https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

How much of a regional Coulomb stress map survives its own uncertainty? An audit with forecast-skill validation for Türkiye

Mehmet Sait Vural

Department of Computer Engineering, Faculty of Engineering and Natural Sciences, Gaziantep Islam Science and Technology University, Gaziantep, Türkiye

Correspondence: mehmetsait.vural@gibtu.edu.tr

ORCID: 0000-0003-2144-5474

Abstract

Static Coulomb stress transfer is increasingly automated over open catalogues and fault databases rather than built by hand for single events. We ask how much of such a result survives its own inputs. For Türkiye we resolve the stress imparted by 46 Mw ≥ 6 earthquakes between 1990 and 2026 onto 6,970 sub-segments of the mapped active-fault network, screening each across 25 parameter combinations. Three silent failure modes emerge. First, 44 records of the GEM Global Active Faults Database store a central rake 180° from the circular midpoint of their own bounds; all lie in two source compilations, one being the principal North Anatolian Fault strand south of Istanbul, where it reverses the stress change from −0.029 to +0.251 MPa. Second, anchoring ruptures on catalogue centroids rather than mapped traces displaces them by a median of 18.4 km; across 27 testable events this flips the sign on a median of 11.9 % of segments against 2.0 % for a rake ±15° and dip ±10° perturbation, and is the larger of the two in 20 of them. Third, sign stability rises from 74.3 % to 89.5 % between 1 and 50 kPa. A pseudo-prospective test against subsequent seismicity gives a segment-weighted AUC of 0.711 for Mw < 7 sources, but clustering leaves an effective sample of 42, one event supplies 38 % of it, an event-equal estimate gives 0.589 ± 0.112, and distance matching leaves 0.655 with an interval spanning 0.5. Screening improves the estimate by +0.016 (95 % paired bootstrap [+0.001, +0.051]; one-sided p = 0.020). The forecast evidence is suggestive rather than strong; the two input-integrity findings do not depend on it. Coulomb maps should be published with an explicit statement of which parts are interpretable.

Keywords: Coulomb stress transfer; earthquake triggering; fault databases; reproducibility; time-dependent seismic hazard; Türkiye; North Anatolian Fault; East Anatolian Fault

1. Introduction

A single record in an open fault database reverses the sign of the Coulomb stress change on the fault that runs past Istanbul. We found it by auditing our own calculation, and it is one of three failure modes that a regional stress-transfer model can

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

carry without any of them being visible in the result. This paper measures how much of such a model survives once all three are exposed.

Static Coulomb failure stress change (ΔCFS) has become the standard framework for describing how one earthquake alters the likelihood of the next. Since the formulation of King, Stein and Lin (1994) and the synthesis of Stein (1999), the approach has explained aftershock distributions, successive ruptures along the North Anatolian Fault, and the migration of seismicity towards unbroken fault segments. Coupled with the rate-and-state formulation of Dieterich (1994), ΔCFS converts into a time-dependent change in earthquake probability, and in that form it now enters operational hazard products.

Two distinct research practices have grown around the method. The first is the bespoke, single-event study, in which the rupture geometry is taken from a published finite-fault inversion and the receiver planes are selected by an expert. Such studies are accurate but do not scale. The second is the regional or automated application, in which source parameters are drawn from a global moment-tensor catalogue and receiver planes from an open active-fault database. This second practice is expanding rapidly: tools such as AutoCoulomb (Wang et al., 2021) exist precisely to support it, and the release of the GEM Global Active Faults Database (Styron and Pagani, 2020) made global receiver geometry available for the first time.

The uncertainty of the Coulomb calculation itself has been examined. Wang et al. (2014) propagated uncertainties in slip distribution, receiver geometry, friction and Skempton coefficient for the 2008 Wenchuan earthquake and concluded that receiver dip dominates, contributing 40–50 % of the total. Mildon et al. (2016) demonstrated that strike-variable receiver geometry matters more than planar approximations allow, and Mildon et al. (2019) extended the argument to pre-stress and fault bends. Cattania et al. (2014) propagated stress uncertainties through physics-based aftershock models. All of these studies, however, treat uncertainty as *continuous perturbation of parameters that are assumed to be correct*.

A second literature tests whether Coulomb calculations actually forecast anything, and it is the one against which Sect. 4.5 must be read. Hardebeck et al. (1998) examined how far static stress change constrains aftershock locations in southern California. Steacy et al. (2005a) collected the field's position on stress transfer and time-dependent hazard, while Steacy et al. (2005b) showed that the receiver planes onto which the perturbation is resolved materially change the answer, which is precisely the sensitivity our screening addresses. Cocco and Rice (2002) established the poroelastic meaning of the apparent friction we vary between 0.0 and 0.8. Woessner et al. (2011) ran a retrospective comparative forecast test on the 1992 Landers sequence, Mancini et al. (2019) improved physics-based aftershock forecasts through the 2016 central Italy cascade, and the Collaboratory for the Study of Earthquake Predictability (Schorlemmer et al., 2018) now provides the standard framework for such evaluation. Our test is deliberately more modest than a CSEP experiment: it asks only whether the sign of the modelled stress change separates segments that gain seismicity from segments that lose it, and it is used to compare a screened field against an unscreened one rather than to rank forecast models.

A different question has not been asked. When the calculation is automated over open databases, the inputs may be not merely uncertain but wrong: mislocated, mislabelled, or internally self-contradictory. Such errors are categorical rather than continuous. They do not average out. Existing fault-database quality work has concentrated on geological and geometric defects: exaggerated fault lengths, improperly reported blind faults, and tracing without field control. Attribute-level kinematic inconsistency has not been examined. The documentation of the GEM database states explicitly that its build performs only range checks on rake, and that values which are erroneous but cannot be fixed raise warnings during the build process yet are still included (Styron and Pagani, 2020).

This paper therefore asks a question that is methodological rather than tectonic: when a regional Coulomb stress calculation is automated over open data, what fraction of the result survives? We take Türkiye as a test case because it offers a dense instrumental catalogue, a well-mapped active-fault network, two well-documented triggering sequences (1999 İzmit–Düzce

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

and 2023 Kahramanmaraş), and an independently published national hazard model against which a derived hazard field can be checked.

Our contributions are the following.

We identify and quantify a reproducible, internally verifiable defect class in the GEM Global Active Faults Database, namely 44 records whose stored central rake contradicts its own uncertainty bounds, and demonstrate that it reverses the sign of a regionally important Coulomb result.

We show that, across the 27 testable Mw ≥ 6 events, displacing the modelled rupture to the catalogue centroid changes the outcome substantially more than a rake ±15° and dip ±10° focal-mechanism perturbation, and we quantify both effects.

We introduce per-segment robustness screening of a regional Coulomb map and show, through a pseudo-prospective test against subsequent seismicity, that screening improves discrimination by a small margin (+0.016, one-sided $p = 0.020$). The underlying skill signal is itself suggestive rather than established: it rests on 16 source events, is carried disproportionately by one of them, and loses significance once loaded and shadowed segments are matched on distance to the rupture.

We propagate the screened stress field through rate-and-state probability change, fault-source and background probabilistic seismic hazard, and a multi-criteria siting analysis, as a proof-of-concept demonstration that the audit chain can be carried through to a siting layer. It is not offered as an operational planning product.

2. Data

2.1 Earthquake catalogues

The primary instrumental catalogue is the AFAD event service, queried for the box 34–43° N, 25–46° E over 1990–2026 in monthly windows, yielding 578,842 events. Independent catalogues were retrieved for cross-checking: the USGS FDSN event service (M ≥ 5.0, same box and period) and the ISC bulletin (M ≥ 3.5), the latter parsed to 23,910 events of which 6,857 carry an Mw-family magnitude.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

The AFAD catalogue exhibits three defects that constrain its use (Table 1). Two complete calendar years, 1993 and 1997, contain no events whatsoever, although the USGS reports five and eight events of M ≥ 5 respectively in those years. The magnitude scale changes twice: duration magnitude accounts for 100 % of reported magnitudes until 2007, local magnitude becomes dominant by 2012, and moment magnitude appears only from 2013. A total of 34.4 % of all hypocentres are reported at exactly 7.00 km depth, with a further 12.9 % within ±0.06 km of that value, which identifies a fixed default assigned when depth is unresolved; consequently the reported depth carries no information for approximately half the catalogue. The completeness magnitude, estimated by the maximum-curvature method with the standard +0.2 correction, falls from 3.1 in the early 1990s to 1.5 from 2014 onwards.

Table 1. Defects identified in the AFAD instrumental catalogue, 1990–2026: missing years, magnitude-scale transitions, fixed-depth fraction, and the time dependence of the completeness magnitude.

Property	Value
Events retrieved	578,842
Calendar years with no events	1993, 1997 (USGS reports 5 and 8 events of M >= 5)
Magnitude scale, 1990-2006	duration magnitude, 100 %
Magnitude scale, dominant from 2012	local magnitude
Moment magnitude first reported	2013
Hypocentres at exactly 7.00 km	34.4 %
Hypocentres within +/-0.06 km of 7.00 km	a further 12.9 %
Completeness magnitude, early 1990s	3.1
Completeness magnitude, 2014 onwards	1.5
Events with M >= 6, AFAD vs USGS	38 vs 55 (31 % shortfall)

For Mw ≥ 6 the AFAD catalogue is incomplete: it reports 38 such events against 55 in the USGS catalogue for the same box and period, a shortfall of 31 %. All target events for the stress calculation were therefore taken from moment-tensor solutions rather than from AFAD.

2.2 Focal mechanisms

Global Centroid Moment Tensor solutions were obtained as NDK files (the combined 1976–2020 file plus monthly files to 2026) and parsed to 62,197 global solutions for 1990–2026. Within the study box, 442 solutions have Mw ≥ 5.0, 46 have Mw ≥ 6.0 and six have Mw ≥ 7.0. Parsed nodal planes reproduce published solutions: the 1999 İzmit event returns Mw 7.57 with a plane striking 091°, dipping 87° and with rake 164° (right-lateral, consistent with the North Anatolian Fault), and the 2023 Pazarcık event returns Mw 7.78 with a plane striking 051°, dipping 70° and rake −4° (left-lateral, consistent with the East Anatolian Fault).

2.3 Active faults

Receiver geometry is taken from the GEM Global Active Faults Database, harmonised release. The study box contains 515 fault features, drawn from the EMME (472), SHARE (39) and Bird (4) source compilations. Each feature carries dip, rake and net slip rate as (best, minimum, maximum) triples together with upper and lower seismogenic depths. Traces were resampled into sub-segments of target length 5 km, giving 6,970 receiver sub-segments and 34,846 km of mapped fault. Each sub-segment carries its own strike, computed from the local trace azimuth.

2.4 Site and exposure data

Topography is GMTED2010 mean elevation at 30 arcsec, mosaicked over the study box. Population exposure is WorldPop 2020 at 1 km resolution for Türkiye, aggregated to the analysis grid.

3. Methods

3.1 Magnitude homogenisation

Because the AFAD magnitude scale changes twice (Sect. 2.1), a conversion to Mw is required. Rather than importing published coefficients, relations were derived from matched event pairs following the same methodology as the published Turkish standard (Tan, 2021): general orthogonal regression, appropriate because both magnitudes carry measurement error, with 95 % confidence intervals from 1,000 bootstrap resamples. AFAD events were matched to GCMT, USGS and ISC solutions within ±30 s and 60 km, giving 6,494 matched pairs (Table 2, Fig. 1). The local-magnitude relation, Mw = 0.9656 ML + 0.2318 (n = 1,513, RMS = 0.258), has a slope consistent with published Turkish relations.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

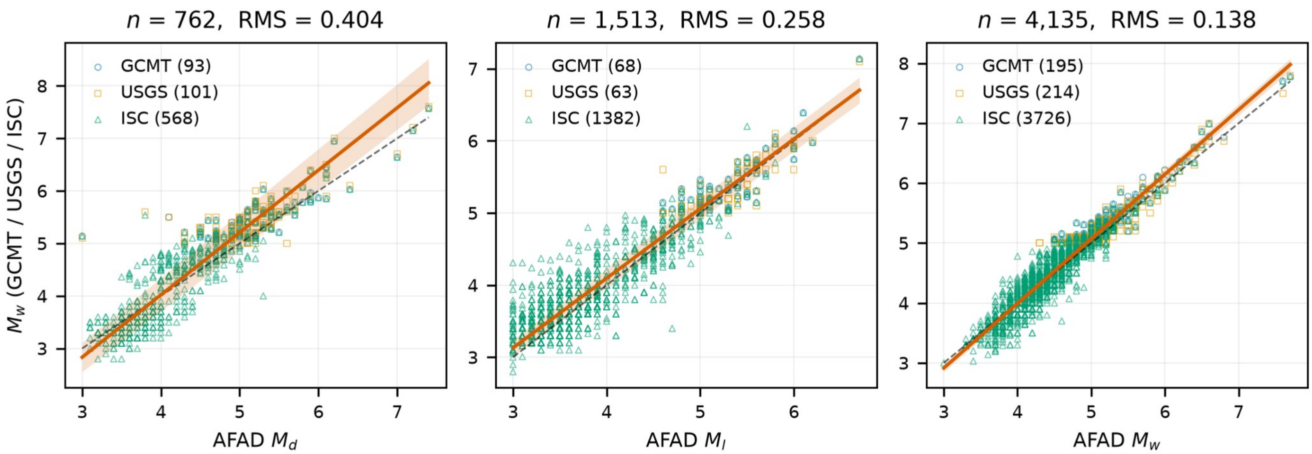

Figure 1. AFAD magnitude scales against Mw from GCMT, USGS and ISC, with general-orthogonal-regression fits and bootstrap confidence bands. Duration, local and moment magnitudes are shown separately; the one-to-one line is given for reference.

Table 2. Magnitude conversion relations derived by general orthogonal regression from 6,494 matched pairs, with 95 % bootstrap confidence intervals on slope and intercept.

AFAD scale	n	Valid range	Slope [95 % CI]	Intercept [95 % CI]	RMS
MD	762	3.0-7.4	1.1874 [1.1533, 1.2288]	-0.7301 [-0.9033, -0.5735]	0.404
ML	1,513	3.0-6.7	0.9656 [0.9484, 0.9825]	+0.2318 [+0.1674, +0.2960]	0.258
MW	4,135	3.0-7.7	1.0777 [1.0692, 1.0862]	-0.3196 [-0.3548, -0.2849]	0.138

3.2 Rupture source model

Contribution (2) of this paper is that rupture position matters more than the mechanism perturbation we evaluate. The source model therefore separates the two components: the mechanism (strike, dip, rake) is taken from GCMT, with the nodal plane selected by best match to nearby mapped faults, whereas the map position is obtained by projecting the hypocentre onto the matched trace. Ruptures follow the mapped trace as a chain of sub-patches rather than as a single planar rectangle. Rupture length follows Wells and Coppersmith (1994), except for four events with published surface-rupture extents (1999 İzmit, Barka et al., 2002; 1999 Düzce, Akyüz et al., 2002; 2023 Pazarcık and Elbistan, Reitman et al., 2023), for which the mapped extent is used. Where the mapped trace supplies less than 60 % of the target length, a single straight rectangle of the full length is substituted, because concentrating a fixed moment on an under-length trace inflates slip. Seismic moment is conserved throughout: slip equals $M_0/(\mu A)$, with A the total patch area. The resulting model comprises 374 patches

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

across 46 events, with mean slip at a median of 0.96 times the Wells and Coppersmith expectation and 45 of 46 events within a factor of 0.4–2.5.

3.3 Coulomb stress engine and its verification

Displacement and its gradient are computed with the Okada (1992) rectangular dislocation solution as implemented in the compiled okada_ext extension of the Pyrocko toolbox (Heimann et al., 2017). Strain follows from the symmetric part of the displacement gradient and stress from Hooke's law with $\lambda = \mu = 32$ GPa. The stress tensor is resolved onto each receiver plane to give the shear traction in the rake direction, $\Delta\tau$, and the tension-positive normal traction, $\Delta\sigma_n$, whence

$$\Delta CFS = \Delta\tau + \mu' \Delta\sigma_n \quad (1)$$

with apparent friction μ'.

One further convention is used throughout and is defined here because several results depend on it. Receiver segments lying within **15 km of the rupture plane** are masked from every calculation. The distance is measured to the rectangle itself, not to its centre: a receiver is projected into along-strike and perpendicular coordinates, the along-strike excess beyond the half-length is taken, and the distance to the rectangle follows. The mask has a fixed width rather than one scaled to rupture length, which means it removes a larger share of the loaded area for a long rupture than for a short one; Sect. 5.3 returns to this. Its purpose is twofold: the elastic solution is singular at the rupture plane, and the source patch has already released its own stress.

Four independent verifications were performed. The free-surface boundary condition is satisfied to a relative error of 1×10^{-15}: at $z = 0$ the tractions on a horizontal plane vanish to that precision. The stress tensor is exactly symmetric. Far-field amplitude decays as r^{-3}, the ratio between 200 km and 400 km being 8.40 against a theoretical 8.0. Finally, the classical lobe pattern of King, Stein and Lin (1994) is reproduced: for a right-lateral vertical strike-slip rupture, ΔCFS is positive beyond both along-strike tips and negative in the flanking off-fault lobes, correct in all four quadrants (Fig. 2).

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

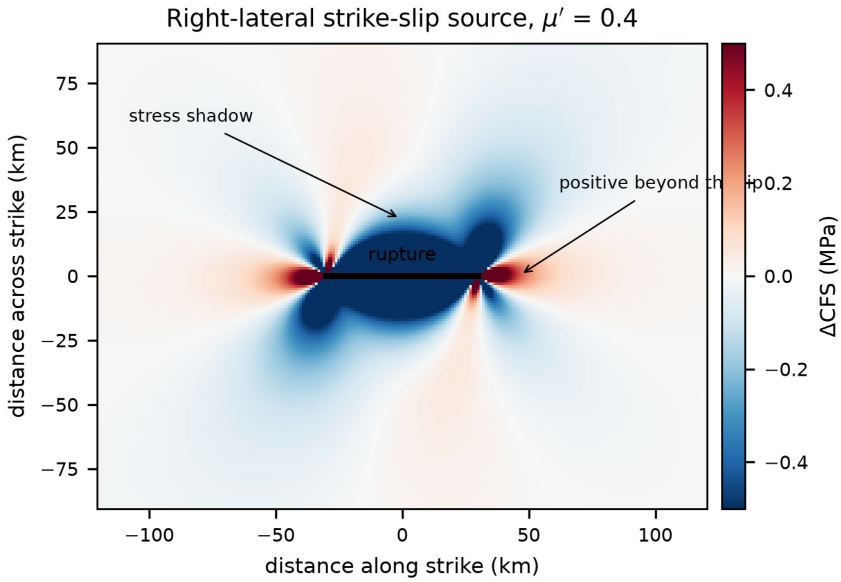

Figure 2. Verification of the Coulomb engine. ΔCFS for a right-lateral vertical strike-slip rupture reproduces the King, Stein and Lin (1994) lobe pattern: positive beyond both along-strike tips, negative in the flanking off-fault lobes.

Literature validation used the 1999 İzmit-to-Düzce transfer. The automated pipeline, without tuning, returns ΔCFS positive on all nine Düzce-fault segments, ranging from +0.144 to +0.029 MPa, and +0.084 MPa at the Düzce hypocentre, against a published range of +0.1 to +0.35 MPa (Hubert-Ferrari et al., 2000; Parsons et al., 2000). The sign and order of magnitude are reproduced; the amplitude is 20–50 % low, as expected of a uniform-slip approximation compared with finite-fault models that concentrate slip near the eastern rupture tip.

3.4 Receiver-rake reconciliation

The GEM database stores rake as a (best, minimum, maximum) triple. The dip triples confirm that ordering unambiguously: (56.5, 35.0, 78.0) has 56.5 exactly equal to (35 + 78)/2. Applying the same reading to rake exposes records whose stored central value lies more than 90° from the circular midpoint of its own bounds, the criterion used throughout; in practice every record that fails does so by close to a full 180°. Because such a record contradicts itself, no external assumption is needed to identify it, and correction consists of rotating the central value by 180°.

A second, weaker test compares rake against the slip_type label. We show in Sect. 4.1 that this test is not reliable: it is confounded by convention differences between source compilations, and slip_type can itself be mislabelled. Only the internal test is used for correction; records lacking bounds are left untouched and flagged.

3.5 Rate-and-state probability change

Following Dieterich (1994) and Toda et al. (2005), a stress step ΔCFS modifies the seismicity rate as $R(t)=r/[(\gamma-1)e^{-t/t_a}+1]$ with $\gamma=\exp(-\Delta CFS/A\sigma)$. Integrating gives the expected number over a forecast window T,

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

$$N(T) = r\, t_a \ln\left[\frac{e^{T/t_a} + \gamma - 1}{\gamma}\right] \tag{2}$$

so that the probability gain N(T)/rT is independent of the poorly constrained background rate r, which is the quantity we report. The stressing rate is $\dot{\tau} = \mu v / (2\pi D)$ from the mapped slip rate v and seismogenic depth D, and Aσ = 0.04 MPa follows Parsons et al. (2000). We also report the clock advance ΔCFS/$\dot{\tau}$ both in years and as a fraction of the segment recurrence interval, since a 1,000-year advance on a 0.5 mm yr−1 fault and a 20-year advance on a 20 mm yr−1 fault are not comparable quantities.

Three data-quality gates are applied and nothing is imputed. A segment receives time-dependent products only if its slip rate is actually reported by GEM, is at least 1 mm yr−1, and the segment lies outside the near-source mask. A fourth gate excludes segments whose clock advance exceeds half a recurrence interval, since these lie in the near-tip zone where the elastic solution is least reliable. Of 6,970 segments, 3,329 are eligible and 3,248 are physically interpretable.

3.6 Position against mechanism

To test whether rupture position has a larger effect than the focal-mechanism perturbation considered here, rather than only in the illustrative İzmit case, both are perturbed for every event meeting the analysis criteria and compared on the same footing. The comparison is bounded in two ways that its wording should reflect: it covers the 27 events with enough loaded fault to be testable, not all 46, and it contrasts position against one specified mechanism perturbation rather than the whole space of focal-mechanism uncertainty. The **position** variant re-anchors the whole rupture on the GCMT centroid instead of the mapped trace; this is not a hypothetical displacement but the actual centroid-to-trace offset for that event, and it is the default that any pipeline inherits by taking the catalogue position at face value. The **mechanism** variant keeps the mapped position and perturbs the orientation by a typical GCMT uncertainty, rake ±15° and dip ±10°, with the worst of the four corners retained so that the comparison favours the mechanism. Changing the dip alters the vertical extent, so the centre depth is recomputed from the fixed top depth. Both variants are compared with the reference model on every receiver segment carrying |ΔCFS| > 5 kPa outside the near-source mask, and we record the fraction of segments whose sign changes and the median absolute change.

3.7 Robustness screening

Because the stress tensor depends on neither the apparent friction nor the receiver-plane orientation, it is computed once and stored, after which every parameter case is resolved directly. Twenty-five cases are evaluated: five apparent-friction values (μ′ = 0.0, 0.2, 0.4, 0.6, 0.8) crossed with five receiver-plane cases (best, rake plus and minus its half-width, and dip at its GEM minimum and maximum). A segment is flagged **sign-stable** when the sign of ΔCFS is identical in all 25 cases.

This is deliberately a **robustness screen rather than a full uncertainty propagation**, and the distinction matters. A propagation would sample dip and rake from continuous distributions by Monte Carlo and return a posterior on ΔCFS. That is the right tool when the input distributions are known. Here they are not: the database supplies bounds rather than distributions, and Sect. 4.1 shows that some stored central values are not merely uncertain but wrong, so a distribution centred on them would be misspecified. The screen instead asks whether the sign changes across a predefined set of parameter combinations spanning the reported bounds. It is a **discrete sensitivity screen, not a guarantee of sign stability throughout the full continuous multidimensional parameter space**: the 25 cases cross five friction values with five receiver-plane variants, and do not evaluate every joint combination of rake and dip within their bounds. Segments classified as sign-stable should therefore be read as stable across the 25 evaluated cases rather than as formally robust to every admissible combination. The screen also returns a binary flag rather than a probability, and must not be read as a credible interval. The choice of endpoints is motivated rather than arbitrary: ΔCFS is a continuous function of friction, rake and dip, so for a response that is monotonic in each of them a sign that survives the endpoints would also survive the values between. We do not claim monotonicity holds everywhere — receiver-plane rotation can be non-monotonic near nodal orientations, which is exactly why the screen is stated as a discrete result — but it is the expected behaviour away from those orientations, and it is why the endpoints of the stated bounds were chosen in preference to a random sample of the interior.

3.8 Time-dependent hazard and siting: a downstream demonstration

What follows is a **proof-of-concept demonstration that the audited stress field can be propagated through a downstream hazard and siting workflow**, not a contribution to probabilistic seismic hazard methodology and not an operational decision product. The hazard model deliberately reuses standard components and treats its epistemic uncertainty far more coarsely than the Coulomb analysis above. The validation in Sect. 4.6 is a sanity check on the ordering of hazard between cities rather than a full hazard-model validation, and the siting layer should be read in the same spirit.

Each mapped fault with a reported slip rate and total length ≥ 20 km becomes a characteristic source: magnitude from length by inverting Wells and Coppersmith (1994), rate from the moment budget $\dot{M}_0 = \mu A v$, and a 30-year Poisson probability modified by the rate-and-state gain. Background seismicity follows Frankel (1995): the declustered catalogue is binned on the hazard grid, smoothed with a 50 km Gaussian kernel, and each cell becomes a point source with a truncated Gutenberg–Richter distribution capped below the fault-source magnitudes to prevent double counting. Declustering uses the Gardner and Knopoff (1974) windows and removes 69.3 % of events, raising the b-value from 0.753 to 0.865 — the expected direction, since aftershock sequences depress b. Ground motion uses the Türkiye-specific model of Kale et al. (2015) with VS30 derived from topographic slope after Wald and Allen (2007).

Siting suitability combines five criteria by the analytic hierarchy process: time-dependent hazard (weight 0.4065), proximity to a mapped active fault (0.2490), site amplification (0.1428), liquefaction susceptibility (0.1428) and slope instability (0.0589). The consistency ratio is 0.0127, well within the admissible limit of 0.10 (Saaty, 1980).

The provenance of these weights should be stated plainly, because it is a limitation. They were not obtained by eliciting judgements from a panel of experts. They encode the author's own reasoning, which ranks the two criteria that describe where an earthquake will strike above the three that describe how the ground will respond, and they are reported so that a reader can disagree with them explicitly. No claim of consensus is made. What carries the result is not the particular weights but their irrelevance to the ordering: perturbing every pairwise comparison log-normally by ±50 % at 95 % confidence over 1,000 draws leaves the city ranking essentially unchanged (Sect. 4.6). A siting analysis intended for actual planning use would require formal elicitation from the relevant engineering and planning authorities, and its weights would very probably differ from ours.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

3.9 Pseudo-prospective skill test

The decisive test asks whether the sign of ΔCFS predicts where seismicity subsequently goes. Each segment serves as its own control: the local earthquake count in the two years before each source event is compared with the count in the two years after, within 15 km. Using the rate change rather than the absolute rate removes spatial variation in background seismicity, and de-meaning the metric within each event removes any region-wide activity change. Sixteen source events have two years of catalogue on both sides, giving 49,157 usable segment–event pairs, of which 40,079 belong to the Mw < 7 subset used for the magnitude-stratified results. Discrimination is measured by the area under the receiver-operating characteristic curve (AUC) between loaded and shadowed segments.

The test is **pseudo-prospective rather than prospective in the CSEP sense**. It is run forward in time, in that seismicity after each source event is never used to build the model of that event, but it is executed retrospectively and the source model for four events uses surface-rupture maps published after the earthquake occurred (Barka et al., 2002; Akyüz et al., 2002; Reitman et al., 2023). A genuinely prospective experiment would have to be declared before the target earthquakes. We use the weaker term throughout.

Four event counts recur below and are fixed here to avoid confusion. Sixteen source events have two years of catalogue on both sides and enter the pooled test; 13 of these retain segments above the 5 kPa threshold outside the mask; 8 retain enough loaded and shadowed segments to be scored individually at 1 kPa; and 6 can be scored individually within the Mw < 7 subset at 5 kPa.

Inference is clustered on the source event. Segment-event pairs are not independent: neighbouring 5 km sub-segments around one source share that source and much of the subsequent seismicity, so treating 49,157 pairs as independent observations would overstate significance severely. The independent unit is the source event, of which there are 16. Two event-level procedures are therefore used. A block bootstrap resamples whole events with replacement (2,000 draws) to give a confidence interval on AUC. A permutation test reverses the stress sign for whole events (2,000 draws), so that the exchangeable unit is the event rather than the segment; p values carry the standard +1 correction and are therefore never reported as zero. The per-event AUC distribution and a sign test are reported alongside, and the intra-event correlation is used to quantify the design effect explicitly. Bootstrap and permutation both use 2,000 draws; the separate within-event shuffle that tests the metric itself uses 300.

Two controls guard against the obvious ways this test could mislead. First, aftershocks concentrate near the rupture and ΔCFS is positive beyond the rupture tips, so a naive comparison might reward proximity rather than positive stress. De-meaning within an event removes region-wide rate changes but not distance decay, so loaded and shadowed segments are additionally **matched on distance to the rupture** in 10 km bands and the AUC is recomputed within strata. Second, whether screening helps is a question about the **difference** between two AUC values, which is not answered by observing that one confidence interval excludes 0.5 and the other does not. The difference is therefore tested directly with a paired bootstrap over whole events, so that the screened and unscreened estimates move together across draws. Because screening was defined in advance to remove parameter-unstable predictions, the alternative for this one comparison is directional, and a one-sided p value is reported alongside the two-sided confidence interval on the effect size. Every other p value in this paper is two-sided unless stated.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

4. Results

4.1 A self-contradictory rake defect in an open fault database

Across the whole GEM database, 2,076 records carry both rake bounds and can therefore be tested internally. Forty-four fail: their stored central rake lies more than 90° from the circular midpoint of their own bounds. All 44 lie in two source compilations — EMME (27 records, 3.47 %) and SHARE (17 records, 2.71 %) — and none occur in UCERF3, Litchfield_NZ_2013, EOS_SE_Asia, GEM_N_Africa, GEM_Central_Am_Carib or GEM_NE_Asia. The signature is uniform (Table 3): every affected record is labelled Dextral or Dextral-Normal, has bounds bracketing ±180°, and a stored central value near 0°. This pattern indicates a wrap or sign error in the harmonisation of two source compilations rather than scattered data-entry mistakes. All 44 records are listed individually, with their catalogue identifiers and locations, in Supplementary Table S1, so that the claim can be checked record by record rather than taken on the strength of the five examples in Table 3.

Table 3. The 44 internally inconsistent rake records in the GEM Global Active Faults Database, by source compilation, with representative examples showing the stored central value against the circular midpoint of its own bounds.

(a) By source compilation (records carrying both rake bounds)

Source compilation	Testable	Failing	Share
EMME	778	27	3.47 %
SHARE	627	17	2.71 %
Litchfield_NZ_2013	635	0	0.00 %
GEM_Central_Am_Carib	10	0	0.00 %
GEM_NE_Asia	10	0	0.00 %
GEM_N_Africa	16	0	0.00 %
Total	**2,076**	**44**	**2.12 %**

(b) Representative failing records

slip_type	Stored average_rake (best, min, max)	Circular midpoint of bounds	Discrepancy
Dextral	(12, 180.0, -156.0)	-168 deg	180 deg
Dextral	(15, 180.0, -150.0)	-165 deg	180 deg
Dextral	(20, 180.0, -140.0)	-160 deg	180 deg
Dextral	(-2, 173.0, -177.0)	+178 deg	180 deg

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

slip_type	Stored average_rake (best, min, max)	Circular midpoint of bounds	Discrepancy
Dextral-Normal	(26, 180.0, -127.0)	-153 deg	180 deg

(c) Why the slip_type test cannot be used for correction

Test	Litchfield_NZ_2013 records failing	Interpretation
rake vs slip_type	631 of 635	different rake sign convention, not an error
rake vs its own bounds	0 of 635	records are internally consistent

The consequence is not cosmetic. ΔCFS resolves shear traction in the rake direction, so rotating rake by 180° reverses its sign. One affected record is the principal North Anatolian Fault strand in the Sea of Marmara, immediately south of Istanbul. Using the stored rake, the Marmara segments return a median ΔCFS of −0.029 MPa following the 1999 İzmit earthquake, that is, a stress shadow. Using the bounds-consistent rake they return +0.251 MPa, a loaded segment, in agreement with the published literature for that fault (Hubert-Ferrari et al., 2000; Murru et al., 2016).

The alternative test, comparing rake against slip_type, must not be used for correction, for two reasons established here. First, it is confounded by convention: 631 of 635 Litchfield_NZ_2013 records fail the slip_type test while none fail the internal test, indicating that the New Zealand compilation uses a different rake sign convention rather than containing errors. Second, slip_type can itself be wrong. Applying a slip_type-driven correction to the Türkiye subset rotates 516 sub-segments, of which 93 are internally consistent and should not have been touched; 30 of these lie on the Malatya–Ovacık Fault Zone (approximately 38.0–39.1° N, 37.5–38.4° E), which GEM labels Dextral but which is documented as sinistral (Westaway and Arger, 2001; Sançar et al., 2019). There the stored rake is correct and the label is wrong, and the naive correction corrupts the result: one of the affected segments carried the maximum ΔCFS of the entire map.

The two defects are therefore distinct and require different diagnostics: a rotated central rake is detectable internally, whereas a mislabelled slip_type is detectable only against external geological information. Applying the internal rule alone corrects 423 of 6,970 Türkiye sub-segments (6.07 %).

4.2 Rupture position dominates focal mechanism

Across the 46 Mw ≥ 6 events, the separation between the GCMT centroid and the PDE hypocentre has a median of 18.4 km, a 90th percentile of 30.0 km and a maximum of 77.0 km. For the 1999 İzmit earthquake the centroid lies at 41.01° N whereas the mapped rupture follows the North Anatolian Fault near 40.73° N, a displacement of 29 km that places part of a symmetric rupture box inside the Sea of Marmara, where no rupture occurred.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

The effect on the computed transfer to Düzce is decisive (Table 4). A rupture centred on the GCMT centroid returns −0.054 MPa; the same rupture placed on the mapped trace returns +0.177 MPa. Substituting the GCMT rake and dip while keeping the mapped position changes the result by only 4 % (+0.170 MPa), and burying the rupture to the centroid depth changes it by 27 % while preserving the sign (+0.129 MPa).

Table 4. Sensitivity of the 1999 İzmit-to-Düzce Coulomb transfer to source geometry: centroid against mapped trace, mechanism substitution, and burial depth.

Receiver plane: Duzce fault, strike 265 deg, dip 60 deg, rake 180 deg; mu' = 0.4, depth 11 km.

Source geometry	dCFS (MPa)	Change vs mapped case
GCMT centroid, scaling-law box	-0.054	sign reversed
Mapped surface rupture (reference)	+0.177	-
Mapped position, GCMT rake and dip	+0.170	-4 %
Mapped position, buried to centroid depth	+0.129	-27 %, sign preserved

Published range for this transfer: +0.1 to +0.35 MPa.

This is one event, so the same comparison was run for the whole catalogue (Fig. 3, Table 5). Twenty-seven of the 46 events retain at least 20 segments above the 5 kPa threshold outside the near-source mask and can therefore be tested; the remainder load too little mapped fault to support the comparison. Across those 27 events the median centroid-to-trace offset is 19.5 km, and re-anchoring the rupture on the centroid changes the sign of ΔCFS on a median of **11.9 %** of segments, against **2.0 %** for a rake ±15° and dip ±10° perturbation. The median absolute change is **7.4 kPa** for position against **2.3 kPa** for mechanism. Position dominates in 20 of 27 events by the sign criterion and in 24 of 27 by the magnitude criterion, and paired two-sided Wilcoxon signed-rank tests reject equality in both ($p = 5.4 \times 10{-}4$ and $p = 1.0 \times 10{-}6$ respectively). Two events are tied on the sign criterion.

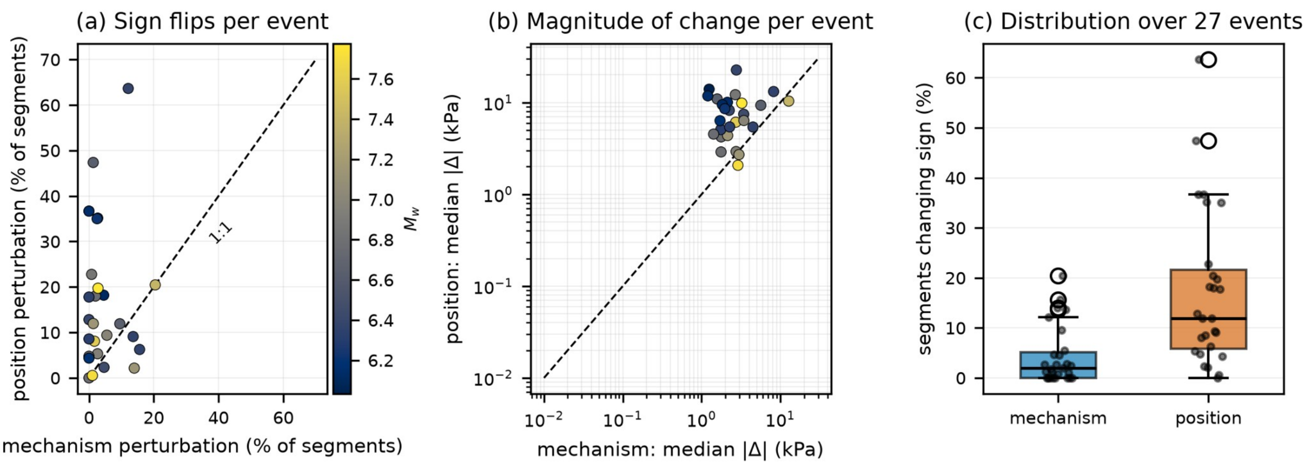

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Figure 3. Rupture position against focal mechanism across 27 Mw ≥ 6 events. (a) Percentage of receiver segments whose ΔCFS changes sign under a position perturbation (re-anchoring on the GCMT centroid) against a mechanism perturbation (rake ±15°, dip ±10°, worst corner); points are coloured by moment magnitude and the dashed line is 1:1. (b) The same comparison for the median absolute change, on logarithmic axes. (c) Distributions of the sign-flip fraction for both perturbations, with individual events overplotted. Position lies above the 1:1 line in 20 of 27 events in (a), with 2 further events tied, and in 24 of 27 in (b).

Table 5. Position against mechanism for every testable event: centroid-to-trace offset, sign-flip fraction and median absolute change under each perturbation.

Position variant: rupture re-anchored on the GCMT centroid. Mechanism variant: rake +/-15 deg and dip +/-10 deg, worst corner. Both compared with the reference model on segments carrying |dCFS| > 5 kPa outside the near-source mask.

Event	Mw	Centroid offset (km)	Segments	Sign flips, position	Sign flips, mechanism	Median change, position (kPa)	Median change, mechanism (kPa)
1991-04-29	6.95	12.7	113	5.3 %	2.7 %	2.9	2.7
1992-03-13	6.64	25.6	76	47.4 %	1.3 %	10.9	1.6
1992-10-23	6.40	19.5	39	12.8 %	0.0 %	8.2	2.2
1992-11-06	6.03	26.8	30	36.7 %	0.0 %	14.0	1.2
1995-10-01	6.38	43.1	33	63.6 %	12.1 %	22.6	2.8
1996-10-09	6.78	7.6	21	4.8 %	0.0 %	4.2	1.8
1999-08-17	7.57	30.4	533	8.1 %	1.7 %	6.1	2.7
1999-11-12	7.15	20.4	226	11.9 %	1.3 %	4.4	2.1
2002-01-22	6.15	29.0	30	36.7 %	0.0 %	11.8	1.2
2002-02-03	6.45	7.6	43	2.3 %	4.7 %	7.4	3.4
2003-05-01	6.32	6.9	35	8.6 %	0.0 %	5.1	1.8
2011-04-01	6.14	21.4	37	35.1 %	2.7 %	10.0	2.1
2011-10-23	7.13	13.1	186	2.2 %	14.0 %	2.7	3.0
2014-05-24	6.86	24.6	123	22.8 %	0.8 %	12.2	2.7
2015-04-16	6.12	13.5	22	18.2 %	4.5 %	9.5	1.8
2017-06-12	6.36	13.8	32	6.2 %	15.6 %	5.4	4.5
2017-07-20	6.64	19.7	42	11.9 %	9.5 %	9.3	5.6
2017-11-12	7.40	14.1	44	20.5 %	20.5 %	10.4	12.7
2020-01-24	6.76	14.0	50	18.0 %	2.0 %	4.5	1.4
2020-10-30	6.99	19.6	181	9.4 %	5.5 %	6.4	3.4

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Event	Mw	Centroid offset (km)	Segments	Sign flips, position	Sign flips, mechanism	Median change, position (kPa)	Median change, mechanism (kPa)
2021-10-12	6.36	29.1	22	9.1 %	13.6 %	13.2	8.2
2023-02-06	7.78	52.1	573	19.7 %	2.8 %	9.8	3.2
2023-02-06	6.78	5.0	59	0.0 %	0.0 %	2.9	1.8
2023-02-06	7.70	7.8	564	0.5 %	1.1 %	2.1	2.9
2023-02-20	6.34	12.2	45	17.8 %	0.0 %	5.4	2.3
2025-05-22	6.22	33.2	40	35.0 %	2.5 %	8.6	2.0
2025-08-10	6.18	18.3	23	4.3 %	0.0 %	6.3	1.7
median		**19.5**		**11.9 %**	**2.0 %**	**7.4**	**2.3**

Position dominates in 20 of 27 events by the sign criterion, with 2 further events tied, and in 24 of 27 by the magnitude criterion. Paired two-sided Wilcoxon signed-rank tests give p = 5.4e-04 and p = 1.0e-06 respectively.

The ordering is therefore consistent across the testable subset rather than an artefact of the İzmit example: under the perturbations evaluated here, map position can reverse the sign, whereas the specified mechanism perturbation rarely does. For automated regional studies this inverts the conventional emphasis, which has centred on receiver geometry (Wang et al., 2014; Mildon et al., 2016). Those studies addressed single events whose source position was fixed by an expert; automation reintroduces position as a free parameter, and of the two uncertainty sources evaluated here it is the larger.

4.3 The screened regional stress field

The cumulative field over 46 sources ranges from −2.271 to +1.556 MPa, with 729 segments loaded above +0.01 MPa and 874 shadowed below −0.01 MPa (Fig. 4). These counts cover the whole network; 224 of those 1,603 segments fall inside the near-source mask of at least one event and are excluded from every interpretation, which is why Table 6 reports 1,379 segments at the same 10 kPa threshold. Sign stability across the 25 parameter cases improves systematically with signal strength (Table 6): 74.3 % of segments with |ΔCFS| > 1 kPa are sign-stable, rising to 82.7 % above 5 kPa, 85.2 % above 10 kPa and 89.5 % above 50 kPa. Rank correlation with the reference case remains high throughout, with a minimum Spearman ρ of 0.798 and a median of 0.921.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

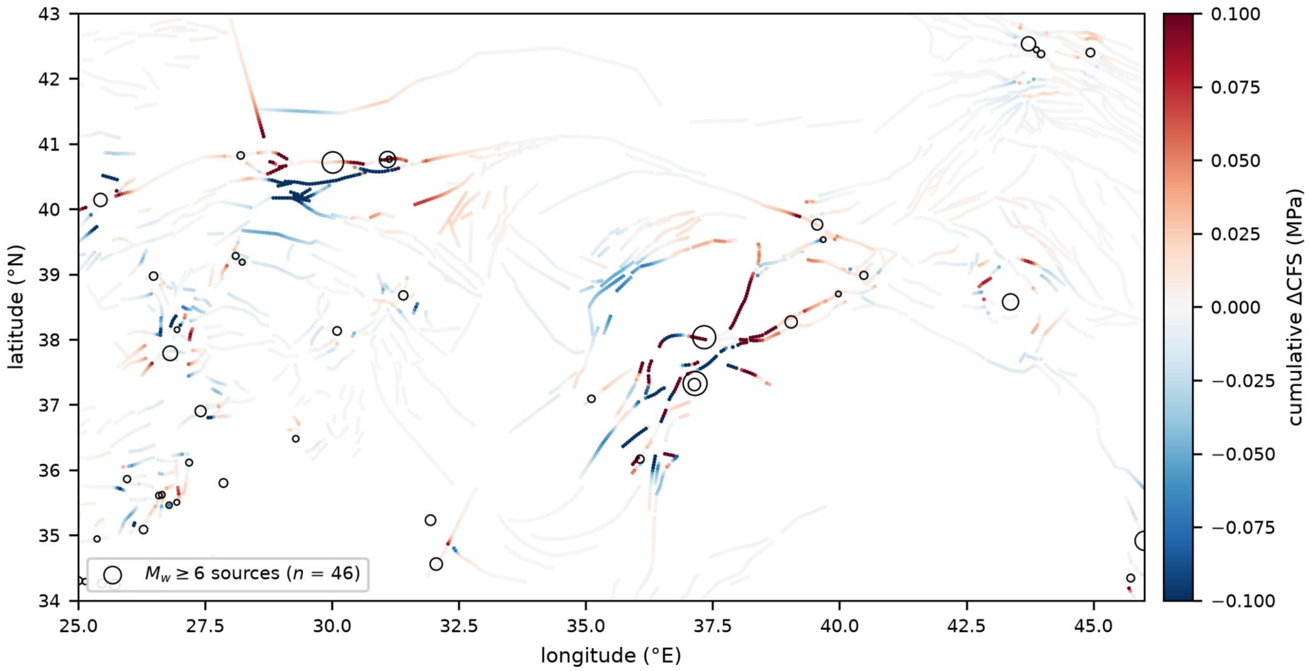

Figure 4. Cumulative ΔCFS resolved on 6,970 mapped fault sub-segments from 46 Mw ≥ 6 earthquakes, 1990–2026, at μ′ = 0.4. Open circles mark source events, scaled by moment magnitude.

Table 6. Sign stability of ΔCFS across 25 parameter combinations, by signal strength, with the corresponding rank correlations against the reference case.

Signal strength	Segments	Sign-stable	Median spread across cases
> 1 kPa	4,339	74.3 %	6.1 kPa
> 5 kPa	2,091	82.7 %	16.7 kPa
> 10 kPa	1,379	85.2 %	26.7 kPa
> 50 kPa	362	89.5 %	95.5 kPa

Spearman rho against the reference case across all 25 cases: minimum 0.798, median 0.921.

Regional outcome

Region	Minimum	Reference	Maximum	Loaded in all 25 cases
Marmara / Prince Islands	+0.123	+0.251	+0.381	yes

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Region	Minimum	Reference	Maximum	Loaded in all 25 cases
East Anatolian Fault, Surgu-Malatya	+0.135	+0.152	+0.167	yes
North Anatolian Fault, Akyazi-Sapanca	-0.563	-0.124	+0.235	**no** - not reported
East Anatolian Fault, Amanos	-0.126	-0.035	+0.086	**no** - not reported

Applied regionally, two of the four regions we would otherwise have reported do not survive the screen; at the segment level the corresponding loss is smaller, 14.8 % of segments above 10 kPa being sign-unstable. The Marmara and Prince Islands segments are loaded in all 25 cases (+0.123 to +0.381 MPa), as are the East Anatolian Fault segments in the Sürgü–Malatya corridor (+0.135 to +0.167 MPa). By contrast, the North Anatolian segments between the İzmit and Düzce ruptures (−0.563 to +0.235 MPa) and the Amanos segments (−0.126 to +0.086 MPa) change sign within the parameter space and are therefore not reported as results. Sensitivity to $A\sigma$ over 0.01–0.10 MPa alters the number of segments exceeding a gain threshold but not which segments are loaded, because the sign is fixed before the rate-and-state step; the Marmara probability gain varies only between 2.10 and 2.15 across that range.

4.4 Time-dependent probability change

For the Marmara segments, which carry a mapped slip rate of 21.4 mm yr^{-1} and hence an aftershock duration t_a of 5.5 years, within the 5–10 year range used in published Turkish studies, the 1999 İzmit earthquake advanced the failure clock by 34.5 years at 29.00° E, with a 30-year probability gain of 2.15. The largest advances lie at the eastern end of the window, nearest the 1999 rupture, falling to 25.6 years and 1.85 at 28.95° E and to 4.9 years and 1.16 at 28.55° E, although the decline is not monotonic with longitude because segment orientation also varies (Table 7). Of the 20 segments in the window, 18 are loaded; 16 of those pass the interpretability gate of Sect. 3.5 and are listed, and 2 fail it.

Table 7. Clock advance and 30-year probability gain for the Sea of Marmara segments following the 1999 İzmit earthquake.

Slip rate 21.4 mm/yr; t_a = 5.5 yr; A-sigma = 0.04 MPa; recurrence 70 yr. Every loaded segment in the window that passes the interpretability gate of Section 3.5 is listed. Two further loaded segments nearer the 1999 rupture fail that gate, carrying clock advances of 106 % and 70 % of a recurrence cycle, and are excluded.

Longitude (deg E)	dCFS (MPa)	Clock advance (yr)	30-yr probability gain
29.001	+0.251	+34.5	2.15
28.950	+0.186	+25.6	1.85

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Longitude (deg E)	dCFS (MPa)	Clock advance (yr)	30-yr probability gain
28.897	+0.054	+7.4	1.25
28.547	+0.035	+4.9	1.16
28.840	+0.033	+4.5	1.15
27.878	+0.033	+5.6	1.18
28.606	+0.023	+3.1	1.10
27.819	+0.022	+3.8	1.12
27.759	+0.017	+3.0	1.10
27.700	+0.014	+2.3	1.08
27.640	+0.011	+1.9	1.06
28.664	+0.009	+1.3	1.04
27.581	+0.009	+1.5	1.05
27.522	+0.006	+1.0	1.03
27.465	+0.003	+0.4	1.01
27.409	+0.002	+0.3	1.01

Two segments closer to the 1999 rupture carry larger stress changes, and they are deliberately not reported as results. At 29.10° E and 29.05° E the implied clock advances are 74.2 and 48.9 years against a recurrence interval of 70 years, that is 106 % and 70 % of a seismic cycle. These exceed the interpretability gate of Sect. 3.5, which excludes any segment whose clock advance passes half a cycle, because a stress step of that size relative to the cycle lies outside the regime where the elastic solution is dependable. Applying the gate consistently costs the paper its most striking single number, and that is the point of having the gate.

For comparison, and consistent with the earlier recalculation of Parsons (2004), Murru et al. (2016) reported an aggregated 30-year probability of M > 7.3 near Istanbul rising from 35 % to 47 %, a gain of 1.34; our per-segment gains of 1.16–2.15 bracket that value.

A physically instructive contrast emerges. Absolute ΔCFS is greatest in the Sea of Marmara, but the largest relative probability gains occur on slow faults surrounding the 2023 Kahramanmaraş ruptures, because the same stress step represents a far larger fraction of a slow seismic cycle.

4.5 Pseudo-prospective forecast skill

The sign of ΔCFS predicts the subsequent seismicity-rate change (Fig. 5, Table 8). The p values in Table 8a and 8b are nominal segment-level values, reported for comparability with the literature; they are not corrected for clustering and should

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

not be read as evidence. The event-clustered values in Table 8d are the ones that carry inference. Pooled over all sources, loaded segments gain activity relative to shadowed segments with AUC = 0.594. Two levels of inference are reported, because segment-event pairs are not independent: neighbouring sub-segments around one source share both that source and much of the subsequent seismicity. At the segment level the association is nominally overwhelming ($p = 8.4 \times 10{-}9$), but the sample that enters the AUC is far smaller and far more clustered than that p value assumes. Computed on the thresholded sample that carries the headline estimate, the 447 Mw < 7 segment-event pairs drawn from 13 events (the all-source thresholded sample is 1,213 pairs), the intra-event correlation is 0.289 over a mean cluster size of 34, giving a design effect of 10.7 and an **effective sample size of 42**. **All significance statements below therefore use the source event as the independent unit**, with a block bootstrap and a permutation test that resample or relabel whole events. On that basis the pooled association gives $p = 0.029$. The result is not an artefact of the metric: permuting ΔCFS within each event gives a null AUC of 0.4996 ± 0.0159, and the observed value lies outside the sampled null distribution ($p < 0.0033$ with 300 permutations and the standard +1 correction).

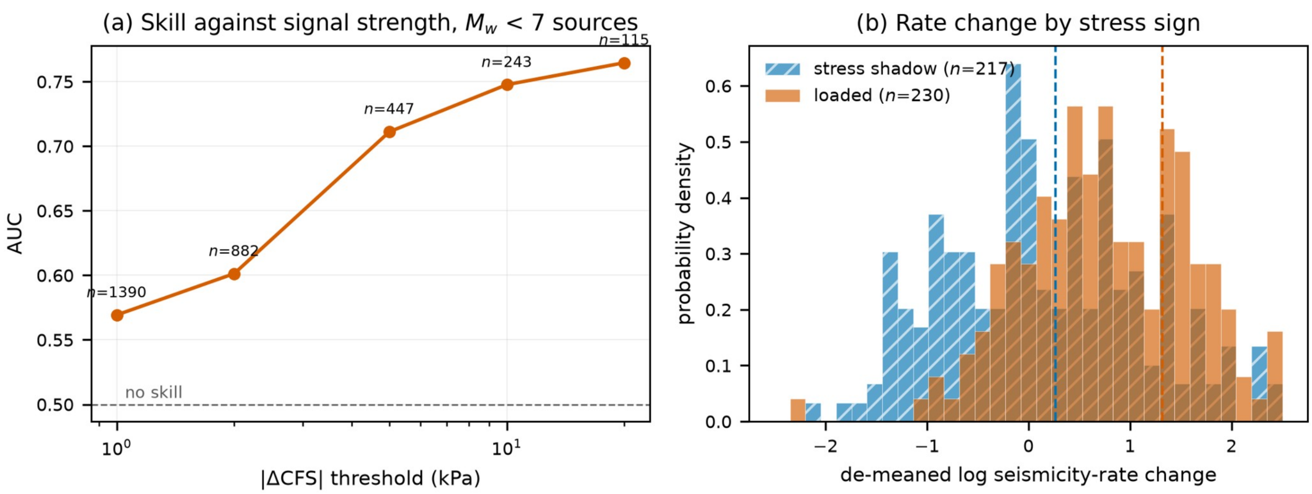

Figure 5. Prospective skill. Left: AUC against the |ΔCFS| threshold for Mw < 7 sources, with sample sizes annotated. Right: distributions of de-meaned log seismicity-rate change for loaded and shadowed segments, with medians marked.

Table 8. Pseudo-prospective forecast skill, stratified by source magnitude and by robustness flag, with the permutation null distribution.

(a) By source magnitude and robustness flag (Signal strength > 5 kPa)

Source subset	Segments	n loaded	n shadowed	AUC	Nominal p
All sources	all	635	578	0.594	8.4e-09
All sources	sign-stable	563	481	0.597	3.3e-08
All sources	sign-unstable	72	97	0.579	4.0e-02

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Source subset	Segments	n loaded	n shadowed	AUC	Nominal p
Mw < 7	all	230	217	0.711	5.7e-15
Mw < 7	sign-stable	208	194	0.728	1.5e-15
Mw < 7	sign-unstable	22	23	0.567	2.2e-01
Mw >= 7	all	405	361	0.524	1.3e-01
Mw >= 7	sign-stable	355	287	0.512	3.1e-01
Mw >= 7	sign-unstable	50	74	0.595	3.7e-02

(b) Skill against signal strength (Mw < 7 sources)

Threshold	n loaded	n shadowed	AUC	Nominal p
> 1 kPa	616	774	0.569	4.4e-06
> 2 kPa	426	456	0.601	9.7e-08
> 5 kPa	230	217	0.711	5.7e-15
> 10 kPa	118	125	0.748	1.3e-11
> 20 kPa	53	62	0.765	5.4e-07

(c) Permutation null (300 within-event shuffles of dCFS, Signal strength > 5 kPa)

Quantity	Value
Observed AUC	0.594
Null AUC, mean +/- SD	0.4996 +/- 0.0159
Null 95 % range	0.4687 - 0.5302
Empirical p (300 permutations, +1 correction)	< 0.0033

(d) Event-clustered inference (independent unit = source event, n = 16)

Subset	AUC	Block-bootstrap 95 % CI	Event-clustered p
All sources, segment-weighted	0.594	0.519 - 0.825	0.029
Mw < 7, segment-weighted	0.711	0.495 - 0.824 percentile; 0.542 - 0.850 BCa	0.043
Mw < 7, **event-equal**	**0.589 +/- 0.112**	(SEM over 6 scorable events)	-

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Subset	AUC	Block-bootstrap 95 % CI	Event-clustered p
Mw < 7, without the 2020 Samos event	0.593	leave-one-out jackknife	-
Mw < 7, sign-stable	0.728	0.530 - 0.835 (excludes 0.5)	0.018

(e) Distance control and paired bootstrap

Test	Result
Median distance to rupture, loaded segments	34.1 km
Median distance to rupture, shadowed segments	43.7 km
Distance-stratified AUC (10 km bands, Mw < 7)	0.655
Bands with AUC > 0.5	7 of 8
Event-clustered 95 % CI on the stratified AUC	0.419 - 0.785 (includes 0.5)
Screening gain, paired bootstrap over events	+0.016 [+0.001, +0.051], p = 0.020

Per-event mean AUC is 0.509 +/- 0.055 (SEM), with 5 of 8 testable events above 0.5 (sign test p = 0.36). The design-effect and single-event figures behind these numbers are given in Sect. 4.5 and are not repeated here.

The pooled AUC also weights events by their segment count, and that weighting matters here. The largest event in the subset, the 2020 Samos earthquake at Mw 6.99 in the GCMT catalogue and therefore below the strict Mw < 7 boundary used throughout, supplies 38 % of the thresholded sample, and a leave-one-event-out jackknife shows the pooled estimate depends on it: removing that one event lowers the pooled AUC from 0.711 to **0.593**, whereas removing any other event moves it by less than 0.04. Scoring each event separately and averaging without weighting gives an **event-equal AUC of 0.589 ± 0.112** over the six events with enough segments on both sides. The weighted and unweighted estimates differ by 0.12, and the unweighted one is the more conservative summary of what 13 events support.

Skill separates sharply by source magnitude. For Mw < 7 sources, AUC = 0.711 (event-clustered p = 0.043); for Mw ≥ 7 sources, AUC = 0.524 (p = 0.13, not significant). Within the Mw < 7 subset, excluding the 2023 sequence raises the estimate to 0.776; this is a property of that subset and does not bear on the Mw ≥ 7 result. This is consistent with the amplitude deficiency noted in Sect. 3.3, and the natural reading is that the uniform-slip rectangular approximation is adequate for moderate events and inadequate for the largest ruptures. We treat that reading as an interpretation to be tested rather than an established result, and state it as a falsifiable prediction in Sect. 5.3.

Skill grows monotonically with signal strength. For Mw < 7 sources, AUC rises from 0.569 at a 1 kPa threshold through 0.711 at 5 kPa and 0.748 at 10 kPa to 0.765 at 20 kPa.

The distance control matters, and it tempers the result. Loaded segments lie closer to the rupture than shadowed ones, at a median of 34.1 km against 43.7 km, so part of the raw association could be proximity rather than stress sign. Matching the two groups within 10 km distance bands lowers the estimate from 0.711 to a stratified **0.655**, with 7 of 8 bands above 0.5. The point estimate remains above 0.5 after matching on distance, which suggests that proximity may not explain the whole

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

association; the event-clustered interval on the stratified estimate is [0.419, 0.785] and includes 0.5 ($p = 0.14$), so an independent contribution of stress sign is **not established**. The cautious summary is that the sign of ΔCFS shows suggestive discrimination beyond the raw proximity association, that part of the raw association is certainly geometric, and that 16 source events are too few to establish the residual effect at conventional significance.

Most importantly, robustness screening improves skill, and it does so in the way that matters under clustered inference. For Mw < 7 sources the unscreened field gives AUC = 0.711 and the screened field AUC = 0.728, a modest gain of 0.017; the sign-unstable segments that screening discards carry AUC = 0.567 and are not distinguishable from chance ($p = 0.22$). Whether that gain is real is a question about the difference between two AUC values, and it is tested directly rather than by comparing where each confidence interval falls. A paired bootstrap over whole events gives a gain of +0.016 with a 95 % interval of [+0.001, +0.051] and a one-sided p of 0.020, so the improvement is significant, though the lower bound is close to zero and the effect is correspondingly modest.

Two procedural notes are needed to read these numbers. With 16 clusters a percentile bootstrap is not reliable, so bias-corrected and accelerated intervals are also reported: for the pooled Mw < 7 estimate the percentile interval is [0.495, 0.824] and the BCa interval is [0.542, 0.850]. The two differ because the bootstrap distribution is skewed. Second, the interval and the permutation p answer different questions and are computed by different procedures, one resampling events and the other relabelling them, so they may disagree without either being wrong.

4.6 Hazard and siting

The full hazard model reproduces the ordering of the national reference map. Against the AFAD TBDY-2018 DD-2 475-year reference peak ground acceleration across 18 cities, the fault-only model attains Spearman $\rho = 0.830$ and the full model, including smoothed background seismicity, attains $\rho = 0.865$. Internal consistency supports the background component: its share of the total is largest exactly where mapped faults are sparse, at 56.7 % in Gaziantep, 55.8 % in Konya and 54.3 % in Van, and smallest where they are dense, at Düzce 2.8 %.

The siting analysis (Fig. 6) ranks Trabzon most and Düzce least suitable (0.802 against 0.158). Istanbul scores 0.298, and the decomposition shows that this is driven by ground conditions rather than by fault proximity: VS30 of 205 m s^{-1} and a liquefaction susceptibility index of 0.92. Of the 80.7 million people mapped by WorldPop, 17.90 million (22.2 %) live in the highest-risk decile.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

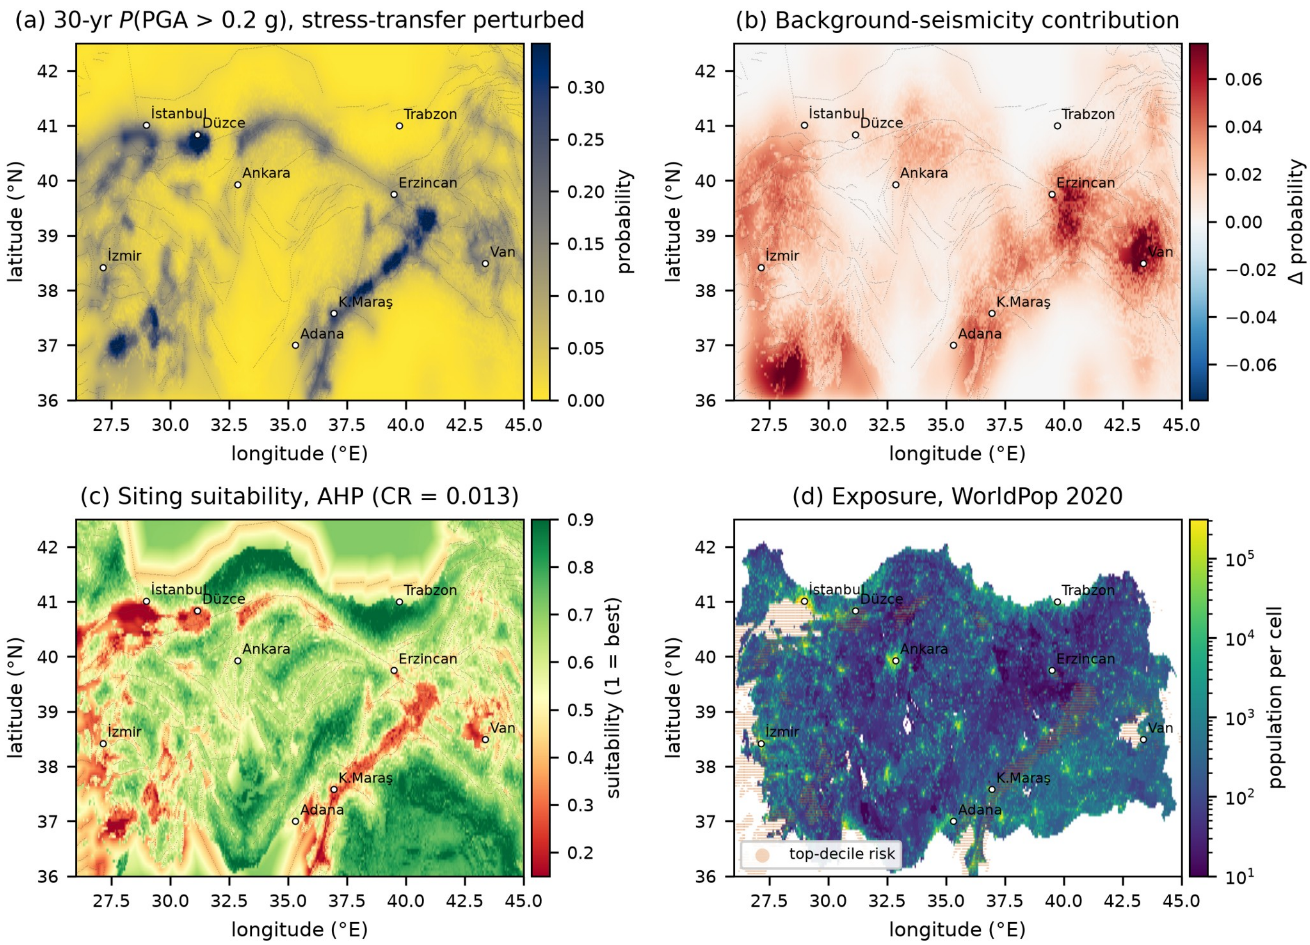

Figure 6. (a) Time-dependent 30-year probability of exceeding 0.2 g. (b) Contribution of smoothed background seismicity. (c) Siting suitability from the analytic hierarchy process. (d) Population exposure with the highest-risk decile overlaid.

The siting ranking is insensitive to the analytic-hierarchy weights. Perturbing every pairwise comparison log-normally (±50 % at 95 % confidence) over 1,000 draws leaves all draws admissible (CR < 0.10) and yields a city-ranking Spearman ρ against the reference of median 0.988 and minimum 0.939, exceeding 0.9 in 100 % of draws. Absolute hazard probabilities are, by contrast, sensitive to the ground-motion model, varying by a factor of 1.64 (Akkar et al., 2014) to 2.00 (Boore et al., 2014) relative to Kale et al. (2015), although the spatial pattern is preserved (ρ = 0.985 and 0.964 respectively).

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

5. Discussion

5.1 What survives automation

Taken together, the results give a quantitative answer to the question posed in Sect. 1. An automated regional Coulomb calculation carries genuine information, since the sign of ΔCFS predicts where seismicity goes with AUC up to 0.765, but that information is conditional upon three things that are not routinely reported.

The first is input integrity. A defect affecting 44 of the 2,076 records that can be tested internally, that is 2.1 %, sounds negligible until one of them is the most consequential fault in a national hazard model. The defect is invisible to the range checks that the database build performs, and would be invisible to any user who took the field at face value. It is, however, trivially detectable by a self-consistency test that requires no geological input, which is our recommendation to the maintainers.

The second is source placement. The finding that map position dominates focal mechanism runs against the emphasis of the existing sensitivity literature, which has concentrated on receiver geometry (Wang et al., 2014; Mildon et al., 2016). The discrepancy is explicable: those studies addressed single events with expert-curated source models, in which position was never in question. Automation reintroduces position as a free parameter, and across the 27 testable events it is the larger of the two uncertainty sources we evaluate. That location error can reverse the sign of a small ΔCFS is known; what we add is that the effect is systematic rather than random, applies to $M_w \geq 6$ sources, and reverses signs an order of magnitude above the conventional 0.01 MPa triggering threshold.

The third is parameter robustness. Two of the four regions we might have reported did not survive parameter variation, including one that carried the largest positive stress change in our first pass; at segment level the screen removes 14.8 % of the signals above 10 kPa. Reporting such a region would have been defensible under current practice, because current practice publishes a single map at a single value of μ'.

5.2 A proposal for reporting

We therefore propose that a regional Coulomb map be published together with a per-segment robustness classification, and that results below a stated stability threshold be marked as uninterpretable rather than shown on the same colour scale as robust results. The proposal is not merely conservative: Sect. 4.5 shows that screening improves discrimination against subsequent seismicity by +0.016 in AUC, with a paired-bootstrap interval of [+0.001, +0.051]. The larger-looking contrast against the discarded unstable segments is not the relevant comparison, for the reason given in Sect. 5.4. Screening therefore improves the product, and its cost is negligible, since the stress tensor need be computed only once and every parameter case resolved from it.

5.3 Where the method fails

The magnitude stratification of Sect. 4.5 is a limitation of the approach and should be stated as such. For $M_w \geq 7$ sources the uniform-slip rectangular chain carries no measurable skill (AUC = 0.524, $p = 0.13$). The evidence is the forward-in-time test itself, together with the moment-versus-amplitude argument of Sect. 5.4. The natural inference is that automated regional Coulomb modelling suits moderate events and needs published finite-fault slip models for the largest ones.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

We state this as an explicit and falsifiable prediction rather than a supposition: **substituting a published finite-fault slip distribution for the uniform-slip chain should raise the Mw ≥ 7 skill above its present AUC of 0.524**. If it does not, the explanation offered here is wrong and the failure must lie elsewhere, most plausibly in the near-field mask, which removes a far larger share of the loaded area for a 345 km rupture than for a 45 km one. The 1999 İzmit and 2023 Kahramanmaraş sequences both have published slip models and are the natural test. We have not performed that test here, and until it is done the mechanism behind the magnitude stratification remains an interpretation rather than a result.

5.4 Limitations

Several limitations qualify the results and are stated plainly.

The forward-in-time skill result rests on 16 independent source events, and the honest reading is that it is marginal rather than overwhelming. Under event-clustered inference the pooled association gives $p = 0.029$, the Mw < 7 subset $p = 0.043$ with a bootstrap interval that includes 0.5, and only the screened subset returns an interval that excludes it ($p = 0.018$). At the individual-event level the mean AUC is 0.509 ± 0.055, five of eight testable events exceed 0.5, and the sign test returns $p = 0.36$. Sixteen events have enough catalogue on both sides of the source to enter the pooled test; only eight retain enough loaded and shadowed segments above threshold to be scored individually, which is why the two counts differ. The distance control in Sect. 4.5 is the more important qualification: once proximity is matched out, the stratified estimate of 0.655 no longer reaches significance under event clustering. We report the result because the point estimate survives the control and the screening gain is separately significant, not because the evidence is strong. Per-event significance is therefore not established, and the pooled result should be read as a modest event-clustered effect rather than a strong one. One event, the 2023 Hatay Mw 6.3, returns AUC = 0.304, in the wrong direction; it occurred inside the Kahramanmaraş aftershock zone, where the pre-event baseline is contaminated.

The robustness-screening gain is +0.016 with a paired-bootstrap interval of [+0.001, +0.051]. The lower bound sits close to zero, so the improvement is real but small. It should not be confused with the 0.161 difference between the screened set and the discarded unstable one, which is not a measure of what screening achieves: the unstable set is discarded precisely because it is uninformative, so comparing against it is circular.

Absolute Coulomb amplitudes scale linearly with the seismic moment adopted. GCMT assigns the 1999 İzmit earthquake Mw 7.57, whereas Mw 7.4 is frequently quoted, a factor of two in moment; the +0.539 MPa we obtain on the single most loaded Prince Islands segment is correspondingly above the published 0.1–0.2 MPa. That figure is the peak on one segment and is not comparable with the +0.123 to +0.381 MPa quoted in Sect. 4.3, which are medians over the group under the extreme parameter cases. The spatial pattern is robust; the amplitude carries the moment uncertainty.

These two observations point in opposite directions and the implication should be stated rather than left to the reader. Section 3.3 reports that the modelled İzmit-to-Düzce transfer is 20–50 % below published values, and attributes the shortfall to uniform slip. Yet the calculation uses roughly twice the seismic moment that those published studies adopted. A model carrying twice the moment and still returning a smaller transfer must be losing more than 20–50 % to its geometry. The deficiency of the uniform-slip approximation is therefore larger than the raw comparison suggests, and this is independent support for the absence of skill at Mw ≥ 7 reported in Sect. 4.5.

The hazard model contains fault sources of length ≥ 20 km plus smoothed background seismicity, but no explicit treatment of offshore sources along the Cyprus arc; Adana remains under-predicted in consequence. The liquefaction layer is a transparent proxy built from VS30, slope and elevation, and requires local geotechnical validation before use in any specific siting decision. The VS30 field itself derives from a 1 km topographic proxy, which smooths coastal plains.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Finally, the magnitude conversion is calibrated over M 3.0–7.7, and the background model is restricted to $M_w \geq 4.0$ accordingly; extrapolation to the completeness magnitude of 1.5 was deliberately avoided.

6. Conclusions

1. The GEM Global Active Faults Database contains 44 records whose stored central rake contradicts the circular midpoint of its own uncertainty bounds. All lie in the EMME and SHARE compilations. One is the principal North Anatolian Fault strand south of Istanbul, where the error reverses the sign of the computed Coulomb stress change from −0.029 MPa to +0.251 MPa. A self-consistency check would detect every such record at no cost.
2. Correction must use the internal-bounds test, not the fault's slip_type label. The label test is confounded by inter-compilation convention differences, and the label can itself be wrong, as at the Malatya–Ovacık Fault Zone.
3. For automated regional modelling, the map position of the rupture dominates the result. Moving the 1999 İzmit rupture from the GCMT centroid to its mapped trace changes the transfer to Düzce from −0.054 to +0.177 MPa, whereas focal-mechanism variation changes it by 4 %; across 27 testable events a position perturbation flips the sign on a median of 11.9 % of segments against 2.0 % for a mechanism perturbation, and is the larger perturbation in 20 of 27 events (paired two-sided Wilcoxon $p = 5.4 \times 10^{-4}$).
4. Sign stability across 25 parameter combinations rises from 74.3 % to 89.5 % as $|\Delta CFS|$ grows from 1 kPa to 50 kPa. Half of our candidate regions failed the screen and are not reported.
5. The screened stress field carries pseudo-prospective skill against subsequent seismicity, but the evidence is suggestive rather than strong and should be read that way. The segment-weighted AUC of 0.711 for $M_w < 7$ sources rests on an effective sample of 42, is carried disproportionately by a single event, and falls to an event-equal 0.589 ± 0.112 and to a distance-stratified 0.655. What is robust is the comparison between screened and unscreened fields: +0.016 with a paired-bootstrap interval of [+0.001, +0.051]. The method carries no measurable skill at all for $M_w \geq 7$ sources under a uniform-slip approximation.
6. Propagated through rate-and-state probability, probabilistic hazard and multi-criteria analysis, the screened field yields a proof-of-concept downstream siting demonstration whose ranking is insensitive to the tested weight perturbations (Spearman $\rho \geq 0.939$ over 1,000 draws), while absolute hazard probabilities remain sensitive to the ground-motion model by a factor of two. It is intended to show that the audited stress field can be propagated, not to provide an operational planning product.

We recommend that regional Coulomb stress maps be published together with an explicit statement of which parts of them are interpretable.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Declarations

Code availability. Analysis scripts are archived in the same Zenodo deposit (Vural, 2026). Every table body is generated from the data by script, and the reference list is rendered directly from Crossref publisher records, so no numerical or bibliographic field in this manuscript was entered by hand. Coulomb calculations use the okada_ext extension of Pyrocko; ground-motion models use the OpenQuake hazard library (Pagani et al., 2014).

Data availability. All input data are public: the AFAD event service; the USGS and ISC FDSN services; the Global CMT catalogue; the GEM Global Active Faults Database; GMTED2010; and WorldPop. Derived products and analysis code are archived at Zenodo (Vural, 2026, https://doi.org/10.5281/zenodo.22134495), which mirrors the version used here and is available to reviewers during the discussion phase. Supplementary Table S1, listing all 44 affected GEM records with their catalogue identifiers, is included in that deposit in machine-readable form (TableS1_GEM_rake_defect_records.csv) so that the central claim of this paper can be checked record by record.

Author contribution. M.S.V. designed the study, developed the analysis code, performed all computations, and wrote the manuscript.

Competing interests. The author declares that they have no conflict of interest.

Acknowledgements. The author thanks AFAD, the Global CMT Project, the International Seismological Centre and the GEM Foundation for maintaining the open data on which this work depends.

The analysis code, the computations reported here, and the first draft of the text were produced with the assistance of a large language model (Claude, Anthropic) used as an interactive programming and drafting tool. The author defined the research question and scope, selected the target journal and the framing, directed every methodological choice, decided which findings were reported and which were withdrawn, and verified the results. All numerical values in the manuscript, its tables and its figures are generated directly from the input data by the scripts archived with this work, and a consistency checker verifies the manuscript text against those data. The author takes full responsibility for the content.

Financial support. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

References

AFAD: Earthquake event service, Disaster and Emergency Management Presidency of Türkiye [data set], https://deprem.afad.gov.tr (last access: 25 August 2026), 2026.

Akkar, S., Sandıkkaya, M. A., and Bommer, J. J.: Empirical ground-motion models for point- and extended-source crustal earthquake scenarios in Europe and the Middle East, B. Earthq. Eng., 12, 359-387, https://doi.org/10.1007/s10518-013-9461-4, 2014.

Akyüz, H. S., Hartleb, R. D., Barka, A. A., Altunel, E., Sunal, G., Meyer, B., and Armijo, R.: Surface Rupture and Slip Distribution of the 12 November 1999 Duzce Earthquake (M 7.1), North Anatolian Fault, Bolu, Turkey, B. Seismol. Soc. Am., 92, 61-66, https://doi.org/10.1785/0120000840, 2002.

Barka, A., Akyüz, H. S., Altunel, E., Sunal, G., Çakır, Z., Dikbaş, A., Yerli, B., Armijo, R., Meyer, B., de Chabalier, J. B., Rockwell, T., Dolan, J. R., Hartleb, R., Dawson, T., Christofferson, S., Tucker, A., Fumal, T., Langridge, R., Stenner, H., Lettis, W., Bachhuber, J., and Page, W.: The Surface Rupture and Slip Distribution of the 17 August 1999 Izmit Earthquake (M 7.4), North Anatolian Fault, B. Seismol. Soc. Am., 92, 43-60, https://doi.org/10.1785/0120000841, 2002.

Boore, D. M., Stewart, J. P., Seyhan, E., and Atkinson, G. M.: NGA-West2 Equations for Predicting PGA, PGV, and 5% Damped PSA for Shallow Crustal Earthquakes, Earthq. Spectra, 30, 1057-1085, https://doi.org/10.1193/070113eqs184m, 2014.

Cattania, C., Hainzl, S., Wang, L., Roth, F., and Enescu, B.: Propagation of Coulomb stress uncertainties in physics-based aftershock models, J. Geophys. Res.-Sol. Ea., 119, 7846-7864, https://doi.org/10.1002/2014jb011183, 2014.

Cocco, M., and Rice, J. R.: Pore pressure and poroelasticity effects in Coulomb stress analysis of earthquake interactions, J. Geophys. Res.-Sol. Ea., 107, https://doi.org/10.1029/2000jb000138, 2002.

Dieterich, J.: A constitutive law for rate of earthquake production and its application to earthquake clustering, J. Geophys. Res.-Sol. Ea., 99, 2601-2618, https://doi.org/10.1029/93jb02581, 1994.

Frankel, A.: Mapping Seismic Hazard in the Central and Eastern United States, Seismol. Res. Lett., 66, 8-21, https://doi.org/10.1785/gssrl.66.4.8, 1995.

Gardner, J. K., and Knopoff, L.: Is the sequence of earthquakes in Southern California, with aftershocks removed, Poissonian?, B. Seismol. Soc. Am., 64, 1363-1367, https://doi.org/10.1785/bssa0640051363, 1974.

GEM Foundation: GEM Global Active Faults Database, harmonised release, GEM Foundation [data set], https://github.com/GEMScienceTools/gem-global-active-faults (last access: 25 August 2026), 2026.

Global CMT Project: Global Centroid Moment Tensor catalogue, Global CMT Project [data set], https://www.globalcmt.org (last access: 25 August 2026), 2026.

Hardebeck, J. L., Nazareth, J. J., and Hauksson, E.: The static stress change triggering model: Constraints from two southern California aftershock sequences, J. Geophys. Res.-Sol. Ea., 103, 24427-24437, https://doi.org/10.1029/98jb00573, 1998.

Heimann, S., Kriegerowski, M., Isken, M., Cesca, S., Daout, S., Grigoli, F., Juretzek, C., Megies, T., Nooshiri, N., Steinberg, A., Sudhaus, H., Vasyura-Bathke, H., Willey, T., and Dahm, T.: Pyrocko - an open-source seismology toolbox and library, GFZ Data Services [code], https://doi.org/10.5880/GFZ.2.1.2017.001, 2017.

Hubert-Ferrari, A., Barka, A., Jacques, E., Nalbant, S. S., Meyer, B., Armijo, R., Tapponnier, P., and King, G. C. P.: Seismic hazard in the Marmara Sea region following the 17 August 1999 Izmit earthquake, Nature, 404, 269-273, https://doi.org/10.1038/35005054, 2000.

International Seismological Centre: On-line Bulletin, ISC [data set], https://doi.org/10.31905/D808B830 (last access: 25 August 2026), 2026.

Kale, Ö., Akkar, S., Ansari, A., and Hamzehloo, H.: A Ground-Motion Predictive Model for Iran and Turkey for Horizontal PGA, PGV, and 5% Damped Response Spectrum: Investigation of Possible Regional Effects, B. Seismol. Soc. Am., 105, 963-980, https://doi.org/10.1785/0120140134, 2015.

King, G. C. P., Stein, R. S., and Lin, J.: Static stress changes and the triggering of earthquakes, B. Seismol. Soc. Am., 84, 935-953, https://doi.org/10.1785/bssa0840030935, 1994.

https://doi.org/10.5194/egusphere-2026-5171
Preprint. Discussion started: 25 September 2026

Mancini, S., Segou, M., Werner, M. J., and Cattania, C.: Improving Physics-Based Aftershock Forecasts During the 2016-2017 Central Italy Earthquake Cascade, J. Geophys. Res.-Sol. Ea., 124, 8626-8643, https://doi.org/10.1029/2019jb017874, 2019.

Mildon, Z. K., Toda, S., Faure Walker, J. P., and Roberts, G. P.: Evaluating models of Coulomb stress transfer: Is variable fault geometry important?, Geophys. Res. Lett., 43, https://doi.org/10.1002/2016gl071128, 2016.

Mildon, Z. K., Roberts, G. P., Faure Walker, J. P., and Toda, S.: Coulomb pre-stress and fault bends are ignored yet vital factors for earthquake triggering and hazard, Nat. Commun., 10, https://doi.org/10.1038/s41467-019-10520-6, 2019.

Murru, M., Akinci, A., Falcone, G., Pucci, S., Console, R., and Parsons, T.: M ≥ 7 earthquake rupture forecast and time-dependent probability for the sea of Marmara region, Turkey, J. Geophys. Res.-Sol. Ea., 121, 2679-2707, https://doi.org/10.1002/2015jb012595, 2016.

Okada, Y.: Internal deformation due to shear and tensile faults in a half-space, B. Seismol. Soc. Am., 82, 1018-1040, https://doi.org/10.1785/bssa0820021018, 1992.

Pagani, M., Monelli, D., Weatherill, G., Danciu, L., Crowley, H., Silva, V., Henshaw, P., Butler, L., Nastasi, M., Panzeri, L., Simionato, M., and Vigano, D.: OpenQuake Engine: An Open Hazard (and Risk) Software for the Global Earthquake Model, Seismol. Res. Lett., 85, 692-702, https://doi.org/10.1785/0220130087, 2014.

Parsons, T., Toda, S., Stein, R. S., Barka, A., and Dieterich, J. H.: Heightened Odds of Large Earthquakes Near Istanbul: An Interaction-Based Probability Calculation, Science, 288, 661-665, https://doi.org/10.1126/science.288.5466.661, 2000.

Parsons, T.: Recalculated probability of M ≥ 7 earthquakes beneath the Sea of Marmara, Turkey, J. Geophys. Res.-Sol. Ea., 109, https://doi.org/10.1029/2003jb002667, 2004.

Reitman, N. G., Briggs, R. W., Barnhart, W. D., Hatem, A. E., Thompson Jobe, J. A., DuRoss, C. B., Gold, R. D., Mejstrik, J. D., Collett, C., Koehler, R. D., and Akçiz, S.: Rapid Surface Rupture Mapping from Satellite Data: The 2023 Kahramanmaraş, Turkey (Türkiye), Earthquake Sequence, Seism. Rec., 3, 289-298, https://doi.org/10.1785/0320230029, 2023.

Saaty, T. L.: The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation, McGraw-Hill, New York, 1980.

Sançar, T., Zabcı, C., Karabacak, V., Yazıcı, M., and Akyüz, H. S.: Geometry and Paleoseismology of the Malatya Fault (Malatya-Ovacık Fault Zone), Eastern Turkey: Implications for intraplate deformation of the Anatolian Scholle, J. Seismol., 23, 319-340, https://doi.org/10.1007/s10950-018-9808-z, 2019.

Schorlemmer, D., Werner, M. J., Marzocchi, W., Jordan, T. H., Ogata, Y., Jackson, D. D., Mak, S., Rhoades, D. A., Gerstenberger, M. C., Hirata, N., Liukis, M., and Maechling, P. J.: The Collaboratory for the Study of Earthquake Predictability: Achievements and Priorities, Seismol. Res. Lett., 89, 1305-1313, https://doi.org/10.1785/0220180053, 2018.

Steacy, S., Gomberg, J., and Cocco, M.: Introduction to special section: Stress transfer, earthquake triggering, and time-dependent seismic hazard, J. Geophys. Res.-Sol. Ea., 110, https://doi.org/10.1029/2005jb003692, 2005a.

Steacy, S., Nalbant, S. S., McCloskey, J., Nostro, C., Scotti, O., and Baumont, D.: Onto what planes should Coulomb stress perturbations be resolved?, J. Geophys. Res.-Sol. Ea., 110, https://doi.org/10.1029/2004jb003356, 2005b.

Stein, R. S.: The role of stress transfer in earthquake occurrence, Nature, 402, 605-609, https://doi.org/10.1038/45144, 1999.

Styron, R., and Pagani, M.: The GEM Global Active Faults Database, Earthq. Spectra, 36, 160-180, https://doi.org/10.1177/8755293020944182, 2020.

Tan, O.: A homogeneous earthquake catalogue for Turkey, Nat. Hazards Earth Syst. Sci., 21, 2059-2073, https://doi.org/10.5194/nhess-21-2059-2021, 2021.

Toda, S., Stein, R. S., Richards-Dinger, K., and Bozkurt, S. B.: Forecasting the evolution of seismicity in southern California: Animations built on earthquake stress transfer, J. Geophys. Res.-Sol. Ea., 110, https://doi.org/10.1029/2004jb003415, 2005.

US Geological Survey and National Geospatial-Intelligence Agency: Global Multi-resolution Terrain Elevation Data 2010 (GMTED2010), USGS [data set], https://doi.org/10.5066/F7J38R2N, 2010.

US Geological Survey: FDSN event web service, USGS [data set], https://earthquake.usgs.gov/fdsnws/event/1/ (last access: 25 August 2026), 2026.

Vural, M. S.: Data, code and supplementary materials for "How much of a regional Coulomb stress map survives its own uncertainty? An audit with forecast-skill validation for Türkiye", Zenodo [code], https://doi.org/10.5281/zenodo.22134495, 2026.

Wald, D. J., and Allen, T. I.: Topographic Slope as a Proxy for Seismic Site Conditions and Amplification, B. Seismol. Soc. Am., 97, 1379-1395, https://doi.org/10.1785/0120060267, 2007.

Wang, J., Xu, C., Freymueller, J. T., Li, Z., and Shen, W.: Sensitivity of Coulomb stress change to the parameters of the Coulomb failure model: A case study using the 2008 Mw 7.9 Wenchuan earthquake, J. Geophys. Res.-Sol. Ea., 119, 3371-3392, https://doi.org/10.1002/2012jb009860, 2014.

Wang, J., Xu, C., Freymueller, J. T., Wen, Y., and Xiao, Z.: AutoCoulomb: An Automated Configurable Program to Calculate Coulomb Stress Changes on Receiver Faults with Any Orientation and its Application to the 2020 Mw 7.8 Simeonof Island, Alaska, Earthquake, Seismol. Res. Lett., 92, 2591-2609, https://doi.org/10.1785/0220200283, 2021.

Wells, D. L., and Coppersmith, K. J.: New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement, B. Seismol. Soc. Am., 84, 974-1002, https://doi.org/10.1785/bssa0840040974, 1994.

Westaway, R., and Arger, J.: Kinematics of the Malatya-Ovacik Fault Zone, Geodin. Acta, 14, 103-131, https://doi.org/10.1080/09853111.2001.11432438, 2001.

Woessner, J., Hainzl, S., Marzocchi, W., Werner, M. J., Lombardi, A. M., Catalli, F., Enescu, B., Cocco, M., Gerstenberger, M. C., and Wiemer, S.: A retrospective comparative forecast test on the 1992 Landers sequence, Journal of Geophysical Research, 116, https://doi.org/10.1029/2010jb007846, 2011.

WorldPop: Global high-resolution population denominators, University of Southampton [data set], https://doi.org/10.5258/SOTON/WP00660, 2020.