نشریه ژئومکانیک و ژئوانرژی (ژئومکانیک نفت)

GEOMECHANICS and GEOENERGY JOURNAL (GGJ)

Original Article

A closed-form screening criterion for stress-dependent permeability effects on waterflood performance

Behnam Ebrahimy*[1], Ehsan Taheri[1]

1- Faculty of Engineering, Tarbiat Modares University, Tehran, Iran

Received: 2026/01/21; Accepted: 2026/03/19
DOI: 10.22107/ggj.2026.594380.1268

Keywords

stress-dependent permeability, hydromechanical coupling, waterflooding, injectivity, reservoir screening

Abstract

Coupled flow–geomechanics simulation is expensive, and it needs to be known whether it changes the forecast. For single-phase flow to a single well the answer has been known since the 1980s: an exponential dependence of permeability on effective stress reduces the problem to one dimensionless group. Whether this extends to two-phase displacement in heterogeneous rock has not been examined. This is tested on a field-scale quarter five-spot, in which a stress-dependent permeability is imposed through a uniaxial-strain poroelastic closure and swept over one homogeneous and four heterogeneous log-normal fields, in 104 simulations. At issue is not the constitutive law, assumed alike by simulation and scaling, but whether the geometric cancellation behind the single-well result is preserved under heterogeneity and a second mobile phase. It is shown to be preserved. The ratio of coupled to uncoupled injector-to-producer pressure difference is collapsed onto a closed-form curve with a median error below one per cent, even though the uncoupled difference itself is shifted by a factor of four across the heterogeneity levels tested; the collapse is insensitive to grid resolution, the pressure solver, and whether permeability is tied to effective stress directly or through porosity. The displacement measures are barely affected: although permeability is varied by an order of magnitude and injection pressure is reduced by up to three quarters, breakthrough time and recovery factor shift by less than one per cent. Under a prescribed rate, the pressure field is rescaled by the coupling without the flow being redirected, and the same dimensionless group is shown to govern a well held at fixed pressure, where the coupling appears as a gain in injectivity rather than as a drop in pressure. Inverting the scaling yields a screening rule that costs nothing to apply and determines whether a coupled simulation is required.

1. Introduction

1.1. Hydromechanical coupling in reservoir rock

Rock deforms when the fluid pressure inside it changes, and the deformation feeds back on the flow. Terzaghi's effective-stress principle and Biot's theory of consolidation quantified that statement eight decades ago [1, 2], and the consequences are now familiar across rock engineering: convergence around a drained tunnel and the seabed subsidence that follows depletion of a weak sandstone reservoir belong to the same class of problem [3, 4]. In petroleum, the feedback that matters most for production forecasting is the dependence of permeability on effective stress. Laboratory measurements on sandstones and carbonates show that permeability falls as the effective confining stress rises, often by more than an order of magnitude over the stress range spanned by a producing field, and that the loss is only partly recovered on unloading [5, 6]. The same mechanism, running in reverse, raises permeability where injection lifts the pore pressure and unloads the rock skeleton.

Whether this feedback needs representing explicitly is a practical question, because doing so is expensive. A fully coupled simulation solves the flow and the mechanical equilibrium equations simultaneously and is unconditionally stable but costly to assemble and precondition; the various sequential schemes trade some of that cost for a stability condition that is not always benign

* **Corresponding Author:** behnam_ebrahimy@modares.ac.ir

[7-9]. Comparative studies of the coupling strategies have been available for two decades [10, 11], and multiscale formulations have since been developed specifically to reduce that cost [12-14]. None of this addresses the question an engineer faces before choosing a simulator: does the coupling change the answer enough to be worth modelling?

1.2. What is already known, and what is missing

For single-phase flow to a single well the question has a clean answer. Pedrosa [15] showed that an exponential permeability–stress law linearises the equation, reducing the problem to one dimensionless group formed from the permeability modulus and drawdown. Kikani and Pedrosa [16] developed the perturbation solution, and Chin and co-workers [17, 18] confirmed against fully coupled simulations that well responses and productivity indices in stress-sensitive formations follow the predicted scaling. It has not, to our knowledge, been carried across to the problem that actually governs field development: two-phase displacement between wells in heterogeneous rock. Several things could break the single-well argument there. The pressure field around an injector in a five-spot is not radial once the permeability is heterogeneous; the mobility ratio makes the effective transmissibility a function of saturation as well as stress; and the quantity an engineer cares about is not only the pressure at the well but the sweep between wells. It is not obvious in advance whether stress sensitivity mainly changes how hard it is to inject, or how well the injected water displaces the oil, and the two have very different practical consequences.

1.3. Scope of the present work

This paper treats that question numerically. We take a two-phase quarter five-spot at field scale, give the rock a stress-dependent permeability through a uniaxial-strain poroelastic closure, and sweep the stress sensitivity over one homogeneous field and four levels of log-normal heterogeneity, three realisations at each level. The performance measures are the injector-to-producer pressure difference at a fixed point of the flood, the time to water breakthrough, and the recovery factor.
It is worth being precise about what such a test can settle. The exponential permeability law is written into the simulations and into the scaling argument alike, so no amount of agreement between them says anything about whether that law describes real rock. What is genuinely open is narrower and more useful: the single-well result rests on a geometric cancellation, and whether that cancellation holds once the medium is heterogeneous, the flow no longer radial, and a second phase mobile is a question the algebra cannot answer.
Two results follow, and they point in opposite directions. The injection pressure obeys the single-well scaling closely: over the whole sweep, including fields with $\sigma_{\ln k}$=2, the ratio of coupled to uncoupled pressure difference collapses onto a closed-form curve with a median departure below 1%, so a criterion derived for radial single-phase flow survives translation to a heterogeneous two-phase displacement. The displacement measures, by contrast, do not respond at all. Breakthrough time and recovery factor move by less than 1 per cent everywhere in the sweep, including where the permeability field has been stretched by an order of magnitude and the injection pressure has fallen by three quarters. Under a prescribed rate the coupling rescales the pressure field without redirecting the flow.
Taken together these give a screening rule that costs nothing to apply: form the dimensionless group from an ordinary uncoupled run, read the expected error off a curve, and let the size of that error decide whether a coupled simulation is justified. The same group serves both ways an injector is operated, on rate or on pressure, a symmetric pair, so the rule applies to both. Section 2 sets out the formulation and the numerical model, together with the verification behind them, and Section 3 presents the sweep and the resulting criterion.

2. Methodology

2.1. Flow model

Two immiscible, incompressible phases occupy the pore space. Capillary and gravitational forces are neglected, so both phases share a single pressure and Darcy's law for phase $\alpha \in o,w$ reads

$$\boldsymbol{u}_\alpha = -\frac{k\, k_{r\alpha}(S_w)}{\mu_\alpha} \nabla p \qquad (1)$$

with k the absolute permeability, $k_{r\alpha}$ the relative permeability and μ_α the phase viscosity. Summing the two phase balances gives the pressure equation,

$$\nabla \cdot (\lambda_t \nabla p) = -q_t, \qquad \lambda_t = k\left(\frac{k_{ro}}{\mu_o} + \frac{k_{rw}}{\mu_w}\right) \tag{2}$$

and the water balance closes the system,

$$\phi \frac{\partial S_w}{\partial t} + \nabla \cdot (f_w \boldsymbol{u}_t) = q_w \tag{3}$$

with the water fractional flow

$$f_w = \frac{k_{rw}/\mu_w}{k_{ro}/\mu_o + k_{rw}/\mu_w} \tag{4}$$

Relative permeabilities follow the Brooks–Corey model [19] with pore-size index λ=3. Equations (2) and (3) are advanced with an IMPES scheme: the pressure is solved implicitly on the fine grid, the total flux is reconstructed, and the saturation is updated explicitly with sub-stepping under a CFL limit of 0.25. That limit, and the pressure step that goes with it, were fixed by the mass balance rather than by taste; at looser settings the injector cell takes in more water in a step than its pore volume holds, the surplus is removed by the saturation cap, and the run quietly loses water.

The residual oil saturation is set to zero throughout, and the reason is numerical rather than petrophysical. The explicit update clips saturations to [S_{wc},,1-S_{or}]; with S_{or}>0 that ceiling sits low enough for the injector cell to reach it, and every subsequent step discards the surplus water, so a long run loses mass at the one cell where the mass balance has to be exact. Rather than tune the clip we removed the ceiling, at the price of a displacement that sweeps to a higher water saturation than a real waterflood would. The cost is confined to absolute levels: recovery factors quoted here run higher than a field would deliver and are comparators between runs rather than forecasts. The conclusions are not exposed to it, because every number in this paper is a ratio of two runs sharing the same S_{or}, and a bias common to both divides out. Restoring a physical S_{or} needs a saturation update that redistributes the clipped volume instead of discarding it, which is a change to the solver rather than to this study.

2.2. Poroelastic closure

The reservoir is taken to be laterally extensive and confined, so that lateral strain vanishes and the total vertical stress is held constant by the overburden. These are the conditions under which the poroelastic problem reduces to the one-dimensional compaction model of Geertsma [20], and they are the standard basis for screening estimates of compaction and subsidence. Under them no mechanical system needs to be solved: the stress and strain follow algebraically from the pressure.

With compression counted positive, the Biot effective vertical stress and its change at constant total stress are

$$\sigma_v' = \sigma_v - \alpha p, \qquad \Delta\sigma_v' = -\alpha\,\Delta p \tag{5}$$

where α is the Biot coefficient and $\Delta p = p - p_0$ is measured from the initial pressure. The volumetric strain under uniaxial-strain conditions is

$$\varepsilon_v = \frac{\Delta\sigma_v'}{K_{\text{oed}}}, \qquad K_{\text{oed}} = \frac{E\,(1-\nu)}{(1+\nu)(1-2\nu)} \tag{6}$$

with E Young's modulus, ν Poisson's ratio and K_{oed} the constrained, or oedometric, modulus. Porosity follows the finite-strain form

$$\phi = 1 - (1-\phi_0)\exp(\varepsilon_v) \tag{7}$$

which is not an empirical fit but the integral of the solid-phase mass balance, and is the relation used in the multiscale geomechanical model of Taheri and co-workers [13, 14, 21]. It reduces to $\varphi \simeq \varphi_0 - (1-\varphi_0)\varepsilon_v$ for small strain. For permeability we use the exponential law fitted to core data by Davies and Davies [5],

$$k = k_0 \exp(-\gamma\,\Delta\sigma_v') = k_0 \exp(\gamma\,\alpha\,\Delta p) \tag{8}$$

in which γ is the permeability modulus. As an independent check we also use the Kozeny–Carman form driven by Eq. (7), which is the permeability–porosity relation adopted in the reference model itself [21],

$$\frac{k}{k_0} = \left(\frac{\phi}{\phi_0}\right)^3 \left(\frac{1-\phi_0}{1-\phi}\right)^2 \tag{9}$$

Because both fluids and the bulk rock volume are treated as incompressible, Eq. (2) carries no $\partial\varphi/\partial t$ storage term. Allowing the pore volume in Eq. (3) to move while that term is absent would create fluid mass, so the transport equation retains φ_0 and Eq. (7) is used for the strain and compaction diagnostics only. This is a real restriction on what the study can conclude, and Section 3.5 quantifies it.

2.3. The controlling group

Combining Eqs. (5) and (8), permeability depends on pressure alone, through the single dimensionless group

$$\Pi = \gamma\,\alpha\,\Delta p \tag{10}$$

which is the natural logarithm of the permeability contrast that the pressure field induces between two points separated by Δp.

The consequence for a rate-controlled well follows from the argument of Pedrosa [15]. For steady flow at prescribed total rate q through a region whose geometry contributes a shape factor G,

$$\frac{q\,\mu}{G} = \int_{p_{\mathrm{out}}}^{p_{\mathrm{in}}} k\,(p)\,\mathrm{d}p = \frac{k_0}{\gamma\alpha}\left[\exp(\gamma\alpha\,\Delta p_{\mathrm{c}}) - 1\right] \quad (11)$$

where the second equality uses $p_{\mathrm{out}}=p_0$, the pressure at which the rock carries its reference permeability k_0. That is how the model is posed here, the producer being held at the initial pressure, but it is a condition and not an identity: Δp in every expression below is measured from the stress state at which k_0 was determined, not from an arbitrary datum. An engineer forming the group for a real field has to be careful on this point, because a core permeability quoted at a laboratory confining stress and a producer sitting well below initial reservoir pressure do not share a reference, and the difference goes straight into the exponent.

The same rate through the same region with γ=0 requires $q\mu/G=k_0\Delta p_{\mathrm{u}}$. Eliminating $q\mu/G$ between the two gives $\Pi_{\mathrm{u}}=\exp(\Pi_{\mathrm{c}})-1$ and hence

$$\frac{\Delta p_{\mathrm{c}}}{\Delta p_{\mathrm{u}}} = \frac{\ln(1+\Pi_{\mathrm{u}})}{\Pi_{\mathrm{u}}}, \qquad \Pi_{\mathrm{u}} = \gamma\,\alpha\,\Delta p_{\mathrm{u}} \quad (12)$$

where subscripts c and u denote the coupled and uncoupled solutions. Two features of Eq. (12) matter here. The shape factor G cancels, which is why the relation is not tied to radial geometry; and the argument Π_{u} is built from the uncoupled pressure difference, so it can be evaluated before any coupled simulation is run. Whether the cancellation survives when the medium is heterogeneous and two phases are flowing is an empirical question, and it is the one the sweep in Section 3 answers.

2.4. Numerical model and case set-up

The flow solver is a Python code written module by module against the finite-volume kernel of the multiscale geomechanical model of Taheri and co-workers [12, 13]. It is not offered here as a reproduction of that model's published output, and no claim in this paper rests on such a match; what the results rest on is the internal evidence set out in Section 2.5, and on the fact that every quantity reported is a ratio between two runs of the same solver on the same grid, in which a systematic bias common to both cancels. Two pressure backends are available: a direct sparse solve on the fine grid, used for the sweep, and a multiscale finite-volume solver, used to confirm that the conclusions do not depend on the pressure solver. The poroelastic closure of Section 2.2 is evaluated at the start of every pressure step from the pressure carried over from the previous step, which is an explicit sequential coupling; with the mechanical problem reduced to Eqs. (5)–(8) there is no mechanical solve to iterate against, and no coupling iteration is needed.

The reference case is a quarter five-spot of 200 m side and 10 m thickness discretised into 55×55 cells (Fig. 1a). Water is injected at a fixed rate in the south-west corner and the north-east corner is held at the initial pressure. Rock and fluid properties are listed in Table 1. The injection rate is calibrated so that the initial injector-to-producer difference is 10 MPa, which places the reference case in the middle of the range encountered in field operations.

Heterogeneity is imposed as a spatially correlated log-normal permeability field with geometric mean k and prescribed $\sigma_{\ln k}$, generated by smoothing white noise with a Gaussian kernel of six-cell correlation length and rescaling so that $\sigma_{\ln k}$ is met exactly (Fig. 1b). Five levels, $\sigma_{\ln k}\in 0,0.5,1,1.5,2$, are used. The homogeneous level is a single field; each of the four heterogeneous levels carries three realisations, drawn from seeds 7, 23 and 91, giving thirteen permeability fields.

Table 1. Reference case properties. The residual oil saturation is zero for the numerical reason given in the text, and recovery factors should be read accordingly.

Property	Symbol	Value
Domain side	L	200 m
Thickness	h	10 m
Grid	–	55 × 55
Reference permeability	$\bar{k}$	50 mD
Reference porosity	ϕ_0	0.20
Oil viscosity	μ_o	5 mPa·s
Water viscosity	μ_u	1 mPa·s
Connate water	Swc	0.20
Residual oil	S_{or}	0
Injection rate	q	1.905×10^{-4} m³/s
Pressure step	Δt	5.25×10^{4} s
Young's modulus	E	0.8 GPa
Poisson's ratio	ν	0.20
Biot coefficient	α	0.9
Initial difference	Δp_0	10 MPa

Eight values of γ from 0 to 3e-7 /Pa are applied to every field, so the sweep is 104 runs in total, of which 91 are coupled. Because the seeds are fixed the whole sweep is reproducible run for run.
Runs are compared at a fixed point of the flood, 0.30 pore volumes injected, rather than at the end of the simulation. Breakthrough moves between cases, so the final state is not the same stage of the displacement in every run.

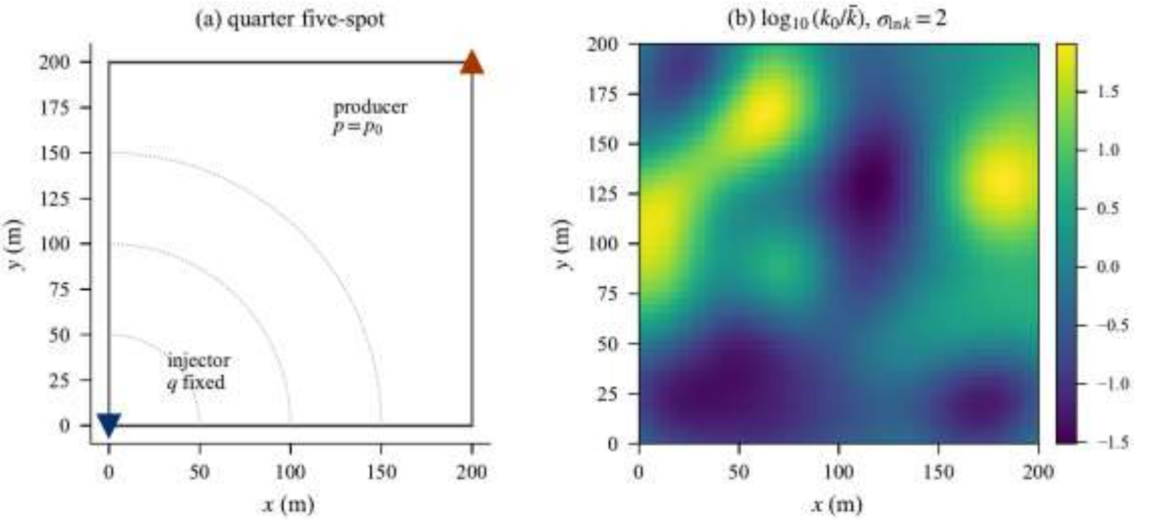

Fig. 1. Problem set-up. (a) Quarter five-spot with a rate-controlled injector and a pressure-controlled producer; dotted arcs mark equal distances from the injector. (b) One realisation of the log-normal permeability field at $\sigma_{\ln k}$=2, relative to the geometric mean

.

2.5. Verification

Six checks stand behind the results that follow. Five of them test the solver against itself; the last puts it against a solution that was known long before the code existed, and it is deliberately placed last because it is the one that could have failed on its own terms.

Zero-coupling reproduction. Setting $\gamma \to 0$ returns the uncoupled saturation field bit for bit, confirming that the closure is inert when it should be and that no part of the coupled path perturbs the solver by accident.

Mass balance. At the step and CFL limit adopted above, the injected volume and the water in place agree to machine precision over a run of 2400 pressure steps, with and without the coupling active.

Grid dependence. Because the wells are single-cell sources with no well model, Δp_u itself depends on the grid, rising from 7.4 MPa at 33×33 to 9.8 MPa at 99×99 and showing no sign of settling; a point source in a finite-volume grid carries a logarithmic singularity, and refinement resolves more of it rather than converging on a limit. This is grid dependence, not grid convergence, and it is worth naming as such. What has to be grid-independent is the ratio at whatever Δp_u the grid produces, and it is: the four grids sit on Eq. (12) to within 1.5%.

Pressure solver. The multiscale finite-volume backend imposes the producer as a pressure anchor rather than a rate sink. The two backends therefore do not solve quite the same well problem and do not share a Δp — the multiscale run sits at 3.2 MPa against 8.2 MPa for the direct solve — so this is not a statement that two discretisations agree on one answer. It is the weaker but still useful statement that a second, structurally different pressure solver, tested at its own Π_u, lands on Eq. (12) to better than 0.1%. The collapse is thus not an artefact of how one solver discretises the pressure equation.

Constitutive law. Replacing Eq. (8) by the Kozeny–Carman form of Eq. (9) removes the fitted permeability modulus entirely and drives the whole effect from rock stiffness instead. Reducing E from 1.6 GPa to 0.1 GPa and collapsing the result on the effective modulus γ_{eff}=dln$k$/d($-\Delta\sigma'_v$) evaluated at the reference state reproduces Eq. (12) to within 6.2%.

Buckley–Leverett. The four checks above establish internal consistency, which is not the same as being right. The transport scheme is therefore run against the one two-phase problem with an exact answer: a 1D incompressible displacement at fixed rate, whose solution is the rarefaction–shock pair of Buckley and Leverett [22] with the shock saturation fixed by Welge's tangent construction [23]. The same upwind face flux and the same explicit saturation update that the five-spot uses every step are driven down a uniform column, with the pressure obtained by integrating the constant total flux back from the outlet. At 0.30 pore volumes injected the computed profile carries a relative L^1 error against the closed form of 2.9% on 100 cells, falling to 0.5% on 800, an observed order of 0.84 that is what a first-order upwind scheme gives on a shock. The front sits at x_D=0.570 against the exact 0.560 on the coarsest grid and at 0.562 on the finest, and the water balance closes to machine precision at every refinement. Table 2 collects the figures. The residual error is numerical diffusion smearing the shock over a few cells, which broadens the front but does not move it, and the displacement measures reported in Section 3.2 are differences between two runs that carry the same smearing.

Table 2. Transport scheme against the Buckley–Leverett solution at 0.30 pore volumes injected. Errors are relative to the analytical profile; the exact front lies at x_D=0.5597.

Cells	L^1 error	L^2 error	Front x_D
100	2.92 %	9.25 %	0.5695
200	1.67 %	6.98 %	0.5653
400	0.94 %	5.20 %	0.5629
800	0.51 %	3.78 %	0.5616

3. Results and discussion

3.1. Injection pressure

Across the 91 coupled runs the permeability modulus spans Π_u from 0.07 to 9.0, which at the upper end raises the permeability by a factor of 10.7 between the producer and the injector. The effect on the pressure required to sustain the prescribed injection rate is large and monotonic. Averaged over the sweep, $\Delta p_c/\Delta p_u$ falls from 0.95±0.03 for Π_u<0.3 to 0.81±0.06 over 0.3≤Π_u<1, to 0.60±0.06 over 1≤Π_u<3, and to 0.39±0.07 beyond that. Ignoring the coupling in the last of these bands would overstate the required injection pressure by roughly a factor of two and a half.

Figure 2a shows every run against the closed-form curve of Eq. (12). The collapse is close: the median departure from the curve is 0.8%, the ninetieth percentile is 5.4%, and the largest single departure is 14%, at the smallest Π_u in the sweep where the pressure difference being compared is itself only a few per cent. Heterogeneity does not degrade the fit: the median departure is 2.1% for the homogeneous field and between 0.6% and 1.2% at every heterogeneous level, with no trend in σ_{lnk}.

That the fit should survive σ_{lnk}=2 is worth dwelling on, because the derivation assumed steady radial flow. It survives because the geometry enters Eq. (11) only through the shape factor G, which is common to the coupled and the uncoupled problem and cancels in the ratio. Heterogeneity moves G a long way. Across the thirteen permeability fields the uncoupled pressure difference runs from 8.2 MPa in the homogeneous case to 33.5 MPa in the most obstructed σ_{lnk}=2 realisation, a factor of four, and the three realisations at that level alone span 8.5, 14.8 and 33.5 MPa — the seed matters as much as the variance. None of that spread reaches the ratio, because a field that is hard to inject into is hard to inject into with and without the coupling. What the criterion needs is not radial symmetry but a prescribed total rate and a pressure that varies monotonically along the flow paths, and a five-spot in log-normal rock satisfies both.

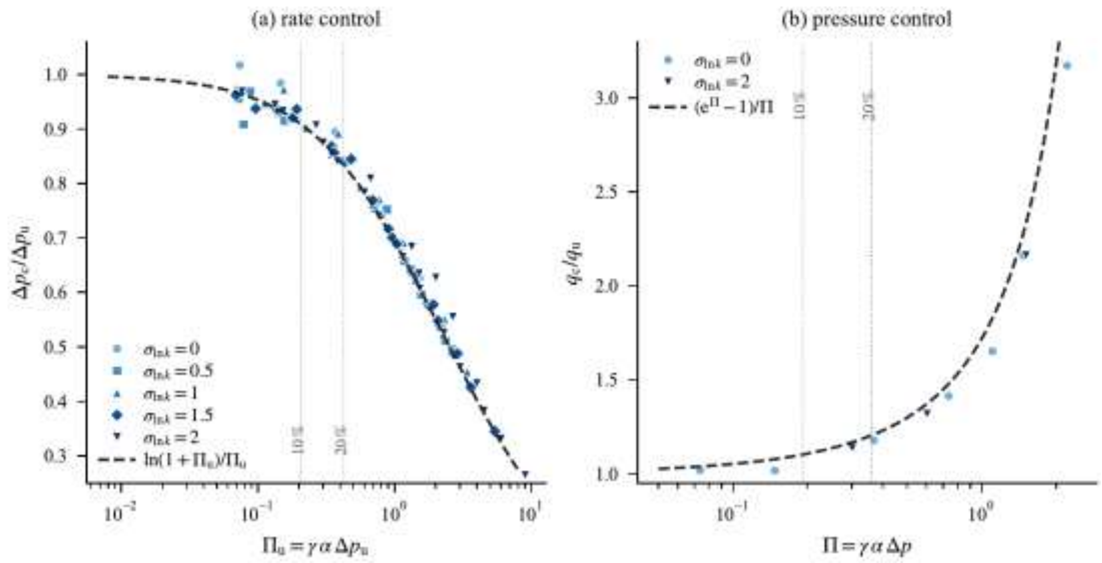

Fig. 2. Screening curves at 0.30 pore volumes injected. (a) Rate-controlled injector: ratio of coupled to uncoupled injector-to-producer pressure difference, for all 91 coupled runs, against the closed form of Eq. (12). (b) Pressure-controlled injector: ratio of coupled to uncoupled injection rate at a prescribed drawdown, measured by the direct experiment of Section 3.4, against Eq. (14). In each panel the dotted lines mark the value of the group at which neglecting the coupling costs 10% and 20%.

3.2. Displacement

The displacement measures behave quite differently. Breakthrough time shifts by a median of 0.1% and by no more than 0.7% in any run; the recovery factor at 0.6 pore volumes shifts by a median of 0.04% and by no more than 0.6%. Every run in the sweep lies inside a band of ±1 per cent on both measures (Fig. 3), including the runs at Π_u=9 where the permeability field is stretched by an order of magnitude and the injection pressure has more than halved.

The saturation fields say why. Figure 4 compares the coupled and uncoupled floods in the most severe case in the study, $\sigma_{\ln k}$=2 at Π_u=2.7. The two fronts are almost superimposed. What the difference map picks out is the front itself, as a thin double band where the coupled flood runs a fraction of a cell ahead of the uncoupled one on one side of a finger and a fraction behind on the other; away from the front the two fields are indistinguishable. Under a prescribed rate the coupling rescales the pressure field along each flow path without greatly altering the relative resistance of one path against another, and the streamline pattern, and with it the sweep, stays where it was. It is the pressure that responds, not the flow geometry.

3.3. A screening criterion

Inverting Eq. (12) turns the observation into a rule. The relative error made in the injection pressure by ignoring stress-dependent permeability is

$$\mathcal{E}=\frac{\Delta p_u-\Delta p_c}{\Delta p_c}=\frac{\Pi_u}{\ln(1+\Pi_u)}-1 \qquad (13)$$

and Table 3 lists the Π_u below which that error stays inside a given tolerance. Because Π_u is built from γ, α and the uncoupled pressure difference, all three are available before any coupled simulation is attempted: run the ordinary model, take the injector-to-producer difference from it, multiply by the permeability modulus and the Biot coefficient, and read off the error.

For orientation, a permeability modulus of 1e-8 /Pa — a factor of 1.1 change in permeability over 10 MPa, typical of a competent, well-cemented sandstone — gives Π_u=0.09 at a 10 MPa drawdown and an error of 4%, which nobody needs to model. At 5e-8 /Pa the error is 21%, and at 1e-7 /Pa, which is within the range reported for weakly cemented and microfractured rock [5, 6], it is 40%. The criterion sorts these three cases without a coupled run in any of them.

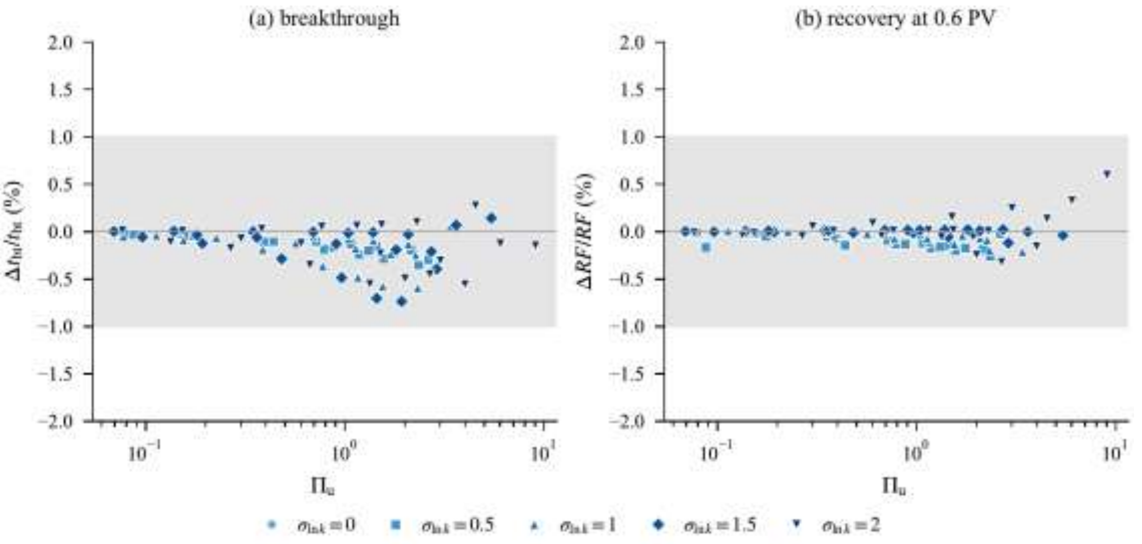

Fig. 3. Change in (a) breakthrough time and (b) recovery factor at 0.6 pore volumes when the coupling is included, relative to the uncoupled run on the same permeability field. Shaded band is ±1 per cent.

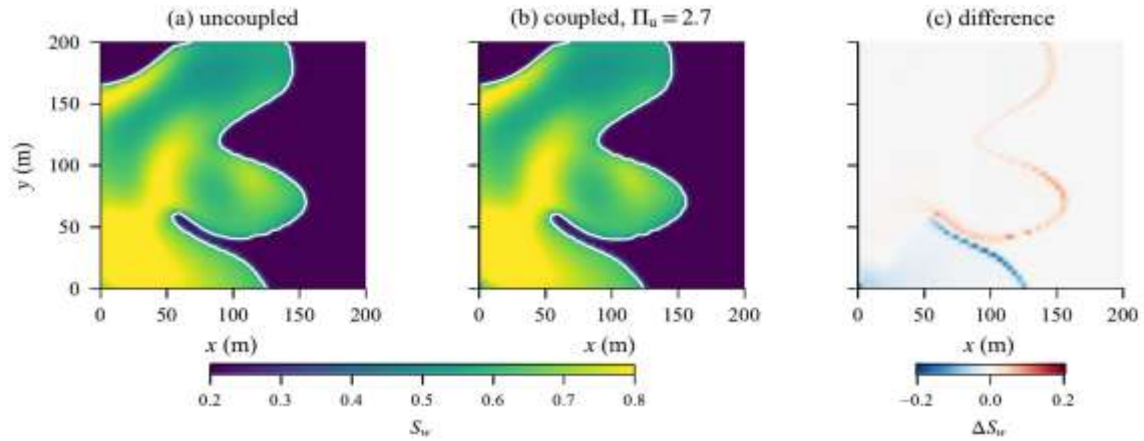

Fig. 4. Water saturation at 0.30 pore volumes injected for the most severe case in the study, $\sigma_{\ln k}$=2. The white line is the S_w=0.5 contour. Panel (c) is the coupled field minus the uncoupled one.

Table 3. Screening thresholds from Eq. (13). Below the listed Π_u, ignoring stress-dependent permeability costs less than the stated error in injection pressure.

Tolerable error in Δp	Π_u below which
2 %	0.04
5 %	0.10
10 %	0.21
20 %	0.43
50 %	1.14

3.4. Pressure-controlled injection

The criterion above was derived and tested for a rate-controlled injector, in which the stress sensitivity is taken up by the pressure. The complementary case — an injector held at a fixed bottom-hole pressure, with the rate left free to respond — follows from the same steady balance, and is worth setting down, because between them the two modes bracket how an injection well is actually run.

Under a prescribed injector-to-producer difference Δp, now common to the coupled and the uncoupled problem, Eq. (11) is read the other way round: the drawdown is fixed and the rate is the unknown. The coupled rate is $q_c\mu/G=(k_0/\gamma\alpha)[\exp(\gamma\alpha\Delta p)-1]$ and the uncoupled rate is $q_u\mu/G=k_0\Delta p$, so that

$$\frac{q_c}{q_u}=\frac{\exp(\Pi)-1}{\Pi}, \qquad \Pi=\gamma\,\alpha\,\Delta p \qquad (14)$$

with the shape factor G cancelling once more. Two things separate this from the rate-controlled result. The group is again $\Pi=\gamma\alpha\Delta p$, but it is now read straight off the prescribed drawdown and needs no distinction between a coupled and an uncoupled pressure; and the ratio lies above unity for every $\Pi>0$, so that the unloading which relieves the injection pressure under rate control reappears here as a gain in injectivity.

The two operating modes make a symmetric pair — $\ln(1+\Pi_u)/\Pi_u\leq 1$ under rate control, $(\exp\Pi-1)/\Pi\geq 1$ under pressure control — and the error each one guards against has the same leading order. Expanded about $\Pi=0$, both depart from unity as $(1)/(2)\Pi$, so the thresholds of Table 3 carry across to injectivity to the same order: below $\Pi\approx 0.2$ the coupling moves injectivity by less than 10%, and above $\Pi\approx 1$ it can no longer be set aside. What separates the two modes is the sign of the effect, and its size at large Π, where Eq. (14) climbs exponentially while its rate-controlled counterpart levels off.

Equation (14) can be tested directly. Because the solver drives the injector at a set rate, a fixed drawdown is reached by adjusting that rate until the coupled injector-to-producer difference at the comparison point matches a prescribed Δp; the injectivity ratio is then that rate divided by the rate the uncoupled well delivers at the same Δp. Each point costs a secant search over full coupled runs, so this experiment is deliberately small: ten points in all, seven on the homogeneous field and three on the $\sigma_{\ln k}=2$, seed 91 realisation, and it is capped at $\Pi=2.5$. The cap is not cosmetic. Beyond it the root-find has to explore drawdowns far outside the screening range, where the exponential map in Eq. (14) magnifies the residual two-phase departures into tens of per cent; and because the heterogeneous field already starts from 33.5 MPa, the cap admits only its three smallest moduli. Every threshold in Table 3 sits at $\Pi<1.2$, comfortably inside what is sampled. Figure 2b gathers the outcome. The measured ratio follows Eq. (14) but sits a little beneath it throughout — by a median of 3% for $\Pi\lesssim 1$, widening to 13% by $\Pi\approx 2$. The closed form is thus an upper bound on the gain: it credits the whole pressure drop with the rise in permeability, whereas in a five-spot the rock near the pressure-anchored producer stays close to its reference value and contributes none, so the real gain is the smaller, and the exponential map lets the shortfall grow with Π. Across the screening range, where $\Pi\lesssim 1$, the bias is small and one-sided, and Eq. (14) is best read there as a conservative estimate of the injectivity gain rather than as an exact equality far beyond it.

3.5. Limitations

Four restrictions bound the result, and the last is the one that matters most.

The wells are single-cell sources with no well model, so Δp_u is itself grid-dependent: it rises from 7.4 MPa on a 33×33 grid to 9.8 MPa on a 99×99 grid, and it is still climbing there. The collapse is unaffected, because it is a statement about the ratio at whatever Δp_u the model produces, and it holds to within 1.5% on all four grids. The criterion is a different matter, because Δp_u is its input: a group formed on a coarse grid and a group formed on a fine one differ by a third here, which is enough to move a case across one of the bands in Table 3. The practical rule is therefore that Π_u must be formed from the same model that will be used for the forecast.

The obvious remedy is a Peaceman well index, and it is worth reporting why we have not adopted one, because the reason is not effort. Referring the pressures to the sandface with endpoint phase mobilities does remove most of the grid dependence, tightening the spread in Δp_u across the four grids from 32% to 7%. It also destroys the collapse, whose departure from Eq. (12) moves from 1.5% to between 18 and 33% with a linear well index, and to between 23 and 55% with a stress-consistent one. The reason is structural rather than numerical. A sandface-to-sandface difference is the sum of three contributions that respond to stress in different ways: the bulk reservoir, which obeys Eq. (12); the injector's near-well region, where the rock is unloaded and the local drop shrinks by around 60%; and the producer's near-well region, where the rock sits below the reference pressure, loses permeability, and the local drop *grows* by around 30%. A linear well index misses the first two effects and a stress-consistent one, in which the near-well drop satisfies the same integrated balance as Eq. (11), makes the third worse. Because the closed form is not linear, a sum of three terms each obeying it does not itself obey it. Recovering a criterion at the sandface is a well-posed problem, but it is the problem of a stress-dependent well index and not a correction to this one; the quantity treated here is a bulk pressure difference and is reported as such.

The residual oil saturation is zero, for the numerical reason set out in Section 2.1. Recovery factors quoted here are consequently higher than a real waterflood would deliver and should be read as comparators between runs rather than as forecasts. The displacement conclusion is a statement about the *difference* between two runs sharing that setting, so it survives; the absolute level does not.

The injector is rate-controlled. The pressure-controlled case is treated in Section 3.4, where a direct experiment confirms the companion form of Eq. (14) as a conservative bound on the injectivity gain; under pressure control the coupling appears as a change in injectivity rather than in pressure.

Most importantly, the pore volume is held fixed in the transport equation. The closure predicts the dilation that accompanies injection — a mean volumetric strain of -3.6×10^{-3} and a porosity rising from 0.200 to 0.204 in the reference case — but because the fluids and the bulk volume are treated as incompressible there is no storage term for that dilation to act through, and admitting it into Eq. (3) would create fluid mass. The volume involved is not negligible: over one run it amounts to 4.8% of the volume injected, an order of magnitude larger than the displacement shifts reported in Section 3.2. The pressure result is untouched by this, since the storage term does not enter the steady balance behind Eq. (11). The displacement result, however, should be read as a statement about the stress-dependent-permeability mechanism alone, and not as a claim that hydromechanical coupling is unimportant for sweep. Settling that question needs a compressible formulation carrying the storage term explicitly, and is the obvious next step.

4. Conclusions

(1) In a rate-controlled waterflood, stress-dependent permeability acts almost entirely on the pressure required to inject, and hardly at all on where the water goes. Over a sweep spanning an order-of-magnitude change in permeability, the injection pressure fell by up to 74% while breakthrough time and recovery moved by less than 1 per cent.

(2) The ratio of coupled to uncoupled pressure difference follows $\ln(1+\Pi_u)/\Pi_u$ with a median error of 0.8%, and does so independently of the degree of heterogeneity, of the grid, of the pressure solver, and of whether permeability is tied to effective stress directly or through porosity. Heterogeneity moves the uncoupled pressure difference by a factor of four across the fields tested without disturbing the collapse, which is the geometric cancellation of Eq. (11) doing its work.

(3) The resulting criterion is cheap to apply: form $\Pi_u=\gamma\alpha\Delta p_u$ from an uncoupled run and read the expected error from Table 3. Below $\Pi_u\approx0.2$ the coupling can be ignored at the 10% level; above $\Pi_u\approx1$ it cannot sensibly be ignored at all. The same group governs a pressure-controlled injector, where the coupling raises injectivity instead of lowering pressure; a direct experiment confirms the companion form $(e^{\Pi}-1)/\Pi$ to within a few per cent over the screening range, on which it is a conservative bound (Section 3.4).

(4) The pore-volume mechanism is outside the scope of all of this, and is large enough to deserve a study of its own.

5. Acknowledgment

Acknowledgment withheld for double-blind review; see the separate author-details file. No external funding supported this work.

The solver, the sweep and verification drivers, and the tabulated output behind every figure and number in this paper are available from the corresponding author on request. The sweep is deterministic: the permeability fields are generated from the three integer seeds quoted in Section 2.4, and rerunning the drivers reproduces the reported values exactly rather than statistically.

6. References

[1] Terzaghi, K. (1943). *Theoretical soil mechanics*. New York, NY: John Wiley & Sons. DOI: 10.1002/9780470172766.

[2] Biot, M. A. (1941). General theory of three-

dimensional consolidation. *Journal of Applied Physics*, 12(2), 155–164. DOI: 10.1063/1.1712886.

[3] Jaeger, J. C., Cook, N. G. W., & Zimmerman, R. W. (2007). *Fundamentals of rock mechanics* (4th ed.). Oxford, UK: Blackwell Publishing.

[4] Rutqvist, J., & Stephansson, O. (2003). The role of hydromechanical coupling in fractured rock engineering. *Hydrogeology Journal*, 11(1), 7–40. DOI: 10.1007/s10040-002-0241-5.

[5] Davies, J. P., & Davies, D. K. (2001). Stress-dependent permeability: Characterization and modeling. *SPE Journal*, 6(2), 224–235. DOI: 10.2118/71750-PA.

[6] Schutjens, P. M. T. M., Hanssen, T. H., Hettema, M. H. H., Merour, J., de Bree, P., Coremans, J. W. A., & Helliesen, G. (2004). Compaction-induced porosity/permeability reduction in sandstone reservoirs: Data and model for elasticity-dominated deformation. *SPE Reservoir Evaluation & Engineering*, 7(3), 202–216. DOI: 10.2118/88441-PA.

[7] Settari, A., & Mourits, F. M. (1998). A coupled reservoir and geomechanical simulation system. *SPE Journal*, 3(3), 219–226. DOI: 10.2118/50939-PA.

[8] Dean, R. H., Gai, X., Stone, C. M., & Minkoff, S. E. (2006). A comparison of techniques for coupling porous flow and geomechanics. *SPE Journal*, 11(1), 132–140. DOI: 10.2118/79709-PA.

[9] Kim, J., Tchelepi, H. A., & Juanes, R. (2011). Stability and convergence of sequential methods for coupled flow and geomechanics: Drained and undrained splits. *Computer Methods in Applied Mechanics and Engineering*, 200(23–24), 2094–2116. DOI: 10.1016/j.cma.2011.02.011.

[10] Longuemare, P., Mainguy, M., Lemonnier, P., Onaisi, A., Gérard, Ch., & Koutsabeloulis, N. (2002). Geomechanics in reservoir simulation: Overview of coupling methods and field case study. *Oil & Gas Science and Technology*, 57(5), 471–483. DOI: 10.2516/ogst:2002031.

[11] Minkoff, S. E., Stone, C. M., Bryant, S., Peszynska, M., & Wheeler, M. F. (2003). Coupled fluid flow and geomechanical deformation modeling. *Journal of Petroleum Science and Engineering*, 38(1–2), 37–56. DOI: 10.1016/S0920-4105(03)00021-4.

[12] Sadrnejad, S. A., Ghasemzadeh, H., & Taheri, E. (2014). Multiscale multiphysic mixed geomechanical model in deformable porous media. *International Journal for Multiscale Computational Engineering*, 12(6), 529–547. DOI: 10.1615/IntJMultCompEng.2014011296.

[13] Taheri, E., Sadrnejad, S. A., & Ghasemzadeh, H. (2015). Multiscale geomechanical model for a deformable oil reservoir with surrounding rock effects. *International Journal for Multiscale Computational Engineering*, 13(6), 533–559. DOI: 10.1615/IntJMultCompEng.2015014333.

[14] Ghasemzadeh, H. (2019). Multiscale multiphysic mixed geomechanical model for deformable porous media considering the effects of surrounding area. *Journal of Petroleum Geomechanics*, 3(1), 79–99. DOI: 10.22107/JPG.2019.88412.

[15] Pedrosa, O. A. (1986, April). Pressure transient response in stress-sensitive formations. Paper SPE-15115 presented at the *SPE California Regional Meeting*, Oakland, CA. DOI: 10.2118/15115-MS.

[16] Kikani, J., & Pedrosa, O. A. (1991). Perturbation analysis of stress-sensitive reservoirs. *SPE Formation Evaluation*, 6(3), 379–386. DOI: 10.2118/20053-PA.

[17] Chin, L. Y., Raghavan, R., & Thomas, L. K. (2000). Fully coupled analysis of well responses in stress-sensitive reservoirs. *SPE Reservoir Evaluation & Engineering*, 3(5), 435–443. DOI: 10.2118/66222-PA.

[18] Raghavan, R., & Chin, L. Y. (2004). Productivity changes in reservoirs with stress-dependent permeability. *SPE Reservoir Evaluation & Engineering*, 7(4), 308–315. DOI: 10.2118/88870-PA.

[19] Brooks, R. H., & Corey, A. T. (1964). *Hydraulic properties of porous media* (Hydrology Paper No. 3). Fort Collins, CO: Colorado State University.

[20] Geertsma, J. (1957). The effect of fluid pressure decline on volumetric changes of porous rocks. *Petroleum Transactions, AIME*, 210, 331–340. DOI: 10.2118/728-G.

[21] Taheri, E. (2015). *Multiscale modeling of oil transport in deformable porous media* (Unpublished doctoral dissertation). K. N. Toosi University of Technology, Faculty of Civil Engineering, Tehran, Iran. (in Persian)

[22] Buckley, S. E., & Leverett, M. C. (1942). Mechanism of fluid displacement in sands. *Transactions of the AIME*, 146(1), 107–116. DOI: 10.2118/942107-G.

[23] Welge, H. J. (1952). A simplified method for computing oil recovery by gas or water drive. *Journal of Petroleum Technology*, 4(4), 91–98. DOI: 10.2118/124-G.