispersive response through the real part of the Johnson velocity at the imaging frequency, which shifts the reflector, while the imaginary part (the intrinsic attenuation the same model predicts) is not carried into propagation or imaging. Our migrated amplitudes are accordingly a frequency-domain dispersion proxy rather than a fully attenuating time-domain waveform. Carrying attenuation as well as dispersion would require a viscoelastic propagator admitting a complex, frequency-dependent modulus. This closure bounds the magnitude of the relaxed-to-tuned difference, but it does not change which case carries the effect, since that ordering is fixed by the position of the relaxation peak relative to the band.

Lateral invariance and the multiple-free pass. The geometry is a laterally invariant layer-cake column, treated in one and a half dimensions: a one-dimensional medium imaged through a two-dimensional prestack record. Lateral homogeneity is an assumption we make so that the reservoir-fluid response can be isolated from structural and lateral-heterogeneity effects. Our migrated amplitudes therefore report what survives propagation and imaging through a layered but laterally uniform medium, and nothing about lateral variation. Internal and surface multiples are not modelled in this pass. The layered-overburden step tests robustness to transmission through the section and to a structured migration background, but not to multiple contamination. This is a material omission for the Dutch setting, where the Cretaceous Chalk Group is a strong multiple generator and overprints reflections at all levels below it.

The controlled salt-below-caprock choice. Our layered-overburden comparison holds the layer immediately above the reservoir identical between the single-block and layered treatments, so that the top-reservoir reflection coefficient is unchanged and any difference in the migrated result comes from transmission through the column above rather than from the reflector itself. A fully faithful Rotliegend column would instead place the Zechstein salt directly on the reservoir, changing the top-reservoir reflection coefficient and with it the absolute monitoring signal. We deliberately do not impose that geometry here, since doing so would confound the two effects we wish to separate.

Elastically static overburden. The overburden is elastically static throughout: we do not model geomechanical changes in the overburden induced by injection and pressure change, such as the strain-related velocity changes described in time-lapse studies by an R-factor. Overburden differences in this work are therefore purely transmission effects, and any geomechanical time-lapse response would superpose on the fluid effect modelled here.

Single geometry and single relaxation spectrum. Our scenarios use the concentric-sphere gas-inner geometry throughout. We quantified the geometry sensitivity at the interface level, where the sphere and slab idealisations differ by of order twenty percent in impedance change for a matched patch size, but a single geometry and a single representative patch scale are carried through the migrated results. A multi-scale patch-size distribution rather than a representative single scale would broaden the dispersion curve (Pride et al., 2004) and is not implemented here.

Patch radius not pinned to reservoir data. The patch radius is not a directly measurable field parameter. The range used in our parameter scans brackets typical mesoscopic scales but is not constrained by facies-scale reservoir data, and the case ordering depends on it. Connecting the patch radius to realistic Dutch injection scenarios remains a separate step.

Initial state and mechanism scope. The initial state throughout is fully brine-saturated, whereas a practical UHS site in a depleted Dutch gas field would begin with residual hydrocarbon gas in place. Squirt flow is likewise excluded: it operates at higher frequencies than the seismic band and is deferred by scope, not by oversight (Pride et al., 2004; Müller et al., 2010). Angle-dependent behaviour is not treated, since our monitoring proxies are taken at normal incidence (Bagheri et al., 2025).

6.5 Directions indicated by the limitations

Each of the limitations above points to a specific extension of the modelling ladder, and the supporting literature for those extensions is collected here. The extensions themselves are stated as future work in Section 7.2.

Moving beyond a single representative patch scale requires a multi-scale patch-size distribution, which broadens the dispersion curve rather than shifting it (Pride et al., 2004). This is the step that would connect the patch radius to facies-scale reservoir data and so convert our conditional result from a statement about parameter space into a statement about named Dutch fields.

Carrying intrinsic attenuation as well as velocity dispersion through propagation and imaging requires a viscoelastic or viscoacoustic treatment admitting a complex, frequency-dependent modulus. Among the available parametrisations of such a modulus, a so-called Cole–Cole representation is an established choice in the seismic modelling literature (Picotti et al., 2010; Picotti & Carcione, 2017). We did not implement it here, and therefore, we cannot make any claim about its accuracy relative to alternatives for hydrogen and brine patchy saturation.

Extending the analysis from normal-incidence monitoring proxies to amplitude-versus-angle attributes would build on the angle-dependent image gathers of Bagheri et al. (2025), which carry additional fluid and saturation discrimination beyond what a zero-offset proxy can resolve.

Making the assessment field-realistic rather than idealised requires two further steps: replacing the prescribed signal-to-noise ratios of Section 5.5.5 with survey-specific noise and repeatability models, and imposing a depth-dependent upper frequency limit representing intrinsic attenuation in the overburden. Neither of these is implemented here.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

We posed the research question whether adding frequency-dependent patchy-saturation physics materially changes the seismic detectability interpretation that would be drawn from a low-frequency Gassmann reference model, under Dutch reservoir conditions. Our conclusions are as follows.

- Patchy hydrogen saturation produces velocity dispersion and attenuation through mesoscopic wave-induced fluid flow. The size of the in-band correction is governed by the position of the peak relaxation frequency relative to the seismic band, and not by the stiffness of the rock frame.

- The position of that relaxation peak is set primarily by the patch radius and the permeability. Of the three Dutch reference cases we modelled, only the larger-patch Bunter-style case places the peak inside the seismic band; the Control and Rotliegend-style cases place it above the band and therefore sample essentially the relaxed limit between 15 and 100 Hz.

- Within the concentric-sphere gas-inner geometry shared by the three models we tested, the scalable branching-function formulation reproduces the exact solution and tracks the external benchmark. This justifies carrying the branching-function formulation into the parameter scans and the migrated scenarios. The agreement is a consistency result within one shared geometry, and it does not transfer to other patch geometries, whose relaxation-time and shape parameters differ.

- Detectability separates into two levels that must be kept separate. The full monitoring signal produced by replacing reservoir brine with patchy hydrogen is large in every case and under every saturation topology we modelled, confirming that hydrogen injection is seismically visible in all three Dutch settings. This level is already predicted by the low-frequency reference model. The second level is the dispersive signature: the frequency-dependent correction to that signal, which the following conclusions concern.

- The two levels of detectability separate further under acquisition noise. The dispersive signature is recoverable at signal-to-noise ratios within field practice only for the Bunter-style case; for the other two cases, it would require ratios one to two orders of magnitude beyond it. The full monitoring signal, by contrast, remains recoverable at a signal-to-noise ratio of five in all three settings. The fragility we identify belongs to the frequency-dependent refinement, not to the detectability of hydrogen injection itself.

- Where the relaxation peak lies inside the seismic band, the frequency-dependent correction is large and opposes the bulk fluid-substitution response. In the Bunter-style case, a low-frequency description overpredicts the imaged monitoring signal by a factor of between 1.3 and 2.0 across the applied band. Where the peak lies above the band, in the Control and Rotliegend-style cases, the same comparison gives factors between 1.00 and 1.02, and the low-frequency description is accurate to within a few percent.

- The answer to our detectability question is therefore conditional rather than absolute. Frequency-dependent patchy saturation changes the quantitative monitoring interpretation in an identifiable region of Dutch reservoir parameter space, characterised by larger patch radii and higher permeability, and is otherwise diluted to negligibility by the imaging workflow.

- Saturation topology modulates the magnitude of the dispersive correction but does not change which case carries it.

- Where the dispersive correction is recoverable at all, its recoverability is stable under a realistic overburden. The overburden enters the migrated amplitude multiplicatively, so it scales the relaxed-to-tuned difference rather than cancelling it: where that difference is large it is insensitive to the overburden treatment, and where it is small the overburden transmission response controls its value. The case ordering is stable under both treatments; the magnitude of the correction is not, and cannot be quoted from a simplified overburden.

- Randomising the patch radius and saturation independently per reservoir layer qualifies but does not overturn this picture. The dispersive correction survives in the Bunter-style ensemble mean, but it is markedly weaker and more uncertain than any single idealised topology suggests. Realistic per-layer heterogeneity smears the in-band dispersive signal without erasing it in the mean.

- These conclusions concern the fluid effect alone. They describe what is available to a seismic survey under ideal acquisition, noise-free recording, and repeatable processing, and they are therefore an upper bound on field detectability rather than a prediction of it.

7.2 Future work

The limitations set out in Chapter 6 map directly onto the most defensible extensions of the modelling ladder. The supporting literature for each of the directions below is given in Section 6.5.

- The first is the salt-directly-on-reservoir geometry for the Rotliegend-style case, together with internal and surface multiples from the layered section. This is a more stringent test of the absolute monitoring signal than the controlled comparison used here, and it is the natural way to make the Rotliegend overburden field-faithful. For the Dutch setting, it would also admit the strong multiple generation associated with the Cretaceous chalk.

- The second is a laterally heterogeneous, fully two-dimensional treatment, which would relax the lateral-invariance assumption and admit the structural and lateral-variation effects that a layer-cake model excludes by design.
- The third is a pinning of the patch radius to reservoir-specific data, by combining facies-scale reservoir simulation output with capillary-entry analysis, and, beyond a single representative scale, a multi-scale patch-size distribution. This is what would convert our conditional result from a statement about parameter space into a statement about specific named Dutch fields.
- The fourth is to replace the dispersion-only propagator with viscoelastic or viscoacoustic modelling that carries attenuation through propagation and imaging, using a complex frequency-dependent modulus. Driving a viscoelastic reflectivity or finite-difference synthetic and the matching migration would carry the intrinsic attenuation alongside the velocity dispersion, so that the time-domain waveform consequences of patchy saturation, and not only the band-by-band reflector shift, could be assessed.
- The fifth is to move from an idealised assessment of the fluid effect toward a field-realistic one. The noise study reported in this work prescribes the signal-to-noise ratio directly; the field-realistic step is to replace that prescription with survey-specific noise and repeatability models, and to impose a depth-dependent upper frequency limit representing intrinsic attenuation in the overburden. This would convert our upper bound on detectability into a statement about the survey conditions under which the dispersive correction is actually recoverable.
- The sixth is a broader treatment of the partial-saturation closure, comparing the uniform-mixing limit and an empirical patchiness parametrisation against the frequency-dependent formulation used here, so that the cost of each simplification can be stated quantitatively rather than argued from first principles.

Two further extensions are supported by our existing framework. A depleted-field initial state with residual methane would make the pre-injection state field-faithful for Dutch depleted gas fields. An angle-dependent treatment, building on angle-dependent image gathers, would extend the analysis from the normal-incidence monitoring proxies used here to amplitude-versus-angle attributes, which carry additional fluid and saturation discrimination.

7.3 Closing statement

We have established that hydrogen injection into Dutch sandstone reservoirs is robustly detectable at the Gassmann level, and that the frequency-dependent patchy-saturation signature is a conditional, case-dependent refinement of that detectability rather than a uniform addition to it. The dispersive correction survives migration, and reaches the order of the full monitoring signal, only where the rock physics places the relaxation peak inside the seismic band, which among our Dutch cases means the larger-patch Bunter-style geometry. Elsewhere it is diluted by the imaging workflow to a few percent. This places our contribution as a quantitative refinement of the monitoring interpretation in a specific and identifiable region of Dutch reservoir parameter space, with the realism extensions above as the route to a field-faithful assessment.

Bibliography

Bagheri, A. R., Verschuur, D. J., & Draganov, D. (2025). Feasibility of seismic monitoring of hydrogen storage and leakage in sandstone reservoirs using angle-dependent image gathers. *Sixth EAGE Global Energy Transition Conference & Exhibition*, *2025*(1), 1–5. https://doi.org/10.3997/2214-4609.202521176

Bagheri, A. R., Draganov, D., & Verschuur, E. (2026). Seismic modelling of underground hydrogen storage [Delphi Consortium, Delft University of Technology]. *Delphi Consortium Characterization & Monitoring (C&M) Volume*, *24*, 3–18.

Batzle, M., & Wang, Z. (1992). Seismic properties of pore fluids. *Geophysics*, *57*(11), 1396–1408. https://doi.org/10.1190/1.1443207

Brie, A., Pampuri, F., Marsala, A. F., & Meazza, O. (1995). Shear sonic interpretation in gas-bearing sands. *SPE Annual Technical Conference and Exhibition*, (30595-MS). https://doi.org/10.2118/30595-MS

Carcione, J. M., Helle, H. B., & Pham, N. H. (2003). White's model for wave propagation in partially saturated rocks: Comparison with poroelastic numerical experiments. *Geophysics*, *68*(4), 1389–1398. https://doi.org/10.1190/1.1598132

Dupuy, B., Romdhane, A., Eliasson, P., Querendez, E., Yan, H., Torres, V. A., & Ghaderi, A. (2017). Quantitative seismic characterization of CO2 at the Sleipner storage site, North Sea. *Interpretation*, *5*(4), SS23–SS42. https://doi.org/10.1190/INT-2017-0013.1

Dutta, N. C., & Odé, H. (1979). Attenuation and dispersion of compressional waves in fluid-filled porous rocks with partial gas saturation (White model)—part I: Biot theory. *Geophysics*, *44*(11), 1777–1788.

Dutta, N. C., & Seriff, A. J. (1979). On White's model of attenuation in rocks with partial gas saturation. *Geophysics*, *44*(11), 1806–1812.

Eid, R., Ziolkowski, A., Naylor, M., & Pickup, G. (2015). Seismic monitoring of CO2 plume growth, evolution and migration in a heterogeneous reservoir: Role, impact and importance of patchy saturation. *International Journal of Greenhouse Gas Control*, *43*, 70–81. https://doi.org/10.1016/j.ijggc.2015.10.019

Gassmann, F. (1951). Über die elastizität poröser medien. *Vierteljahrsschrift der Naturforschenden Gesellschaft in Zürich*, *96*, 1–23.

Gei, D., Carcione, J. M., & Picotti, S. (2024). The physics of seismic propagation in aquifers partially saturated with H2. comparison with the cases of CO2 and CH4. *Bulletin of Geophysics and Oceanography*, *65*(4), 611–632. https://doi.org/10.4430/bgo00470

He, Y.-X., Wang, S. X., & Wu, X. Y. (2020). Effects of patch size changes on the time-lapse seismic reflections of highly attenuating reservoirs. *Journal of Applied Geophysics*, *172*, 103878. https://doi.org/10.1016/j.jappgeo.2019.103878

Heinemann, N., Alcalde, J., Miocic, J. M., Hangx, S. J. T., Kallmeyer, J., Ostertag-Henning, C., Hassanpouryouzband, A., Thaysen, E. M., Strobel, G. J., Schmidt-Hattenberger, C., Edlmann, K., Wilkinson, M., Bentham, M., Haszeldine, R. S., Carbonell, R., & Rudloff, A. (2021). Enabling large-scale hydrogen storage in porous media – the scientific challenges. *Energy & Environmental Science*, *14*, 853–864. https://doi.org/10.1039/d0ee03536j

Johnson, D. L. (2001). Theory of frequency dependent acoustics in patchy-saturated porous media. *Journal of the Acoustical Society of America*, *110*(2), 682–694. https://doi.org/10.1121/1.1381021

Kazemeini, S. H., Juhlin, C., & Fomel, S. (2010). Monitoring CO2 response on surface seismic data; a rock physics and seismic modeling feasibility study at the CO2 sequestration site, Ketzin, Germany. *Journal of Applied Geophysics*, *71*, 109–124. https://doi.org/10.1016/j.jappgeo.2010.05.004

Kennett, B. L. N. (1983). *Seismic wave propagation in stratified media*. Cambridge University Press.

Lebedev, M., Toms-Stewart, J., Clennell, B., Pervukhina, M., Shulakova, V., Paterson, L., Müller, T. M., Gurevich, B., & Wenzlau, F. (2009). Direct laboratory observation of patchy saturation and its effects on ultrasonic velocities. *The Leading Edge*, *28*(1), 24–27. https://doi.org/10.1190/1.3064142

Loyola, A., Voskov, D., Farajzadeh, R., Bisdom, K., de Borst, K., & Geiger, S. (2026). Geological controls on underground hydrogen storage in depleted gas fields. *Marine Geoscience and Energy Resources*, *193*, 207819. https://doi.org/10.1016/j.marger.2026.207819

MacBeth, C., Toh, S. Y., Kopydlowska, B., & Jafarizadeh, F. (2024). Evaluation of signals from monitoring CO2 injection in the North Sea using 4D seismic. *85th EAGE Annual Conference & Exhibition*, *2024*(1), 1–5. https://doi.org/10.3997/2214-4609.2024101143

MacBeth, C. (2004). A classification for the pressure-sensitivity properties of a sandstone rock frame. *Geophysics*, *69*(2), 497–510. https://doi.org/10.1190/1.1707070

Mavko, G., & Mukerji, T. (1998). Bounds on low-frequency seismic velocities in partially saturated rocks. *Geophysics*, *63*(3), 918–924. https://doi.org/10.1190/1.1444402

Müller, T. M., Gurevich, B., & Lebedev, M. (2010). Seismic wave attenuation and dispersion resulting from wave-induced flow in porous rocks — a review. *Geophysics*, *75*(5), 75A147–75A164. https://doi.org/10.1190/1.3463417

Papageorgiou, G., Amalokwu, K., & Chapman, M. (2016). Theoretical derivation of a Brie-like fluid mixing law. *Geophysical Prospecting*, *64*, 1048–1053. https://doi.org/10.1111/1365-2478.12380

Peng, D. Y., & Robinson, D. B. (1976). A new two-constant equation of state. *Industrial & Engineering Chemistry Fundamentals*, *15*(1), 59–64. https://doi.org/10.1021/i160057a011

Picotti, S., & Carcione, J. M. (2017). Numerical simulation of wave-induced fluid flow seismic attenuation based on the Cole–Cole model. *Journal of the Acoustical Society of America*, *142*(1), 134–145. https://doi.org/10.1121/1.4990965

Picotti, S., Carcione, J. M., Rubino, J. G., Santos, J. E., & Cavallini, F. (2010). A viscoelastic representation of wave attenuation in porous media. *Computers & Geosciences*, *36*, 44–53. https://doi.org/10.1016/j.cageo.2009.07.003

Pride, S. R., Berryman, J. G., & Harris, J. M. (2004). Seismic attenuation due to wave-induced flow. *Journal of Geophysical Research: Solid Earth*, *109*, B01201. https://doi.org/10.1029/2003JB002639

Queißer, M., & Singh, S. C. (2013). Full waveform inversion in the time lapse mode applied to CO2 storage at Sleipner. *Geophysical Prospecting*, *61*(3), 537–555. https://doi.org/10.1111/j.1365-2478.2012.01072.x

Reijnen-Mooij, G., van Klaveren, S., Jaarsma, B., Haagenson, R., Draganov, D., & Hajibeygi, H. (2025). *Summary of main characteristics of the depleted gas fields in the Netherlands that could be candidates for UHS* (Deliverable D5.4.1). GroenvermogenNL WP2 R&D Project HyTROS.

Rozhko, A. Y. (2020). Effective fluid bulk modulus in the partially saturated rock and the amplitude dispersion effects. *Journal of Geophysical Research: Solid Earth*, *125*(3), e2019JB018693. https://doi.org/10.1029/2019JB018693

Rubino, J. G., & Holliger, K. (2012). Seismic attenuation and velocity dispersion in heterogeneous partially saturated porous rocks. *Geophysical Journal International*, *188*(3), 1088–1102. https://doi.org/10.1111/j.1365-246X.2011.05291.x

Smith, T. M., Sondergeld, C. H., & Rai, C. S. (2003). Gassmann fluid substitutions: A tutorial. *Geophysics*, *68*(2), 430–440. https://doi.org/10.1190/1.1567211

Stekelenburg, J. (2026, March). *Fluid substitution for H_2 injection detectability: Preliminary research for Master Thesis* (Research Module Report). RWTH Aachen University / Delft University of Technology.

Terzaghi, K. (1943). *Theoretical soil mechanics* [DOI is Wiley's retrospective electronic-book DOI for the 1943 work.]. John Wiley & Sons, Inc. https://doi.org/10.1002/9780470172766

Toh, S. Y., MacBeth, C., & Landa, J. (2025). Feasibility of monitoring CO2 injection in depleted UK Continental Shelf and Norwegian Continental Shelf gas fields using 4D seismic. *Geoenergy*, *3*(1), geoenergy2025–003. https://doi.org/10.1144/geoenergy2025-003

Wang, G., Pickup, G. E., Sorbie, K. S., & MacKay, E. J. (2022). Driving factors for purity of withdrawn hydrogen: A numerical study of underground hydrogen storage with various cushion gases [Paper SPE-209625-MS; conference held in Madrid, Spain, 6–9 June 2022]. *SPE Europe Energy Conference Featured at the 83rd EAGE Annual Conference and Exhibition 2022*. https://doi.org/10.2118/209625-MS

White, J. E. (1975). Computed seismic speeds and attenuation in rocks with partial gas saturation. *Geophysics*, *40*(2), 224–232.

Wood, A. B. (1955). *A textbook of sound: Being an account of the physics of vibrations with special reference to recent theoretical and technical developments* (3rd revised) [xvi + 610 pages]. G. Bell; Sons, Ltd.

Zivar, D., Kumar, S., & Foroozesh, J. (2021). Underground hydrogen storage: A comprehensive review. *International Journal of Hydrogen Energy*, *46*, 23436–23462. https://doi.org/10.1016/j.ijhydene.2020.08.138

Appendix A

Supplementary Baseline Equations

This appendix collects the low-frequency fluid-property and monitoring-proxy expressions retained from the baseline workflow. They remain useful as reference formulas, but they are not the Johnson branching-function dynamic-modulus closure used in Chapter 4.

A.1 Fluid-property equations

Brine bulk modulus. Given brine density ρ_b and brine P-wave velocity V_b from the Batzle–Wang correlations (Batzle & Wang, 1992), the brine bulk modulus is

$$K_b = \rho_b V_b^2. \tag{A.1}$$

Gas properties. After solving the Peng–Robinson equation of state (Peng & Robinson, 1976) for the compressibility factor Z, molar volume and density follow as

$$V_m = \frac{ZRT}{P}, \qquad \rho_g = \frac{M}{V_m}, \tag{A.2}$$

and the gas bulk modulus follows from the compressibility:

$$C_T = -\frac{1}{V_m}\frac{dV_m}{dP}, \qquad C_S = \frac{C_T}{\gamma}, \qquad K_g = \frac{1}{C_S}. \tag{A.3}$$

A.2 Monitoring proxies

The low-frequency observables retained from the baseline are

$$Z = \rho_{\text{bulk}} V_p, \tag{A.4}$$

$$R = \frac{Z_2 - Z_1}{Z_2 + Z_1}, \tag{A.5}$$

and

$$\Delta t_{\text{2way}} = 2h\left(\frac{1}{V_{p,\text{mon}}} - \frac{1}{V_{p,\text{base}}}\right). \tag{A.6}$$

A.3 Wood–Brie closure sensitivity

The static partial-saturation closure used in the original Stekelenburg (2026) baseline can be either Wood mixing (Eq. (2.11)) or the empirical Brie relation (Eq. (2.13)) with a stiffening exponent e. Wood mixing is the active closure used in Chapter 3 and is the relaxed limit that the frequency-dependent models of Chapter 4 must recover as $\omega \to 0$. Brie mixing is retained only as a baseline-reference comparison, because the e-controlled stiffening makes it useful for showing how a stronger mixing law would moderate the low-saturation response.

Figure A.1 reproduces the saturation dependence of the three Research Module time-lapse proxies at the Control reference parameter point under Wood and Brie partial-saturation mixing. The Wood curves show the steep low-saturation response discussed in Chapter 3; the Brie curves show how a non-zero stiffening parameter delays and softens that response. The sign and qualitative shape of the response are preserved across the two closures, which is what justifies retaining Wood as the active relaxed-limit reference in this thesis.

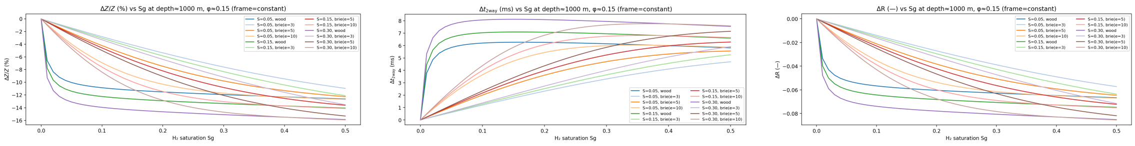

(a) (b) (c)

Figure A.1: Relative impedance change, two-way traveltime change, and reflection-coefficient change are shown against hydrogen saturation for Wood and Brie mixing at the Control point. (a) $\Delta Z/Z$. (b) $\Delta t_{2\text{way}}$. (c) ΔR. Reproduced from Stekelenburg (2026).

A.4 Porosity sensitivity by formation

The Upper and Middle Bunter porosity sweeps complement the Rotliegend exemplar in Figure 3.7. The Middle Bunter interval extends to lower porosity and a deeper reference depth than the Upper Bunter interval, broadening the elastic-property envelope while preserving the same qualitative saturation response.

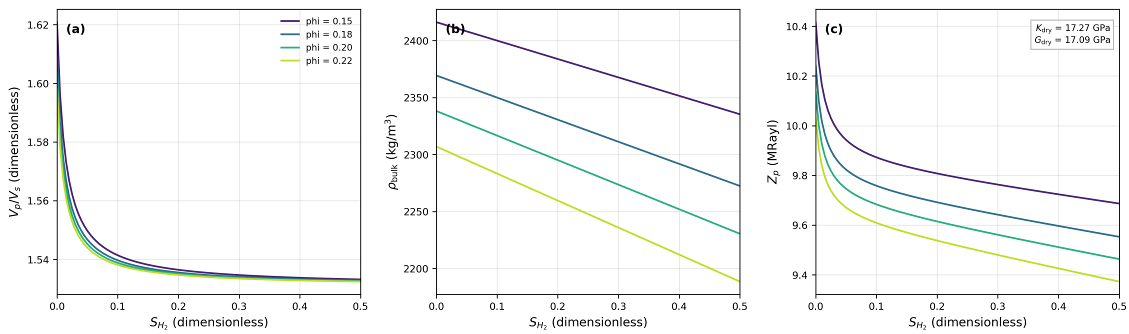

Figure A.2: V_p/V_s, bulk density ρ_{bulk}, and P-wave impedance Z_p are shown against hydrogen saturation for Upper Bunter porosities $\phi \in \{0.15, 0.18, 0.20, 0.22\}$ at $z =$ 2200 m and $S_b = 0.15$. (a) V_p/V_s. (b) ρ_{bulk}. (c) Z_p.

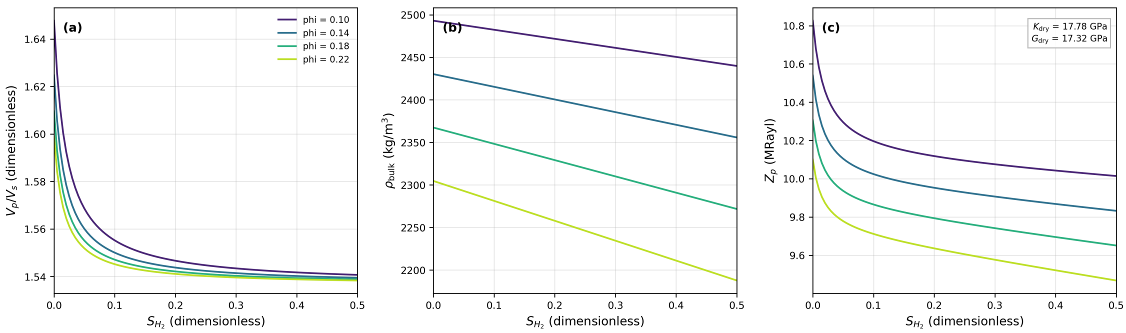

Figure A.3: V_p/V_s, bulk density ρ_{bulk}, and P-wave impedance Z_p are shown against hydrogen saturation for Middle Bunter porosities $\phi \in \{0.10, 0.14, 0.18, 0.22\}$ at $z = 2900$ m and $S_b = 0.15$. (a) V_p/V_s. (b) ρ_{bulk}. (c) Z_p.

Appendix B

Dry Rock-Frame Model Equations

This appendix documents the two dry-frame parameterisations retained as inputs to the frequency-dependent workflow.

Symbols used in this appendix: $p(z)$ and $p_c(z)$ are the pore and confining pressures at depth z, p_d the differential pressure, p_0 the surface pressure, g_p the hydrostatic pressure gradient, ρ_{ob} the overburden density, and g the gravitational acceleration. In the Krief–Athy frame, ϕ_0 is the surface porosity, p_ϕ the compaction pressure constant, K_s and G_s the mineral bulk and shear moduli, and A and B the Krief exponent parameter and the shear proportionality constant of Gei et al. (2024). The MacBeth parameters K_∞, G_∞, E_K, E_G, P_K, P_G, S_K, and S_G are defined where they appear in Section B.3.

B.1 Depth-to-pressure mapping

Following Terzaghi (1943), we define differential pressure as

$$p_d(z) = p_c(z) - p(z), \tag{B.1}$$

The confining pressure is approximated by the overburden pressure in this one-dimensional setting. The corresponding depth mapping is

$$p(z) = p_0 + g_p\, z, \qquad p_c(z) = p_0 + \rho_{\mathrm{ob}}\, g\, z. \tag{B.2}$$

B.2 Gei-style Krief–Athy frame

Following Gei et al. (2024), we prescribe porosity as

$$\phi(z) = \phi_0 \exp\left(-\frac{p_d(z)}{p_\phi}\right), \tag{B.3}$$

the dry bulk modulus is

$$K_{\mathrm{dry}}(\phi) = K_s(1-\phi)^{A/(1-\phi)}, \tag{B.4}$$

and the dry shear modulus is

$$G_{\mathrm{dry}}(\phi) = \frac{B\, G_s}{K_s}\, K_{\mathrm{dry}}(\phi). \tag{B.5}$$

B.3 MacBeth-style pressure-sensitive frame

The default dry rock-frame model used in the reservoir cases of Chapter 4 is the three-parameter excess-compliance relation of MacBeth (2004). The functional form is motivated by the closure of a distribution of compliant cracks and grain contacts under increasing differential pressure, modelled as a single decaying exponential in the excess compliance. Following MacBeth (2004, Eqs. 3–4),

$$K_{\mathrm{dry}}(p_d) = \frac{K_\infty}{1 + E_K\, \exp(-p_d/P_K)}, \tag{B.6}$$

$$G_{\mathrm{dry}}(p_d) = \frac{G_\infty}{1 + E_G\, \exp(-p_d/P_G)}, \tag{B.7}$$

where p_d is the differential pressure of Eq. (B.2), K_∞ and G_∞ are the high-pressure asymptotes of the dry bulk and shear moduli, E_K and E_G are dimensionless excess-compliance amplitudes, and P_K and P_G are the characteristic pressure constants. The amplitudes relate to the overall stress-sensitivity indices through (MacBeth, 2004, Eqs. 5–6)

$$S_K = \frac{E_K}{1 + E_K}, \qquad S_G = \frac{E_G}{1 + E_G}, \tag{B.8}$$

so the implementation parametrises each modulus by the triplet (K_∞, S_K, P_K) and recovers $E_K = S_K/(1 - S_K)$ internally. MacBeth (2004) uses confining pressure as the loading variable in his database; the database measurements are differential-pressure tests with controlled pore pressure, so applying Eqs. (B.6)–(B.7) with p_d as the argument is consistent with both the original calibration and the basin-style p_d convention of Gei et al. (2024) adopted elsewhere in this thesis.

Parameter values for the two reservoir cases of Chapter 4 are drawn directly from MacBeth (2004, Table 2). The Rotliegend-style case uses the Rotliegend (southern North Sea) Facies B group restricted to confining pressures above the in-situ effective stress (the italicised re-fit row of Table 2), which matches the case $p_d \approx 33$ MPa and the deeper-Rotliegend porosity range 0.10–0.20 reported by Reijnen-Mooij et al. (2025). The Bunter-style case uses the Berea Sandstone group, which is the closest clean continental fluvial analogue available in MacBeth's database; the Berea porosity range 0.16–0.22 brackets the HyTROS Upper Bunter porosity range 0.20–0.27 (Reijnen-Mooij et al., 2025), and Berea's depositional setting (clean Devonian fluvial-deltaic quartz arenite) is petrologically similar to clean Triassic Bunter sandstones. No Triassic Buntsandstein samples are present in MacBeth's database, so this analogue choice is necessarily approximate; the Forties Sandstone groups either fall outside the HyTROS porosity range (Facies B at $\phi = 0.13$) or correspond to anomalously soft Paleocene marine turbidites (Facies A at $K_\infty = 9.15$ GPa), making them less defensible analogues than Berea. The numerical triplets used in this thesis are listed with the other modelling cases in Table 3.1.

This formulation supersedes the approximation $K_{\mathrm{dry}}(p_d) = K_{\mathrm{ref}}\,(p_d/p_{d,\mathrm{ref}})^{n_K}$ used in earlier iterations of this work and in Stekelenburg (2026). The power-law form captures the qualitative stiffening of K_{dry} with p_d but does not reproduce the high-pressure plateau characteristic of the soft-pore-closure regime. It is retained as a legacy option in the codebase for reproducibility checks. The Krief–Athy frame of Section B.2 is retained as the option used for the Gei et al. (2024) comparison only; it is not the default for the reservoir cases because its compaction-driven $\phi(z)$ couples porosity to depth in a way that conflates frame-stiffness uncertainty with porosity uncertainty.

Note on P_G uncertainty for the Rotliegend Facies B re-fit. The shear-modulus characteristic pressure P_G in the Rotliegend Facies B re-fit row of MacBeth (2004, Table 2) is large ($P_G = 56.47$ MPa) and weakly constrained by the high-pressure shear data. At the case pressure $p_d = 33$ MPa, G_{dry} has therefore not reached its high-pressure asymptote. This is a feature of the data, not an implementation choice; the bulk-modulus parameter $P_K = 15.75$ MPa is much better constrained.

Appendix C

Johnson Model Details

This appendix gives the support formulas used to compute the two parameters of the Johnson branching-function model from the asymptotic construction implemented in the repository.

Symbols used in this appendix: K_0, K_{dry}, and G_{dry} are the mineral, dry-frame bulk, and dry-frame shear moduli; ϕ the porosity and κ the permeability; $K_{f,j}$ and η_j the fluid bulk modulus and viscosity of region j, with $j = 1$ the gas-saturated and $j = 2$ the brine-saturated region, and S_j the corresponding saturation; a and b the inner and outer radii of the concentric-sphere cell and L_1 and L_2 the slab thicknesses; K_W the Wood effective fluid modulus and K_{GW} and K_{GH} the Gassmann–Wood and Gassmann–Hill limits of Chapter 2; D_j the Biot diffusivity and P_j, Q_j, R_j the Biot coefficients of region j; Δp_f the interregion pore-pressure difference, P_e the applied effective stress, and S/V the patch surface-to-volume ratio.

C.1 Notation mapping

The code uses the notation

$$K_s \to K_0, \qquad K_b \to K_{\text{dry}}, \qquad N \to G_{\text{dry}}, \tag{C.1}$$

relative to Johnson's original notation (Johnson, 2001). This mapping is applied consistently throughout the implementation.

C.2 Biot coefficient and Wood effective fluid

The Biot coefficient and modulus are

$$\alpha = 1 - \frac{K_{\text{dry}}}{K_0}, \qquad M = \left(\frac{\phi}{K_f} + \frac{\alpha - \phi}{K_0} \right)^{-1}. \tag{C.2}$$

The Wood effective fluid modulus is

$$K_W = \left(\frac{S_b}{K_b} + \frac{S_g}{K_g} \right)^{-1}. \tag{C.3}$$

C.3 Biot P,Q,R coefficients

The code evaluates the Biot coefficients in Johnson's Eq. (5) form as

$$D_{PQR} = 1 - \phi - \frac{K_{\text{dry}}}{K_0} + \phi \frac{K_0}{K_f}, \tag{C.4}$$

$$P = \frac{(1-\phi)\left(1-\phi-\frac{K_{\mathrm{dry}}}{K_0}\right)K_0 + \phi\left(\frac{K_0}{K_f}\right)K_{\mathrm{dry}}}{D_{PQR}} + \frac{4}{3}G_{\mathrm{dry}}, \tag{C.5}$$

$$Q = \frac{\left(1-\phi-\frac{K_{\mathrm{dry}}}{K_0}\right)\phi K_0}{D_{PQR}}, \qquad R = \frac{\phi^2 K_0}{D_{PQR}}. \tag{C.6}$$

C.4 Diffusivity

The exact Biot diffusivity used in the asymptotic support formulas is

$$D = \frac{\kappa}{\eta\phi^2}\frac{PR-Q^2}{P+2Q+R}. \tag{C.7}$$

The simplified brine-side reference diffusivity used only in the simplified Johnson path is

$$D_0 = \frac{\kappa}{\eta_b}\frac{M_b L_{\mathrm{dry}}}{H_2}, \tag{C.8}$$

with

$$M_b = \left(\frac{\phi}{K_b} + \frac{\alpha-\phi}{K_0}\right)^{-1}, \qquad L_{\mathrm{dry}} = K_{\mathrm{dry}} + \frac{4}{3}G_{\mathrm{dry}}, \qquad H_2 = K_{\mathrm{sat}}(K_b) + \frac{4}{3}G_{\mathrm{dry}}. \tag{C.9}$$

C.5 Piecewise coupling term

The piecewise coupling term used in the low-frequency asymptotics is

$$g_j = \frac{\left(1-\frac{K_{\mathrm{dry}}}{K_0}\right)\left(\frac{1}{K_W}-\frac{1}{K_{f,j}}\right)}{1-\frac{K_{\mathrm{dry}}}{K_0}-\phi\frac{K_{\mathrm{dry}}}{K_0}+\phi\frac{K_{\mathrm{dry}}}{K_W}}. \tag{C.10}$$

C.6 Low-frequency parameter T

For concentric spheres, the implementation uses

$$\begin{aligned} T = \frac{K_{\mathrm{GW}}\phi^2}{30\kappa b^3}\Big[& \left(3\eta_2 g_2^2 + 5(\eta_1-\eta_2)g_1 g_2 - 3\eta_1 g_1^2\right)a^5 \\ & - 15\eta_2 g_2(g_2-g_1)a^3 b^2 \\ & + 5g_2\left(3\eta_2 g_2 - (2\eta_2+\eta_1)g_1\right)a^2 b^3 \\ & - 3\eta_2 g_2^2 b^5\Big]. \end{aligned} \tag{C.11}$$

For periodic slabs, the implementation uses

$$T = -\frac{K_{\mathrm{GW}}\phi^2}{24\kappa(L_1+L_2)}\left(\eta_1 g_1^2 L_1^3 + 3\eta_1 g_1 g_2 L_1^2 L_2 + 3\eta_2 g_1 g_2 L_1 L_2^2 + \eta_2 g_2^2 L_2^3\right). \tag{C.12}$$

Verification of the periodic-slab shape polynomial. The periodic-slab expression for T (Johnson Eq. 41) was verified independently during implementation. Evaluating the polynomial exactly as printed does not reproduce the slab shape parameter $\zeta = 0.77$ annotated in Johnson (2001, Fig. 1); it yields $\zeta \approx 0.19$. To resolve this, T was re-derived from first principles from Johnson's upstream general result, by numerically solving the auxiliary potential problem (Johnson Eqs. 21–22) for the periodic-slab geometry without reference to the closed-form Eq. 41. The first-principles value reproduces $\zeta = 0.77$ to better than one percent and differs from the literal Eq. 41 value by a constant factor of four, independent of gas saturation, indicating a discrepancy in the overall denominator coefficient rather than in the polynomial structure. The same first-principles procedure reproduces the concentric-sphere result (Eq. 40, $\zeta = 2.35$) and the gas-outer-shell case ($\zeta = 0.43$) exactly, confirming the method on the cases for which the published shape parameters are reproduced directly. The implementation therefore uses the first-principles coefficient, which reproduces all three benchmark geometries of Johnson (2001, Fig. 1).

Uncertainty note. The sphere expression for T is the highest-risk notation-mapping formula in the current implementation. We have checked it against the Johnson paper and benchmarked it indirectly through the reproduced validation curves, but the report should still avoid overstating certainty beyond that.

C.7 High-frequency coefficient G

The code constructs the high-frequency coefficient from

$$G = \sqrt{D^*} \left(\frac{\Delta p_f}{P_e} \right)^2 \frac{S}{V}, \tag{C.13}$$

with

$$D^* = \left[\frac{\kappa K_{\mathrm{GH}}}{\eta_1 \sqrt{D_1} + \eta_2 \sqrt{D_2}} \right]^2, \tag{C.14}$$

and

$$\frac{\Delta p_f}{P_e} = \frac{(R_2 + Q_2)\left(K_1 + \frac{4}{3}G_{\mathrm{dry}}\right) - (R_1 + Q_1)\left(K_2 + \frac{4}{3}G_{\mathrm{dry}}\right)}{\phi S_1 K_1 \left(K_2 + \frac{4}{3}G_{\mathrm{dry}}\right) + \phi(1 - S_1) K_2 \left(K_1 + \frac{4}{3}G_{\mathrm{dry}}\right)}. \tag{C.15}$$

C.8 Relaxation time and shape parameter

The Johnson model parameters are then

$$\tau = \left[\frac{K_{\mathrm{GH}} - K_{\mathrm{GW}}}{K_{\mathrm{GH}} G} \right]^2, \tag{C.16}$$

$$\zeta = \frac{K_{\mathrm{GH}} - K_{\mathrm{GW}}}{2K_{\mathrm{GW}}} \frac{\tau}{T}. \tag{C.17}$$

C.9 Sign convention

For the Johnson branching-function model of Eq. (2.17), the time-harmonic convention is $\exp(+i\omega t)$. A passive dissipative medium must therefore satisfy

$$\mathrm{Im}[K(\omega)] \geq 0. \tag{C.18}$$

Johnson's paper uses the opposite time dependence. The implementation therefore states the sign convention explicitly and confirms it through the benchmark reproductions rather than by relying on notation alone.

C.10 Permeability sweep at other saturations

These permeability-sweep variants complement the $S_{H_2} = 20\%$ exemplar in Figure 4.5 and retain its fixed $a = 5$ cm patch radius and constant-frame assumption.

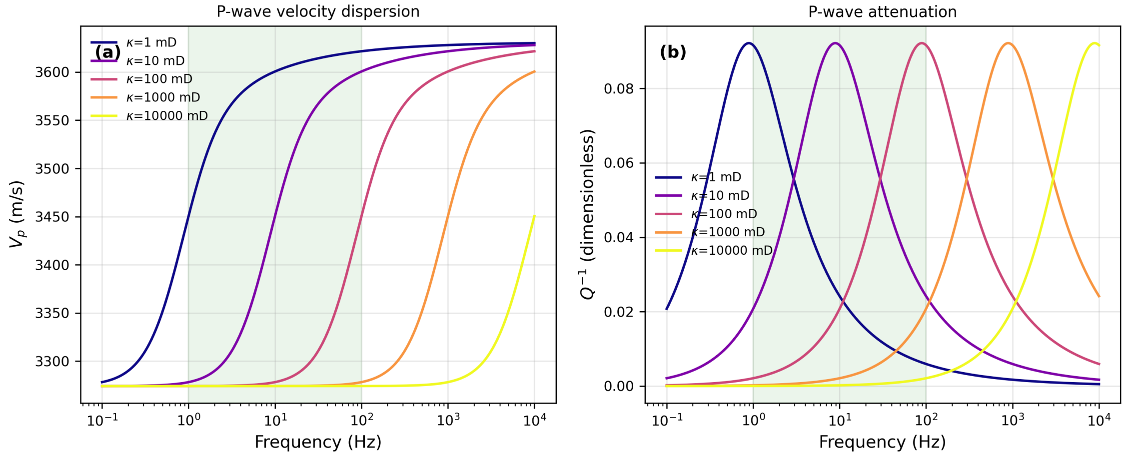

Figure C.1: V_P and Q^{-1} are shown against frequency for permeability from 1 mD to 10 000 mD at $S_{H_2} = 10\%$, $a = 5$ cm, and a constant frame. (a) V_P. (b) Q^{-1}.

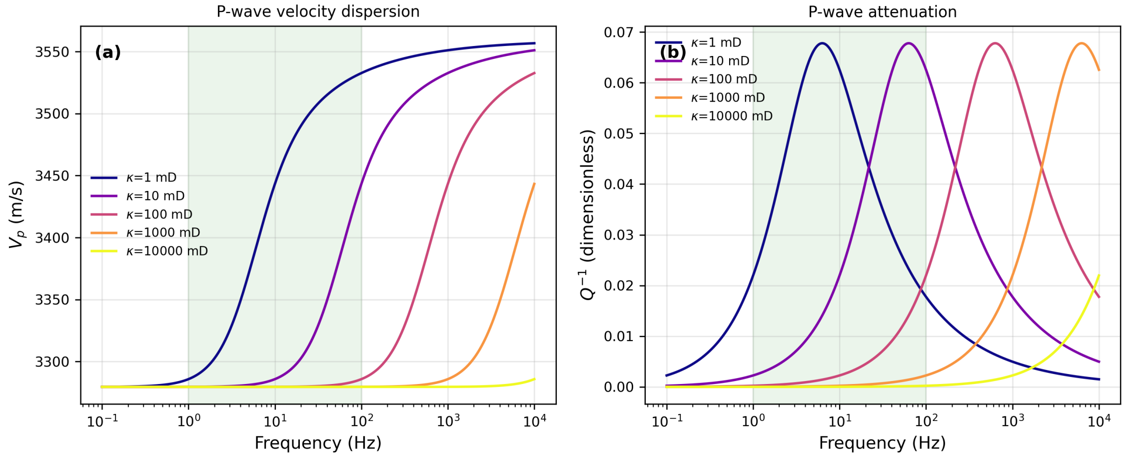

Figure C.2: V_P and Q^{-1} are shown against frequency for permeability from 1 mD to 10 000 mD at $S_{H_2} = 30\%$, $a = 5$ cm, and a constant frame. (a) V_P. (b) Q^{-1}. The $\kappa = 10\,000$ mD peak lies above the frequency axis.

Appendix D

Exact Reference Solutions for Simple Geometries

This appendix records the exact reference-solution structure used to benchmark the Johnson model in simple geometries.

Symbols follow Appendix C; in addition, σ_{33} is the applied axial stress, $\langle e_{zz} \rangle$ the period-averaged axial strain, R_b the outer radius of the concentric-sphere cell, $u(R_b)$ the radial displacement at that radius, and P_e the applied effective stress.

D.1 Periodic slabs

The periodic-slab reference solution follows the Johnson (2001, Appendix B) construction for two alternating fully saturated layers. The implementation solves the quasistatic system in each layer subject to continuity of displacement, relative displacement, and pore pressure at the interface. We then recover the effective bulk modulus from the average strain over one period,

$$K(\omega) = -\frac{\sigma_{33}}{\langle e_{zz} \rangle} - \frac{4}{3} G_{\text{dry}}, \tag{D.1}$$

with the stress choice used in the implementation fixed consistently across the cell solve.

D.2 Concentric spheres

The concentric-sphere reference solution follows the Johnson (2001, Appendix B) construction for an inner sphere and an outer shell. The implementation enforces continuity of radial displacement, relative displacement, pore pressure, and radial stress at the inner interface, together with the sealed outer boundary and applied stress condition used in the Appendix B derivation. We then recover the effective bulk modulus from the outer radial displacement,

$$K(\omega) = -\frac{R_b P_e}{3u(R_b)}. \tag{D.2}$$

D.3 Implementation support note

The concentric-sphere reference solution in the code uses a scaled spherical Hankel basis in the shell region. This is an algebraically equivalent reformulation of the original Johnson (2001, Appendix B) system, and we use it only to stabilise the numerical solve. It does not change the governing benchmark problem being solved.

Appendix E

White Concentric-Sphere Benchmark Formulas

This appendix collects the support formulas used in the White concentric-sphere benchmark implementation. The code follows the compact notation of Carcione et al. (2003).

Symbols follow Appendix C; in addition, K_j is the single-fluid Gassmann modulus of region j, $K_{A,j}$ and $K_{E,j}$ the storage and diffusion moduli, γ_j the diffusion wavenumber, Z_j the impedance factor of region j, S_1 the inner-sphere volume fraction, and a and b the inner and outer radii.

E.1 Region-wise Gassmann and storage moduli

For each region,

$$K_j = K_{\text{sat}}(K_{f,j}), \tag{E.1}$$

and the storage modulus is

$$K_{A,j}^{-1} = \frac{\phi}{K_{f,j}} + \frac{1-\phi}{K_0} - \frac{K_{\text{dry}}}{K_0^2}. \tag{E.2}$$

The diffusion modulus is

$$K_{E,j} = K_{A,j}\left[1 - \frac{K_{f,j}\,(1 - K_j/K_0)\,(1 - K_{\text{dry}}/K_0)}{\phi K_j\,(1 - K_{f,j}/K_0)}\right]. \tag{E.3}$$

E.2 Coupling coefficients and high-frequency limit

Let S_1 be the inner-sphere volume fraction. The coupling denominator in the implementation is

$$D_R = K_2\,(3K_1 + 4G_{\text{dry}}) + 4G_{\text{dry}}(K_1 - K_2)S_1. \tag{E.4}$$

The coupling coefficients are

$$R_1 = \frac{(K_1 - K_{\text{dry}})(3K_2 + 4G_{\text{dry}})}{D_R}, \qquad R_2 = \frac{(K_2 - K_{\text{dry}})(3K_1 + 4G_{\text{dry}})}{D_R}, \tag{E.5}$$

and the high-frequency limit is

$$K_\infty = \frac{D_R}{(3K_1 + 4G_{\text{dry}}) - 3(K_1 - K_2)S_1}. \tag{E.6}$$

E.3 Diffusion wavenumbers and impedance factors

The diffusion wavenumbers are

$$\gamma_j = \sqrt{\frac{i\omega\eta_j}{\kappa K_{E,j}}}. \tag{E.7}$$

With inner radius a and outer radius b, the impedance factors used in the implementation are

$$Z_1 = \frac{1 - e^{-2\gamma_1 a}}{(\gamma_1 a - 1) + (\gamma_1 a + 1)e^{-2\gamma_1 a}}, \tag{E.8}$$

$$Z_2 = \frac{(\gamma_2 b + 1)e^{-2\gamma_2 (b-a)} + (\gamma_2 b - 1)}{(\gamma_2 b + 1)(\gamma_2 a - 1)e^{-2\gamma_2 (b-a)} - (\gamma_2 b - 1)(\gamma_2 a + 1)}. \tag{E.9}$$

E.4 Transfer function and bulk modulus

The transfer function is

$$W = \frac{3ia\kappa(R_1 - R_2)\left(K_{A,1}/K_1 - K_{A,2}/K_2\right)}{b^3\omega\left(\eta_1 Z_1 - \eta_2 Z_2\right)}. \tag{E.10}$$

The final White benchmark modulus is then

$$K(\omega) = \frac{K_\infty}{1 - K_\infty W}. \tag{E.11}$$

E.5 Scope note

The classical White benchmark corresponds to an inner gas sphere and an outer brine shell. The current code also supports a phase-swapped extension for comparison with the gas-outer-shell Johnson case. That phase-swapped branch is exploratory, so we do not describe it as the classical White benchmark in the main text.

Appendix F

Additional Saturation Scenarios

This appendix documents the three saturation topologies not carried in the body of Chapter 5, together with the layer-count convergence test for the discretised scenarios. The modelling workflow, the three acoustic states, and the two carrying differences are exactly those of Section 5.4; only the saturation topology changes. Except where a table is explicitly labelled as a layered-overburden counterpart, all amplitude differences are migrated top-reservoir reflector values in the same arbitrary but internally consistent units as the Chapter 5 results tabulated below, under the simplified single-caprock overburden. TmB is the all-brine-to-tuned difference and TmR is the relaxed-to-tuned difference.

Table F.1: Effective coherently-stacked trace count at the top-reservoir sample, $N_{\mathrm{eff}} = A_{\text{full-stack}}/A_{p=0}$, the ratio of the full-stack migrated amplitude to the single-ray-parameter ($p = 0$) migrated image amplitude at the same sample, per case and applied frequency, under the single-block and layered Dutch overburden treatments. Each stack is formed from 601 ray-parameter traces; the departure of N_{eff} from that constant, and its variation across cases and frequencies, shows that the migration stack does not act as a constant multiplicative gain.

Case	f (Hz)	N_{eff} (single-block)	N_{eff} (layered)
Control	15	257.5	263.4
Control	30	223.9	239.6
Control	50	149.4	179.2
Control	100	60.2	94.8
Rotliegend-style	15	153.8	127.9
Rotliegend-style	30	99.4	106.2
Rotliegend-style	50	54.8	73.9
Rotliegend-style	100	32.2	36.5
Bunter-style	15	139.3	151.0
Bunter-style	30	135.7	136.4
Bunter-style	50	100.0	99.9
Bunter-style	100	51.8	59.1

F.1 Chapter 5 deterministic scenario results

The tables in this section collect the full numerical Chapter 5 results: migrated amplitude differences for the uniform and gravity-segregated deterministic scenarios, the reservoir-only two-way travel-time shifts, the overburden-treatment comparison, and the random-patchiness ensemble summary. The corresponding result figures and interpretation remain in Chapter 5; the layer-cake construction schematic is retained here as supporting model documentation. Model-column schematics for the three cases and the migrated images behind the principal tables are collected here as well.

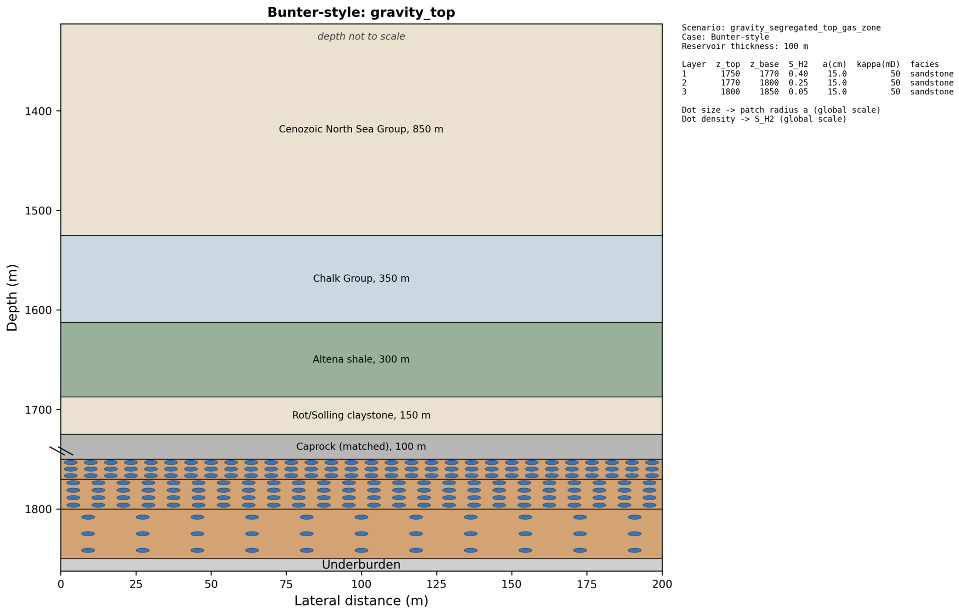

Figure F.1: Depth, stratigraphic-unit thickness, and reservoir hydrogen saturation are shown for the Bunter-style gravity-segregated case. (left) Layered Dutch overburden and reservoir at a true depth of 1750 m. (right) Reservoir-layer depths, saturations, patch radii, permeabilities, and facies. Glyph size encodes patch radius and glyph density encodes saturation; overburden depth is compressed.

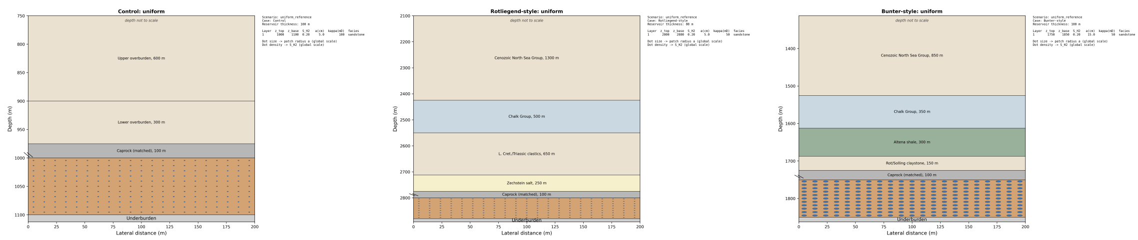

(a) (b) (c)

Figure F.2: Layered-overburden model columns for the uniform scenario. Each column spans the surface to the true reservoir depth of its case and terminates in the matched caprock immediately above the reservoir; overburden depth is compressed for display. (a) Control at 1000 m. (b) Rotliegend-style at 2800 m. (c) Bunter-style at 1750 m.

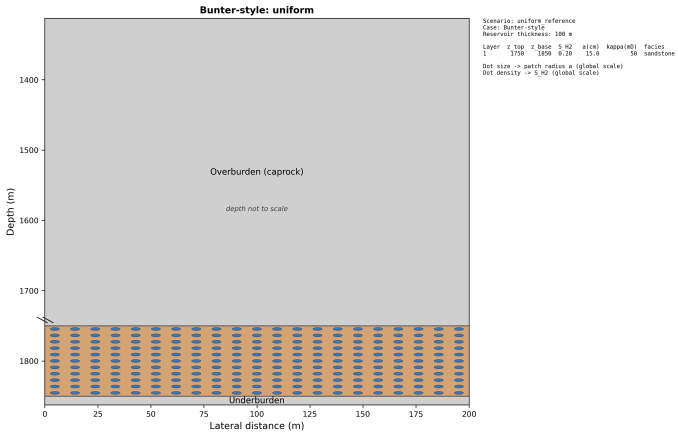

Figure F.3: Single-block counterpart of the Bunter-style column in Figure F.2: a single homogeneous block at the matched-caprock properties spans the surface to the reservoir top, so the top-reservoir reflection coefficient is identical between the two overburden treatments. Overburden depth is compressed for display.

Table F.2: Migrated top-reservoir reflector amplitude differences for the uniform scenario under the simplified single-block overburden, at true reservoir depth. TmB is the all-brine-to-tuned difference (monitoring signal); TmR is the relaxed-to-tuned difference (frequency-dependent part). Each is followed by its value as a percentage of the absolute all-brine reference amplitude at the same top-reservoir sample. The final column is |TmR/TmB| as a percentage. Absolute values are in arbitrary but internally consistent migrated-amplitude units.

Case	f (Hz)	TmB	TmB (%)	TmR	TmR (%)	|TmR/TmB| (%)
Control	15	−31.0	−79.9	0.0650	0.2	0.2
Control	30	−24.1	−72.5	0.192	0.6	0.8
Control	50	−12.7	−57.5	0.290	1.3	2.3
Control	100	0.0345	0.4	0.0461	0.5	133.5
Rotliegend-style	15	−8.02	−23.4	0.0144	0.0	0.2
Rotliegend-style	30	−3.17	−16.9	0.0201	0.1	0.6
Rotliegend-style	50	−0.722	−7.0	0.00931	0.1	1.3
Rotliegend-style	100	−1.21	−19.4	0.0955	1.5	7.9
Bunter-style	15	−7.64	−26.2	2.38	8.1	31.1
Bunter-style	30	−3.84	−14.3	2.75	10.3	71.6
Bunter-style	50	−1.50	−7.6	1.46	7.4	97.0
Bunter-style	100	−0.0780	−0.7	−0.318	−3.0	407.8

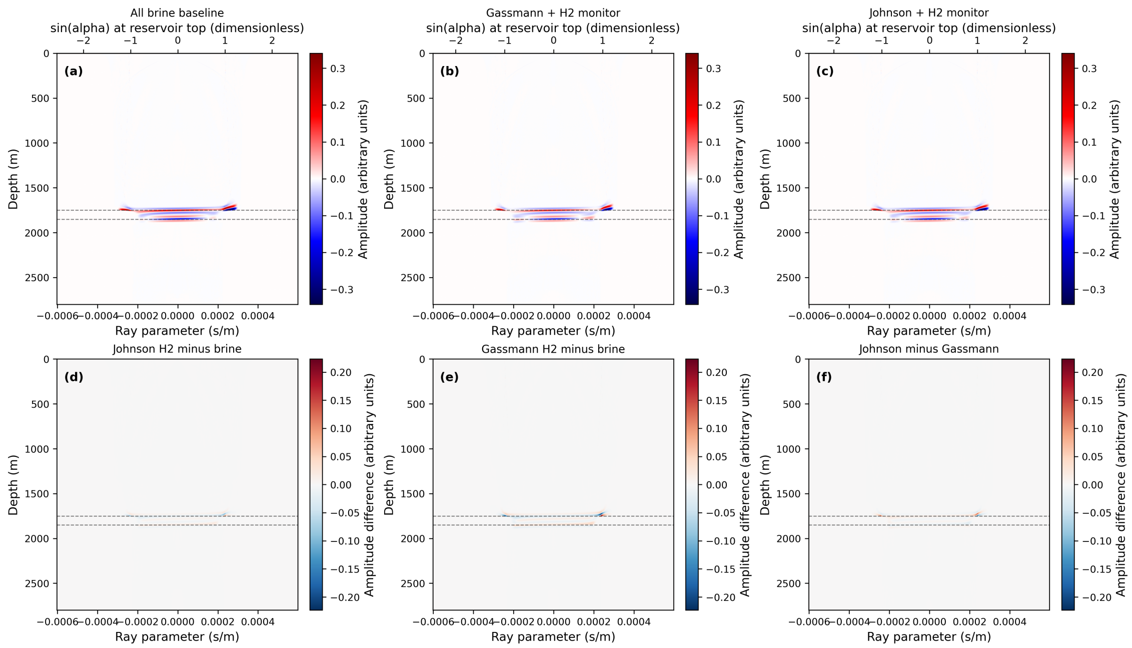

Figure F.4: Migrated angle-domain images for the Bunter-style uniform scenario under the single-block overburden at 30 Hz. The horizontal coordinate of each panel is the migrated ray-parameter axis, expressed as the sine of the propagation angle, not lateral distance; the medium itself is laterally invariant. The panels show the all-brine, relaxed, and tuned states and their differences. The tabulated amplitude differences of Table F.2 correspond to the unnormalised sum over this axis, read at the top-reservoir reflector.

Table F.3: Two-way travel-time shift $\Delta t_{2\text{way}}$ through the reservoir interval, in milliseconds, between the all-brine state and the tuned (Johnson) state, for the uniform and gravity-segregated topologies at the four applied frequencies. Positive values indicate a slower reservoir after hydrogen injection, and so a push-down of underlying reflectors on a monitor survey. The quantity is an integral over the full reservoir interval and is independent of the overburden treatment by construction, so a single set of values applies to both the single-block and layered columns. For scale, Toh et al. (2025) adopt 0.5 ms as the minimum detectable shift for towed-streamer acquisition; values below that figure are set in italics.

		$\Delta t_{2\text{way}}$ (ms)			
Scenario	Case	15 Hz	30 Hz	50 Hz	100 Hz
Uniform	Control	6.59	6.55	6.44	5.99
	Rotliegend-style	1.58	1.57	1.54	1.43
	Bunter-style	1.17	0.69	*0.45*	*0.27*
Gravity-segregated	Control	5.89	5.12	4.44	3.80
	Rotliegend-style	1.37	1.21	1.05	0.86
	Bunter-style	0.76	0.58	*0.44*	*0.28*

Table F.4: Migrated top-reservoir reflector amplitude differences for the gravity-segregated scenario under the simplified single-block overburden, at true reservoir depth. Columns as in Table F.2.

Case	f (Hz)	TmB	TmB (%)	TmR	TmR (%)	|TmR/TmB| (%)
Control	15	−32.6	−84.1	−0.120	−0.3	0.4
Control	30	−25.9	−78.1	−0.318	−1.0	1.2
Control	50	−14.0	−63.3	0.0529	0.2	0.4
Control	100	−1.02	−11.0	−0.219	−2.4	21.5
Rotliegend-style	15	−8.21	−24.0	0.134	0.4	1.6
Rotliegend-style	30	−4.08	−21.7	−0.146	−0.8	3.6
Rotliegend-style	50	−1.19	−11.5	−0.0294	−0.3	2.5
Rotliegend-style	100	−1.24	−19.9	0.131	2.1	10.5
Bunter-style	15	−9.84	−33.7	1.61	5.5	16.4
Bunter-style	30	−7.49	−27.9	0.423	1.6	5.6
Bunter-style	50	−3.59	−18.1	0.0507	0.3	1.4
Bunter-style	100	0.160	1.5	0.241	2.3	150.9

Table F.5: Frequency-dependent part of the migrated monitoring signal (TmR, relaxed to tuned, at the top reflector) under the single-block overburden and the layered Dutch overburden, per case and applied frequency, at true reservoir depth. Each absolute column is followed by its value as a percentage of the absolute all-brine reference amplitude at the same top-reservoir sample.

		Uniform				Gravity-segregated			
Case	f (Hz)	Single	Single (%)	Layered	Layered (%)	Single	Single (%)	Layered	Layered (%)
Control	15	0.0650	0.2	0.0629	0.2	−0.120	−0.3	−0.175	−0.5
Control	30	0.192	0.6	0.200	0.6	−0.318	−1.0	−0.571	−1.7
Control	50	0.290	1.3	0.371	1.4	0.0529	0.2	−0.284	−1.1
Control	100	0.0461	0.5	0.434	3.2	−0.219	−2.4	−0.379	−2.8
Rotliegend-style	15	0.0144	0.0	0.00882	0.0	0.134	0.4	0.0370	0.1
Rotliegend-style	30	0.0201	0.1	0.0202	0.1	−0.146	−0.8	−0.171	−1.0
Rotliegend-style	50	0.00931	0.1	0.0341	0.3	−0.0294	−0.3	−0.0492	−0.4
Rotliegend-style	100	0.0955	1.5	0.0745	1.2	0.131	2.1	0.141	2.3
Bunter-style	15	2.38	8.1	2.21	7.6	1.61	5.5	1.54	5.3
Bunter-style	30	2.75	10.3	2.58	10.4	0.423	1.6	0.375	1.5
Bunter-style	50	1.46	7.4	1.30	7.1	0.0507	0.3	−0.0723	−0.4
Bunter-style	100	−0.318	−3.0	−0.270	−2.4	0.241	2.3	0.194	1.7

Table F.6: Ensemble random-patchiness scenario under the single-block (simplified) overburden, at true reservoir depth, complete over ten seeds. Across-seed mean and range across realisations, reported as the tenth to ninetieth percentiles of the monitoring (TmB) and frequency-dependent (TmR) migrated top-reservoir reflector amplitude differences, per case and applied frequency. Same units as Table F.2.

Case	f (Hz)	TmB mean	TmB [P10, P90]	TmR mean	TmR [P10, P90]
Control	15	−30.01	[−31.37, −27.48]	1.67	[0.38, 3.71]
Control	30	−21.88	[−24.49, −17.23]	2.88	[0.04, 7.30]
Control	50	−12.81	[−14.02, −11.65]	0.93	[−0.45, 2.22]
Control	100	−1.54	[−4.81, 1.79]	−0.81	[−4.26, 2.22]
Rotliegend-style	15	−8.48	[−8.85, −7.87]	0.34	[0.08, 0.57]
Rotliegend-style	30	−3.09	[−3.62, −2.33]	0.37	[−0.08, 0.96]
Rotliegend-style	50	−0.84	[−1.03, −0.60]	0.08	[−0.04, 0.21]
Rotliegend-style	100	−1.00	[−1.31, −0.73]	0.26	[0.02, 0.46]
Bunter-style	15	−10.60	[−11.32, −9.19]	0.65	[0.25, 1.20]
Bunter-style	30	−6.12	[−6.85, −4.70]	0.98	[−0.01, 2.00]
Bunter-style	50	−2.92	[−3.45, −2.40]	0.37	[−0.20, 0.78]
Bunter-style	100	−0.09	[−1.36, 0.98]	−0.12	[−1.26, 0.69]

F.2 Capillary transition zone

The capillary transition topology imposes a smooth downward decay of hydrogen saturation, representing a capillary-controlled transition between a gas-charged upper interval and a brine-dominated base. It is the smoothest of the topologies and is therefore used for the convergence test of Section F.5. Its migrated response, Table F.7, follows the same case ordering as the body scenarios: a few percent frequency-dependent correction for Control and Rotliegend-style, and a substantial correction for Bunter-style.

Table F.7: Capillary transition scenario. Columns as in Table F.2.

Case	f (Hz)	TmB	TmR	$\lvert$TmR/TmB$\rvert$ (%)
Control	15	−31.5	−0.0193	0.1
Control	30	−25.2	−0.228	0.9
Control	50	−13.5	0.110	0.8
Control	100	−0.0997	0.379	380.0
Rotliegend-style	15	−7.90	0.0526	0.7
Rotliegend-style	30	−3.52	−0.0510	1.4
Rotliegend-style	50	−0.936	−0.0250	2.7
Rotliegend-style	100	−1.14	0.207	18.0
Bunter-style	15	−6.41	3.92	61.1
Bunter-style	30	−5.16	2.05	39.8
Bunter-style	50	−2.52	0.764	30.4
Bunter-style	100	−0.202	−0.240	119.2

F.3 Baffled column

The baffled topology introduces a low-permeability interval that discontinuously separates two storage zones, representing a flow baffle within the reservoir. Its migrated response is given in Table F.8. The case ordering is preserved; the internal contrast at the baffle additionally produces a deeper reflector whose amplitude difference can carry the opposite sign to the top reflector, which is why only the top-reservoir reflector is reported here as the consistent monitoring quantity.

Table F.8: Baffled column scenario. Columns as in Table F.2.

Case	f (Hz)	TmB	TmR	|TmR/TmB| (%)
Control	15	−28.1	−0.0496	0.2
Control	30	−22.2	−0.241	1.1
Control	50	−14.2	0.234	1.7
Control	100	2.43	−0.238	9.8
Rotliegend-style	15	−7.37	0.134	1.8
Rotliegend-style	30	−3.01	−0.142	4.7
Rotliegend-style	50	−0.842	−0.0183	2.2
Rotliegend-style	100	−0.898	0.156	17.4
Bunter-style	15	−8.39	1.37	16.3
Bunter-style	30	−4.40	1.39	31.7
Bunter-style	50	−1.93	1.40	72.7
Bunter-style	100	0.152	−0.757	498.4

F.4 Rotliegend-style continuous column

The Rotliegend-style continuous topology represents a laterally continuous reservoir body with moderate vertical saturation variation. Its migrated response, Table F.9, again preserves the case ordering established in the body.

Table F.9: Rotliegend-style continuous column scenario. Columns as in Table F.2.

Case	f (Hz)	TmB	TmR	|TmR/TmB| (%)
Control	15	−31.5	0.00922	0.0
Control	30	−24.6	0.0239	0.1
Control	50	−13.2	0.0757	0.6
Control	100	−0.349	−0.131	37.6
Rotliegend-style	15	−8.18	0.0111	0.1
Rotliegend-style	30	−3.38	0.000104	0.0
Rotliegend-style	50	−0.853	−0.00354	0.4
Rotliegend-style	100	−1.28	0.0600	4.7
Bunter-style	15	−8.46	1.99	23.6
Bunter-style	30	−5.05	1.89	37.4
Bunter-style	50	−2.08	1.04	50.2
Bunter-style	100	−0.0377	−0.174	460.7

F.5 Layer-count convergence

The continuous saturation profiles of the capillary and gravity-segregated scenarios are represented internally by a finite number of constant-saturation layers. The capillary transition is the smoothest profile and therefore the most demanding to discretise, so it was used to test the sensitivity of the modelled response to the layer count. The top-reservoir impedance change at 30 Hz was evaluated for layer counts of 5, 10, 20, 40 and 80, and the layer count was selected as the first for which doubling it changed the result by less than one percent. Doubling from five to ten layers changes $\Delta Z/Z$ by 1.45 percent, and from ten to twenty by 0.75 percent, so ten layers is used for the capillary results of Section F.2. Further refinement to eighty layers changes the result by a further 1.3 percent in total, so the discretisation is a small but not strictly negligible contribution to the absolute value; it does not affect the comparisons drawn between scenarios, which all use the same discretisation.

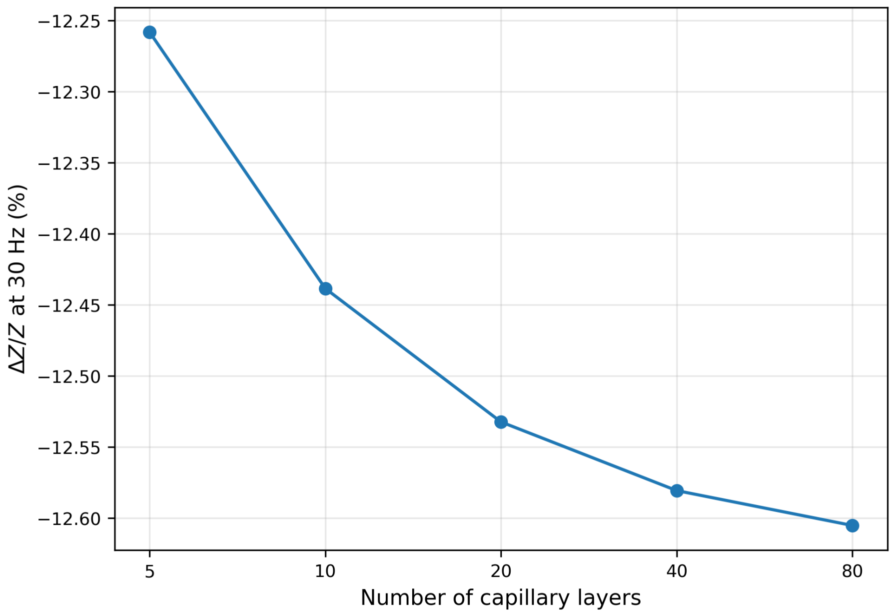

Figure F.5: Migrated top-reservoir reflector amplitude is shown against the number of constant-saturation layers used to discretise the capillary transition.

F.6 Layered-overburden amplitude differences

Chapter 5 interprets the layered-overburden result through the relaxed-to-tuned difference, placed beside the single-block value in Table F.5. For completeness, the full migrated amplitude differences under the layered Dutch overburden are tabulated here for the two primary scenarios in Tables F.10 and F.11, in the same units and with the same columns as the single-block tables collected above. The all-brine-to-tuned signal TmB carries the overburden modulation discussed in Section 5.5.4; the relaxed-to-tuned signal TmR is close to its single-block value for the Bunter-style case, though not for the two cases in which it is small, for the reasons set out there.

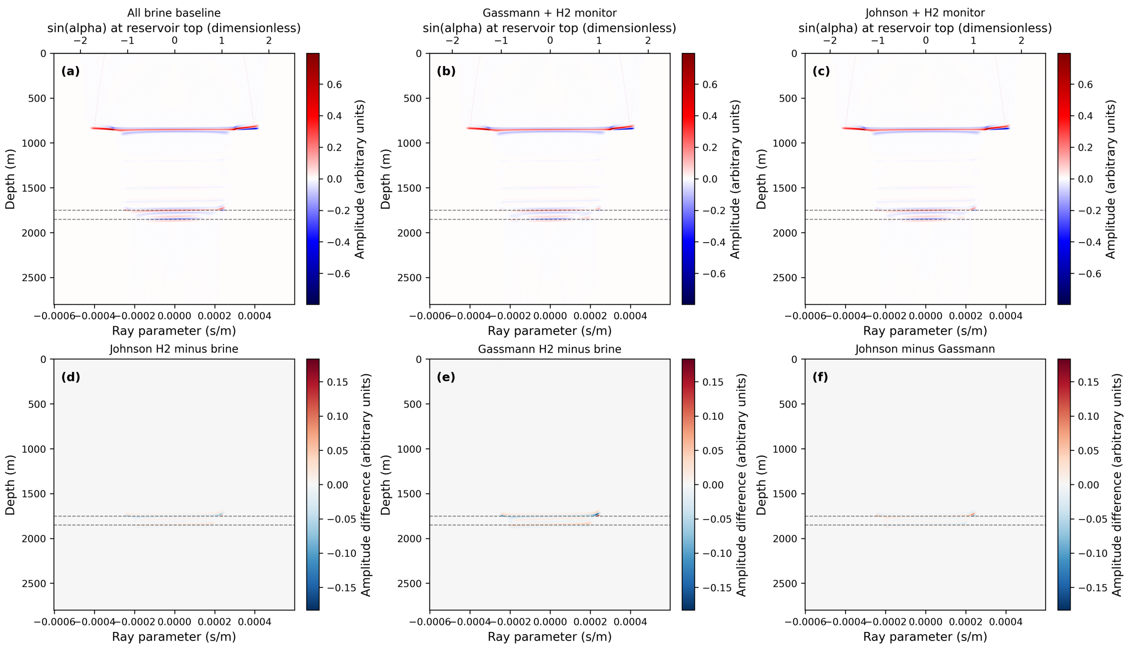

Figure F.6: Migrated angle-domain images for the Bunter-style uniform scenario under the layered Dutch overburden at 30 Hz. The horizontal coordinate of each panel is the migrated ray-parameter axis, expressed as the sine of the propagation angle, not lateral distance; the medium itself is laterally invariant. The panels show the all-brine, relaxed, and tuned states and their differences. The corresponding tabulated amplitude differences of Table F.10 correspond to the unnormalised sum over this axis, read at the top-reservoir reflector.

Table F.10: Uniform scenario under the layered Dutch overburden, at true reservoir depth. Columns as in Table F.2.

Case	f (Hz)	TmB	TmR	|TmR/TmB| (%)
Control	15	−30.9	0.0629	0.2
Control	30	−25.7	0.200	0.8
Control	50	−16.6	0.371	2.2
Control	100	−4.50	0.434	9.6
Rotliegend-style	15	−5.07	0.00882	0.2
Rotliegend-style	30	−2.98	0.0202	0.7
Rotliegend-style	50	−1.84	0.0341	1.9
Rotliegend-style	100	−1.01	0.0745	7.3
Bunter-style	15	−6.92	2.21	31.9
Bunter-style	30	−3.62	2.58	71.3
Bunter-style	50	−1.65	1.30	78.8
Bunter-style	100	−0.264	−0.270	102.6

Table F.11: Gravity-segregated scenario under the layered Dutch overburden, at true reservoir depth. Columns as in Table F.2.

Case	f (Hz)	TmB	TmR	|TmR/TmB| (%)
Control	15	−32.5	−0.175	0.5
Control	30	−27.9	−0.571	2.0
Control	50	−18.4	−0.284	1.5
Control	100	−6.12	−0.379	6.2
Rotliegend-style	15	−5.37	0.0370	0.7
Rotliegend-style	30	−3.90	−0.171	4.4
Rotliegend-style	50	−2.38	−0.0492	2.1
Rotliegend-style	100	−1.11	0.141	12.7
Bunter-style	15	−8.77	1.54	17.5
Bunter-style	30	−7.12	0.375	5.3
Bunter-style	50	−3.96	−0.0723	1.8
Bunter-style	100	−0.263	0.194	74.1

F.7 Ensemble per-seed spread

The single-block ensemble band of Table F.6 summarises ten random realisations by their mean and percentile band. The full per-seed spread for the Bunter-style case, the case on which the in-band correction turns, is given in Table F.12 so the distribution behind the band is visible. Values are the frequency-dependent difference TmR for each seed.

Table F.12: Per-seed frequency-dependent migrated reflector difference (TmR) for the Bunter-style case under the single-block overburden, ten seeds, at true reservoir depth. Same units as Table F.6.

Seed	TmR (15 Hz)	TmR (30 Hz)	TmR (50 Hz)	TmR (100 Hz)
0	0.471	−0.151	0.00806	0.203
1	0.824	1.29	0.488	−1.22
2	0.251	0.414	−0.195	0.608
3	0.203	0.384	−0.195	0.458
4	0.904	1.01	0.160	0.672
5	0.418	1.65	1.13	−0.940
6	0.286	0.00515	0.296	0.327
7	1.17	2.37	0.692	−0.606
8	1.50	1.96	0.555	−1.59
9	0.478	0.922	0.738	0.889
mean	0.651	0.984	0.367	−0.120

Appendix G

Use of AI Tools

I used ChatGPT and Codex (OpenAI) and Claude (Anthropic) for writing support (structure, phrasing, condensing, and LaTeX formatting), as coding assistants (prototype scripts, troubleshooting, and improving readability), and for consistency checking across the document and its references. I reviewed, modified, and validated all generated text and code. All modelling assumptions, final implementations, simulations, figures, and interpretation are my own.