Participation Is Not Protection: Sovereign Catastrophe Cover and Reserve Adequacy in MENAAP*

Ilker Aslan, PhD
Kalaidos University of Applied Sciences, Zürich

November 2026

Abstract. For a sovereign whose reconstruction is priced in a currency it does not issue, a catastrophe is a balance-of-payments event before it is a fiscal event. Reserves and parametric cover are then competing ways of meeting one foreign currency requirement, which makes the choice between them a reserve management question before it is an insurance one. No literature prices them against each other; this paper does, for earthquake as well as for flood, storm and drought, and asks how much of the observed behaviour the comparison explains. Cover is cheaper when the hazard frequency times the market's loading over expected loss falls below the carry cost of reserves: the loading, not the gross price, because a reserve unit spent in the disaster state is lost while the liability behind it stands. Neither basis risk nor fungibility rescues it. Estimating the pricing function on twenty-three sovereign catastrophe bond tranches puts the elasticity of the multiple in expected loss at -0.46 pooled and -0.43 with the cycle in the specification, while the level rises 55 per cent from the soft market of 2017–21 to the hard market since 2022; at a one-in-a-hundred-year frequency the threshold is 150 to 340 basis points soft and 300 to 600 hard.

Applied to the twenty MENAAP economies against a hand-built inventory of their risk financing, it explains little. Egypt, Pakistan, Jordan, Tunisia and Morocco each satisfied it with their standing to contract intact, at eleven observations surviving a cluster bootstrap in every draw, and none bought cover. Morocco, the only placement of macroeconomic consequence, met the condition in 2001 and renewed in 2023 at a spread far below it; its 2020 placement falls in a gap where no spread is observed. Nor does the obvious explanation survive: at the only documented fixed cost of access the entry condition is satisfied several times over, so the model predicts entry. Running it on the sovereigns that did buy settles nothing either way, three of the four sponsors lying outside the comparison's domain.

What follows is therefore a measurement agenda, not an instrument one, and it is addressed to central banks and to the external sector assessment rather than to the adaptation finance architecture. Reserve adequacy metrics carry no term for catastrophe exposure, so two sovereigns with identical ratios and different cover are recorded as equally protected. Three changes follow, none needing a new instrument, a new facility or a new fund: report pre-arranged cover as a memorandum item in the reserve adequacy assessment; credit it against a medium-term accumulation path rather than at par against a short-term reserve floor, conditional on settlement-timing disclosure; and report liquid external assets by holder and encumbrance, which the Omani case here shows surveillance cannot presently resolve.

Keywords: reserve adequacy; international reserves; disaster risk financing; adaptation finance; climate and seismic catastrophe risk; catastrophe bonds; parametric insurance; basis risk; sovereign spreads; MENAAP
JEL classification: F31, F34, G22, H84, Q54

*Lecturer, Kalaidos University of Applied Sciences, Zürich. ORCID 0000-0001-9241-9420. Correspondence: ilker.aslan@alumni.unifr.ch.

A data appendix is attached as a separate supplement. It sources the inventory of Table 5 cell by cell and the monetary-position figures of Table 6 line by line, documents the pricing sample of Table 2 and the spread readings of Tables 3 and 4b, distinguishes confirmed absences from documents that could not be retrieved, names the vintage of each series where vintages differ, and adjudicates every conflicting figure with the rule applied. Appendix A.11 lists the quantities in this paper that are our own computation rather than a published figure, together with the inputs and assumptions each rests on; Tables 1, 4, 4a and 4b and the arithmetic of Section 7.4 are computed objects and are listed there. References to Appendix A throughout are to that supplement.

1. Introduction

A catastrophe in an economy that cannot issue the currency its obligations are written in is a balance-of-payments event before it is a fiscal event. Reconstruction is import-intensive and immediate. The insurance recovery, where one exists, is largely denominated in foreign currency, because catastrophe-exposed property risk written by domestic insurers in these markets is overwhelmingly ceded onward. The qualification is line-specific rather than market-wide, and Section 2.2 makes it. Where the sovereign issues its own currency the resulting mismatch resolves through the monetary authority's balance sheet; where it does not, it is a claim on reserves.

That observation makes foreign exchange reserves and pre-arranged catastrophe cover competing ways of meeting one requirement. Neither of the two literatures that study them treats them that way. The reserve adequacy literature prices reserves carefully against the carry cost and contains no disaster state. The sovereign disaster risk financing literature treats cover as a fiscal instrument benchmarked against a contingent liability, and the leading treatment of optimal sovereign insurance states explicitly that it sets self-insurance aside. Neither asks whether a dollar of cover is cheaper than a dollar of reserves held against the same contingency. The applied literature does compare insurance against alternative financing, Cebotari and Youssef (2020) and Borensztein, Cavallo and Valenzuela (2009) among others, and the sovereign risk pool literature prices multiples against sovereign borrowing costs as a matter of routine. What is missing is the comparison against the carry cost of the reserve stock itself, as an instrument the sovereign is already holding for the same purpose.

This paper asks it, and the answer is largely negative.

Section 3 sets the two instruments on one balance sheet. Cover is the cheaper way to deliver foreign currency in the disaster state when the hazard frequency, multiplied by the market's loading over expected loss, falls below the carry cost of reserves — the loading rather than the gross price, because a reserve unit spent in the disaster state is lost while the liability behind it stands. Two features that an informal treatment disposes of in a line turn out to bite, and both against cover. Basis risk does not cancel: a weaker trigger buys a cheaper contract, but the market prices lower expected loss at a higher multiple, so the cost per unit of expected delivery rises. And the fungibility of reserves nets to zero only at an interior optimum, so where a buffer is thin or a reserve floor binds — which is wherever the question is live here — the credit is positive and reserves are more attractive still. What matters is not the form of the condition, which is the one an underwriter would write down, but that it survives both corrections and that its right-hand side is a sovereign spread rather than a budget line.

Section 4 inverts the condition. Calibrated on twenty-three sovereign catastrophe bond tranches issued since 2017, none in the region, substitution on a remote parametric layer requires a spread of at least 150 to 340 basis points at soft-market prices and 300 to 600 at hard-market prices, taken at a one-in-a-hundred-year frequency and a perfect trigger, and more at any realistic one. Two things follow. The threshold is not a number: the fitted price of catastrophe risk rises about 55 per cent between the regimes, moving it from 193 to 354 basis points, so a ministry that satisfied the condition at one renewal can fail it at the next with nothing changed on its

own balance sheet. And price is only one of three conditions a sovereign must meet to transfer risk, the others being accessibility and the standing to contract; Section 4.1 states them as a definition and the circumstances that raise a spread often impair standing, so one might expect the three rarely to hold together; Section 4.3 finds that they held together at eleven dated observations, which is what makes the non-purchases a puzzle rather than an artefact.

Sections 5 and 6 bring the evidence. We assembled an inventory of sovereign catastrophe risk financing for the twenty economies of the World Bank's MENAAP region, from national legislation, regulator publications, project documents, staff reports and the catastrophe bond issuance record, sourced in Appendix A. The region holds two placed contracts and one of macroeconomic consequence: Morocco's parametric earthquake cover at a limit of USD 275 million, which the World Bank records as having disbursed in full after the 2023 Al Haouz earthquake, and Djibouti's multi-peril policy with the African Risk Capacity at a combined limit of USD 2.2 million, a limit too small to move a balance of payments.

Setting the inventory against the band produces the paper's sharpest finding, and it is not the finding the framework was built to produce. Table 4 applies the three conditions by a rule stated once and applied to every dated observation without exception. Egypt, Pakistan, Jordan, Tunisia and Morocco each satisfied all three at one or more observations between 2001 and the end of 2021 while quoted in international markets with their contracting standing intact: eleven observations on the strict screen set out in the note to Table 4, or seventeen if Egypt's crisis-period option-adjusted spreads and its post-2011 index readings are included. None bought cover at any of those dates. Because the threshold depends on a hazard frequency and a trigger reliability that nobody publishes, Table 4a reports the classification across both: five economies at the reference parameters, never fewer than two anywhere on the surface, and Morocco's placement dates failing throughout.

That result constrains what the rest of the paper can claim. The natural reading of those non-purchases is that they were crisis episodes, in which the spread that cleared the threshold also recorded the circumstances foreclosing the transaction. The dates do not support it: on the strict screen every qualifying reading is an ordinary market observation, and the crisis-period readings are precisely the ones the screen removes. The price condition therefore does not explain the region's absence of sovereign catastrophe cover. Where it could be evaluated and was satisfied, nothing was bought.

Morocco is the sharpest version of this, with one qualification the paper is careful about. It met all three conditions in 2001, at a spread of 518 basis points. Its 2020 placement falls inside the 2015 to 2021 window for which no spread observation exists for any economy here, so that date cannot be judged and is not judged. Its 2023 renewal can be, at 109 basis points against a hard-market floor of 300, and every other reading in the surrounding period runs between 73 and 292, below the applicable range throughout. Djibouti, which a price-only reading places deepest inside the band, is on the argument of Section 3.2 outside the domain of the comparison: it has no market access, so neither instrument is available at the margin and Proposition 1 does not apply. It bought anyway, because donors paid, at a limit of five hundredths of one per cent of GDP.

The first question a reader will put to that finding is whether the sovereigns that did buy were inside the band when they bought, and Section 4.4 runs the screen on the eight remote parametric layers in the calibration sample. It reports an attempt rather than a result. Two of the seven placements with a sourced spread clear the reference threshold, one is indeterminate and four do not, but three of the four sponsors are inflation-targeting floaters for which Section 3.1's premise does not hold and the participation condition fails before the price condition is reached; the one sponsor plainly inside the domain has a single evaluable placement. Ranked by spread the observations separate almost monotonically in the direction Proposition 1 implies, which is weak evidence for the condition's ordering rather than against it, and every negative verdict reverses at a reserve-drawdown probability the paper concedes it cannot estimate. We

report the exercise because the question is unavoidable and because what the available data cannot identify is itself worth stating. It leaves Section 4.3's finding standing as a regional one.

Six sovereigns could not have transacted at all, a seventh, Iraq, only marginally, and four hold buffers ample enough that *cover* is worth little to them, which is not the same as saying that protection is: for those four, reserves are the cheap instrument and they hold them in quantity. Bahrain and Oman are in neither group. Bahrain, which income and peg alone would place among the ample-buffer sovereigns, holds the thinnest import cover in the Gulf at 2.4 months with government debt at 134 per cent of GDP; its absence of cover is nonetheless what the price condition predicts once its post-2021 spreads are judged against the hard-market calibration rather than the soft one. The reading that follows relocates the problem rather than reducing it. For the four Gulf states with ample buffers the absence of cover is what the pricing condition predicts. For the five economies that satisfied the condition and did not act on it, it is not. The obvious candidate, the fixed cost of market access, can be tested against the entry condition and fails: at the only fixed-cost figure on the record the condition is satisfied by a factor of three to forty, so the model predicts entry. The remaining five candidates in Section 3.4 — unplaceable risk, religious objection to conventional insurance and securitisation, charity hazard, the budgetary salience of a premium paid against a contingent benefit, and weakened incentives for ex ante risk reduction — are untested, because the data that would separate them are not published.

Section 7 draws the consequence for policy. The disaster risk financing agenda records which sovereigns lack instruments, builds programmes to supply them, and counts them in its results frameworks. Those metrics count participation. Protection is a different quantity and the agenda does not measure it. Four of the six Gulf states appear in regional gap statistics as underprotected when on the evidence here they are correctly uninsured; the other two, Bahrain and Oman, are not ample-buffer sovereigns and should never have been counted as such. Djibouti appears as a success, holding a parametric policy and a pool membership whose premium was paid by donors, at a limit too small to move its balance of payments. Pakistan, the one sovereign in the region whose external position a catastrophe actually overwhelmed, is the case where the economics most clearly favoured transfer and where the transaction was least available. The instrument's distribution across the region does not match where the pricing condition says it is worth holding. The household protection gap is large and nothing here bears on it.

Three changes follow for surveillance, and Section 7.3 sets them out. Pre-arranged catastrophe cover should be reported as a memorandum item in the reserve adequacy assessment, in the external sector annex of the Article IV staff report and in the selected-indicators table, with notional, trigger, term and counterparty. It should be credited against a medium-term reserve accumulation path rather than at par against a short-term net international reserve floor, and settlement-timing disclosure should be the condition of that recognition, since no contract in this region publishes one. And liquid external assets should be reported by holder and by encumbrance: Section 5.3 could not place Oman's buffer because official reserves, the investment authority's total assets and the Fund's liquid external measure give three different answers and the second rests on a press release rather than a published statistical return. None of the three requires a new instrument, a new facility or a new fund.

The evidence is institutional and descriptive, and with two placed contracts, only one of them of macroeconomic consequence, no causal claim is available. The shadow value governing whether protection is worth holding at all is assumed, not derived, so it orders sovereigns only qualitatively. No continuous sovereign spread series spanning the whole period exists in the public record for these economies, so Table 3 assembles two partial sources in separate panels, with a third set of point observations from staff reports recorded in Table 4 and Appendix A.5 rather than tabulated, and the comparison of threshold against actual is made at dated observations rather than against a continuous series. The transactability condition is coarse, a yes or no rather than a price of access.

The paper proceeds as follows. Section 2 positions the argument against the two literatures. Section 3 derives the substitution condition and the entry condition. Section 4 inverts the first into threshold spreads. Section 5 presents the inventory. Section 6 develops the Pakistan, Morocco and Türkiye cases. Section 7 draws the policy implications. Section 8 concludes.

2. Two literatures that do not meet

2.1 Reserves as self-insurance

The modern treatment of reserve accumulation is precautionary. Reserves buy insurance against a sudden stop in external financing, and they are costly: the social cost of a unit of reserves is conventionally measured as the spread between the sovereign's external borrowing rate and the yield on the reserve asset (Rodrik, 2006). Optimal reserve models equate the marginal cost of carry against the marginal benefit in reduced crisis probability and reduced output loss conditional on crisis (Jeanne and Rancière, 2011), and the precautionary motive has been distinguished empirically from the mercantilist one (Aizenman and Lee, 2007). Subsequent work has set reserves in the trilemma (Obstfeld, Shambaugh and Taylor, 2010), identified their macroprudential function (Jeanne, 2016), and embedded accumulation in a sovereign default model where reserves and debt are held simultaneously because reserves stabilise the economy and reduce rollover risk (Bianchi and Sosa-Padilla, 2024).

What none of this literature contains is a disaster state. Its shock is financial: a reversal of capital flows, a rollover failure, a run. The possibility that the economy sustains a large real shock generating a foreign currency requirement of its own, instead of a withdrawal of foreign currency supply, sits outside the standard setup, even though the balance-of-payments consequence is similar and the instrument used to meet it is identical.

2.2 Disasters, insurance and macroeconomic adjustment

A parallel literature studies the macroeconomics of disasters, which impose measurable and persistent costs on output (Noy, 2009; Cavallo *et al.*, 2013; Hsiang and Jina, 2014). Within it, two strands bear on the question here. The first asks whether the exchange rate regime conditions adjustment to real shocks, using disasters as plausibly exogenous variation, and finds that flexible arrangements adjust faster (Ramcharan, 2007); recent work confirms this with panel local projections on 177 countries, estimating depreciation of the nominal effective exchange rate of up to 15 per cent over the twenty-four months following a large event in non-pegged emerging and developing economies (Nguyen and Nguyen, 2025). The second, and the closer antecedent to this paper, examines whether *insurance* alters the macroeconomic cost of a disaster, and finds that it is the uninsured portion of losses that drives the output cost: adequately insured economies recover, while uninsured ones do not (von Peter, von Dahlen and Saxena, 2024).

Whether an insurance recovery in these markets arrives in foreign currency is an empirical question and the answer is line-specific. Regional aggregates suggest high retention: Moroccan direct insurers ceded about 14 per cent of non-life premium in 2023, a ratio we compute ourselves from the regulator's published amounts, the regulator publishing none, and the MENA Reinsurance Barometer, in editions dating from 2014 and 2015, put regional non-life retention near 70 per cent. Those aggregates are dominated by motor and medical business, and the catastrophe-exposed lines run the other way, which the regulators' own line-level tables show directly. Saudi Arabia's regulator reports property and fire retention of 17.2 per cent in 2020 and engineering retention of 22.3 per cent, against a market aggregate of 82.5 per cent. Bahrain's central bank reports fire, property and liability retention of 11 per cent in 2024 and engineering retention of 7 per cent. The IAIS assessment of Oman records that about half of total premiums are ceded and that most reinsurance is placed with the international market or within the group. Two qualifications attach. Local-cession requirements in several markets keep part of the ceded

premium onshore. The only compilation of these rules we located dates from 2014, giving a 30 per cent Saudi rule and a 50 per cent Algerian cession to the state reinsurer; the Saudi rule has since been revised and we were unable to verify the current position for any market in the region, so these figures establish that such rules exist rather than their present level. And no regulator in the region publishes the currency of settlement, so the final step from cession abroad to a foreign currency recovery is an inference rather than an observation.

That second finding is the empirical foundation on which this paper's question rests. If insurance coverage changes the macroeconomic consequences of a disaster, then the degree of coverage is a macroeconomic policy variable, and it should be evaluated alongside the other macroeconomic policy variables, reserves included. That step has not been taken.

2.3 Sovereign disaster risk financing

The applied literature on sovereign catastrophe risk financing is largely practitioner-led and has built the instruments now in use: national catastrophe pools, compulsory extensions to property policies, sovereign parametric contracts, catastrophe bonds, contingent credit lines and regional risk pools (Cummins and Mahul, 2009; Clarke and Dercon, 2016). Its framework is the risk layering stack, under which the frequent and small is retained and the rare and large transferred, and its benchmark is the fiscal contingent liability, in the tradition of Borensztein, Cavallo and Valenzuela (2009) on debt sustainability under catastrophic risk. Pricing in the instruments themselves has been studied separately (Lane and Mahul, 2008; Braun, 2016), and the supply-side determination of catastrophe loadings through the reinsurance capacity cycle, which bears on the multiple in Section 3, in Aslan (forthcoming).

This literature is explicit that it brackets self-insurance. The leading treatment of optimal sovereign insurance states that it focuses on risk transfer instruments and sets aside instruments under which risk is retained, including self-insurance and contingent financing (Cebotari and Youssef, 2020). The exclusion is methodologically reasonable within that paper's scope, and it is exactly the boundary this paper crosses.

2.4 The unoccupied comparison

The contribution is to put catastrophe cover and foreign exchange reserves on the same balance sheet, for the same contingency, and to price them against each other.

The starting point is Aslan (2026), which establishes that a sovereign's capacity to absorb a catastrophe on its own balance sheet depends on the currency denomination of its liabilities, and compares Turkey with Japan on that basis. That paper is about fiscal capacity and takes the instrument set as given. What is new here is the external margin and the instrument comparison: the substitution condition of Proposition 1, with its basis-risk and fungibility terms; the separation of instrument choice from participation, which is what allows the shadow value to cancel from the first and govern the second; the entry condition and its implication for the composition, not the volume, of concessional support; and the inventory in Section 5, which is original data.

The consequences for the two literatures run in three directions. The exchange rate regime and the currency composition of public debt become determinants of optimal insurance purchase, which they are not in the disaster risk financing literature. Catastrophe exposure becomes a determinant of reserve adequacy, which it is not in the reserve literature. And the regional protection gap, which Swiss Re has put at the lowest catastrophe insurance protection of any region, at 3.8 per cent of economic losses, in remarks reported in the trade press (Strebel, 2024), becomes a question about the composition of external buffers, not solely one of insurance market development.

3. When is cover cheaper than reserves?

3.1 Setting

A small open economy faces catastrophe risk. With annual probability p a catastrophe occurs that is large enough to require the sovereign to source foreign currency, either because reconstruction is import-intensive or because the fiscal response cannot be financed domestically. Write N for the foreign currency requirement in that state.

The sovereign's capacity to meet the requirement without recourse to the market depends on its monetary position: the share of public debt denominated in foreign currency, in the sense of Eichengreen, Hausmann and Panizza (2005); the rigidity of the exchange rate arrangement, classified *de facto* following Ilzetzki, Reinhart and Rogoff (2019); and import dependence in reconstruction. Section 3.3 returns to this. The present subsection asks a narrower question that does not require it: given that the sovereign wants protection, which instrument delivers it more cheaply?

Two instruments are available.

Reserves. The sovereign holds official reserves at a carry cost δ per unit per period, the spread between its external borrowing rate and the yield on the reserve asset (Rodrik, 2006), which in practice tracks the sovereign spread. A unit of reserves is available in every state and is drawn in the disaster state, with probability p.

Parametric cover. The sovereign buys notional K at a rate on line π per unit per period. The contract pays on an index, not on assessed loss. Write q for the probability that the index triggers conditional on a catastrophe of the relevant severity, so the unconditional probability of payout is pq. This coincides with the contract's published expected loss rate only where the index is calibrated to the event that would force the sovereign to the external market, which is not something a modelled loss establishes; Section 4.5 sets out why p cannot be read off a published expected loss, and Section 4.4 declines to do so. Reinsurers and catastrophe bond investors price off that expected loss rate at a multiple $M > 1$, so

$$\pi = p\,q\,M.$$

The multiple is the observable the market quotes, and it is cyclical (Lane and Mahul, 2008; Braun, 2016).

Contingent credit, in the form of a World Bank Catastrophe Deferred Drawdown Option or equivalent, is also pre-arranged, also fast-disbursing and also a foreign currency claim on a counterparty, and it belongs in any complete account of the sovereign's choice. It is excluded from the comparison here for a substantive reason. A drawdown is a loan: it supplies liquidity, not resources, and the sovereign repays it with interest. Cover delivers foreign currency the sovereign does not owe afterwards; reserves deliver currency the sovereign has already paid for, which Section 3.2 now prices as the principal term. Contingent credit delivers currency the sovereign owes from the moment of drawdown and repays with interest, which is a third position rather than a variant of either, and makes it a complement to both rather than a rival to either. The complete treatment is a three-instrument problem in which contingent credit bridges the timing gap while one of the other two bears the loss, and we do not attempt it. Two features of the region limit what is lost by the omission. Only one MENAAP sovereign has ever arranged such a line, Morocco's, and it was restructured and drawn for a pandemic rather than a catastrophe; and the six Gulf states are ineligible for the instrument on income grounds, so for them the three-instrument problem collapses back to the two compared here. For the lower-income economies it does not, and a reader who concludes that contingent credit dominates parametric cover for those sovereigns is drawing an inference this paper does not contradict.

Restricting attention to parametric contracts is deliberate. An indemnity recovery settles after loss adjustment, typically in months, and is therefore not a substitute for reserves at the

horizon over which a reconstruction import bill falls due; it is a substitute for post-disaster borrowing. Only an index-triggered contract settles fast enough for the substitution question to arise.

3.2 The condition

Compare the two instruments by the cost of delivering a unit of foreign currency in the state where it is wanted.

A unit of cover costs π per period and pays with probability pq, so its cost per unit of expected foreign currency delivered is $\pi/(pq)$, which by the pricing relation is simply M. The sovereign's cost per unit of expected delivery is therefore the multiple itself.

A unit of reserves has two costs, and the comparison usually written down counts only one of them. The first is the carry, δ per period, paid whether or not the event occurs. The second is the principal. A reserve unit is an asset held against a liability; the stock is retained if nothing happens, but in the disaster state it is spent while the liability stands, so with probability p the sovereign loses the unit itself. Counting only the carry treats the reserve stock as if it replenished itself. The expected cost of a unit of reserves is therefore $\delta + p$, and its cost per unit of expected foreign currency delivered is

$$\frac{\delta + p}{p} = 1 + \frac{\delta}{p}.$$

Cover's cost per unit of expected delivery, M, is a gross price in the same sense: it already contains the unit of expected payout plus the market's loading, since $M = 1 + \text{loading}/(pq)$. Comparing M against δ/p sets a gross price against a net carry and makes reserves look cheaper than they are by exactly one unit of expected delivery.

Proposition 1. *Parametric cover delivers foreign currency in the disaster state more cheaply than reserve accumulation when*

$$p\,(M - 1) < \delta,$$

that is, when the annual probability of a qualifying catastrophe multiplied by the market's loading over expected loss falls below the sovereign spread. Equivalently, since $\pi = pqM$, cover is cheaper when the rate on line falls below $q(\delta + p)$.

The condition compares like with like: what the sovereign pays the market *above* the expected payout, against what it pays to carry a reserve stock. Two checks confirm the form. At $M = 1$, actuarially fair cover, the left-hand side is zero and cover is preferred at any positive spread, which is right: fair insurance dominates costly self-insurance. At $\delta = 0$, costless reserves, the condition fails for any $M > 1$, which is also right. The earlier form $pM < \delta$ fails the first of these, telling a sovereign at 60 basis points facing a one-in-a-hundred-year event to prefer reserves to free-of-loading cover.

The condition differs in two respects from what an informal treatment produces, and in a third from what its form suggests.

A parametric contract fails to pay on some damaging events, and the natural first thought is that this cancels from the comparison: the sovereign does not pay for payouts it does not receive, since the contract's expected loss rate pq and therefore its premium fall in the same proportion as its expected payout. Written that way the cost per unit of expected delivery is $\pi/(pq) = M$, and q disappears.

That reasoning holds only if M is invariant to q, and it is not. Degrading a trigger lowers the expected loss rate, and the market prices lower expected loss at a higher multiple. Section 4.2 documents the relation directly in the sovereign catastrophe bond population: remote layers with expected losses at or below 1.6 per cent price at multiples between 1.92 and 4.75, while working layers at or above 5 per cent price between 1.20 and 2.37. Writing the multiple as a

decreasing function of the expected loss rate, $M(pq)$ with $M' < 0$, the cost per unit of expected delivery is $M(pq)$, and it rises as the trigger worsens.

Proposition 1 (restated). *Parametric cover delivers foreign currency in the disaster state more cheaply than reserve accumulation when*

$$p\,[M(pq) - 1] < \delta.$$

Basis risk enters through the multiple rather than as a separate discount. A less reliable trigger buys a cheaper contract and a worse one, and the market charges more per unit of expected delivery for the difference.

What does not follow, and what an informal treatment sometimes asserts, is that a reader may take π as the market quote and then discount the payout by q a second time. That double-counts. The correct statement is that q is already inside π, and that its effect on the comparison runs through the pricing function $M(\cdot)$, in the direction that disfavours cover.

Reserves serve many contingencies, and it is tempting to credit the catastrophe-motivated unit with a share of those other services and so reduce its effective carry cost below δ. At an interior optimum that credit is zero. A sovereign holding reserves at the level that optimises its non-catastrophe objectives has exhausted the marginal value of reserves in those uses, so a further unit held because of catastrophe exposure yields no additional non-catastrophe service and carries the full δ.

The premise fails where a constraint binds, and in this region it usually does. A sovereign below an adequacy benchmark, or under a programme with a net international reserve floor, values the marginal reserve unit in its non-catastrophe uses at a shadow price above δ. The fungibility credit is then strictly positive, the effective cost of the catastrophe-motivated unit is below δ, and reserves are correspondingly more attractive than Proposition 1 as stated allows. Pakistan at 17.9 per cent of the adequacy metric is not a sovereign contemplating a top-up beyond an optimum; it is a sovereign for which every reserve unit is scarce in every use. The condition below should therefore be read as generous to cover wherever the buffer is thin, which is to say wherever the question is live. We return to this in Section 7.3, where the same shadow price bears on whether cover can be credited against a reserve floor.

Cover is cheaper when the rate on line falls below the trigger probability times the sovereign spread plus the hazard frequency, that is when $\pi < q(\delta + p)$. The hazard frequency belongs on the right-hand side because the reserve unit is consumed in the disaster state while the liability behind it stands; dropping it recovers the retracted form. The condition's form is standard; what matters is that it survives both corrections, and that both run in the same direction, which is against cover. One correction we have not applied runs the other way, and we note it rather than leave a reader to find it: the comparison assumes the marginal reserve unit is drawn with certainty in the disaster state, and if it is drawn with probability $z < 1$ the reserve arm becomes $1 + \delta/(pz)$, which is larger, so cover becomes relatively more attractive. We have no basis for estimating z and set it to one throughout. The flow comparison also assumes the reserve stock is rebuilt at par each period, since otherwise a unit could be consumed only once; the depletion term below charges the spread differential on that rebuilding and so does not double-count it.

Depletion. One asymmetry does not cancel. Reserves drawn in a disaster must be rebuilt, and rebuilt when borrowing conditions have deteriorated. If rebuilding proceeds at $\delta' > \delta$ over T periods, the reserve side carries an additional expected cost of $p(\delta' - \delta)T$ per unit, and the condition becomes $p\,[M(pq) - 1] < \delta + p(\delta' - \delta)T$. The threshold spread at which the two instruments break even is therefore $p\,[M(pq) - 1] - p(\delta' - \delta)T$, which is lower than the threshold without depletion: the band in which cover is preferred widens, and Section 4.1 reports the magnitude.

The rationed sovereign, and why the condition does not extend to one. For a sovereign with no access to international capital markets δ is not observable, and it is tempting to handle this by letting δ grow without bound, so that Proposition 1 holds *a fortiori*: if cover

is cheaper at a spread of 600 basis points it is cheaper still where no spread exists. We set the argument out because it is the natural extension to make, and it is wrong.

The error is a conflation of two distinct costs. A unit of reserves has an acquisition cost, the foreign currency that must be obtained to build the stock, and a carry cost δ, the spread between what the sovereign pays on the liability and what it earns on the asset. The principal is recoverable; the carry is not. Proposition 1 compares carry costs, and that is legitimate when both instruments are accessible at the margin. Rationing is not a high carry cost. It is a quantity constraint on acquisition: the sovereign cannot obtain the marginal unit at any price. Letting $\delta \to \infty$ writes a quantity constraint as a price and then reads a price comparison off it.

Doing so also produces an asymmetry that cannot be defended. The same foreign currency scarcity that makes the reserve unit hard to acquire makes the premium hard to fund, because the premium is paid in that currency too and, unlike the reserve principal, is never returned. An argument that lets scarcity raise δ without limit while leaving the premium side untouched has scored one consequence of scarcity and ignored the other. The fungibility correction above runs the same way: a sovereign below an adequacy benchmark or under a reserve floor values the marginal reserve unit above δ in its other uses, which makes reserves *more* attractive, and it is precisely the rationed sovereigns for which that correction bites hardest.

The honest statement is therefore narrower. For a quantity-rationed sovereign the comparison in Proposition 1 is not decided in favour of cover; it is undefined, because neither instrument is available at the margin and the premium must be funded from the same currency the sovereign cannot obtain. Section 3.3 formalises this by introducing the shadow price of foreign currency in the pre-event state, and Section 4.1 makes accessibility a domain restriction on the condition rather than an extension of it. Three of the economies in Section 5 are in that position, and the consequence is that the paper has less to say about them, not more.

3.3 Whether to hold protection at all

Proposition 1 compares instruments conditional on the sovereign wanting protection. Whether it wants any is a separate question, and the two should not be conflated.

Let $v \geq 1$ denote the shadow value of a unit of foreign currency in the disaster state, which depends on monetary sovereignty σ and on the existing buffer ρ measured against an adequacy benchmark. Two properties are imposed rather than derived.

Assumption 1. *v is decreasing in monetary sovereignty and decreasing in the buffer ratio, approaching unity as either grows large.*

The first reflects that a sovereign issuing the currency its liabilities are written in can meet them without sourcing foreign exchange. The second reflects that a sovereign already holding liquid external assets well above the adequacy benchmark absorbs the requirement out of stock. We label this an assumption, not a result, because no primitive environment is specified here from which the signs could be derived, and the ordering it generates should be read as qualitative. Micro-founding v in a two-period model with a borrowing constraint and an import-compression cost would sharpen what follows and is left for future work.

Both instruments are paid for in foreign currency before the event, and that currency has a shadow price of its own. Let $\lambda \geq 1$ denote the shadow value of a unit of foreign currency in the pre-event state, so that $\lambda = 1$ for a sovereign that can obtain foreign currency freely and λ rises as it cannot. Writing the comparison with λ in it is what the previous section's failed *a fortiori* argument omitted.

Consider each instrument per period. A unit of cover notional costs $\pi = pqM$ in foreign currency, so $\lambda\pi$ in utility, and delivers one unit valued at v with probability pq; it is worth buying when $pqv > \lambda pqM$, that is when $v/\lambda > M$. A unit of reserves costs δ in carry and, with probability p, the principal itself, so $\lambda(\delta + p)$ in expectation, and delivers one unit valued at v with probability p; it is worth holding when $pv > \lambda(\delta + p)$, that is when $v/\lambda > 1 + \delta/p$. Protection of some kind is therefore worth holding when

$$\frac{v}{\lambda} > \min\left\{M,\ 1 + \frac{\delta}{p}\right\},$$

and Proposition 1 is the statement of which arm of that minimum binds: cover is the cheaper instrument exactly when $M(pq) < 1 + \delta/p$, which is the condition above.

Note what λ does and does not do. It multiplies both instruments equally, so it cancels from the comparison and Proposition 1 is unaffected by it. That is a robustness result worth stating: the instrument-choice condition does not depend on how scarce foreign currency is in the pre-event state. But λ does not cancel from the participation condition, where it appears as the ratio v/λ. The relevant quantity is therefore not the absolute value of disaster-state foreign currency but its value *relative* to pre-event foreign currency, and for a sovereign that is foreign-currency constrained in both states the ratio can be close to unity even though both v and λ are large.

This is what the failed *a fortiori* argument of Section 3.2 missed. Rationing raises λ as surely as it raises the shadow cost of reserves, and the ratio is what governs whether any protection is worth buying. A sovereign that cannot obtain foreign currency does not thereby become a natural buyer of cover; it becomes a sovereign for which both instruments recede together, and Section 4.1 treats accessibility as a restriction on the domain of the comparison rather than as a limiting case within it. The ordering the assumption generates is that protection is worth most to a sovereign combining a rigid exchange rate arrangement, a high share of foreign-currency liabilities and high import dependence with a thin buffer. One qualification is needed and it limits where the ordering may be used. Assumption 1 constrains v alone, while the operative quantity established above is v/λ, and nothing is assumed about how λ varies with the same characteristics. If λ rises in buffer thinness at least as fast as v does, which is the natural reading for a sovereign whose foreign currency is scarce in both states, then v/λ need not be monotone and the ordering does not follow at the thin-buffer end. We therefore use the ordering only where λ is close to unity, which is the ample-buffer tail, and we do not rely on it to establish that protection is most valuable to the constrained sovereigns. Section 5.3's partition is read accordingly.

The magnitudes of the two arms are very different, and the difference governs how the rest of the paper should be read. Table 2 puts M between 1.92 and 4.75 on the remote layers a sovereign would actually buy, so $v/\lambda > M$ requires the shadow value of disaster-state foreign exchange to be between roughly two and five times its pre-event value. That is a demanding condition and we do not establish whether any sovereign in this region meets it. The reserve arm is less demanding but, once the principal is counted, not trivially satisfied: for Saudi Arabia at 60 basis points and a one-in-a-hundred-year qualifying event, $1 + \delta/p$ is about 1.6, so holding reserves against the catastrophe motive requires the disaster state to be worth at least sixty per cent more than the pre-event state. Both arms therefore impose a real hurdle, and neither can be evaluated from published data. Counting the carry alone makes the reserve arm clear by construction; that is an artefact of omitting the principal.

Two consequences follow. First, for a low-spread sovereign the live question is never whether to hold protection but which instrument to hold, and the answer is reserves. Second, and more awkwardly for the instrument this paper is about, $v > M$ is necessary for cover to be worth buying *independently of the spread*. If v lies below M for most sovereigns, cover is not worth buying anywhere, whatever Proposition 1 says, and the threshold exercise in Section 4 characterises only the instrument-choice margin conditional on protection being worth holding at all. Clearing the threshold is therefore necessary and not sufficient. We draw the consequence explicitly rather than leaving it implicit: the threshold exercise that follows is a conditional statement, and the paper's negative conclusion about this region is over-determined, resting both on the price comparison and on a participation condition we cannot evaluate. Quantifying v/λ for even one sovereign would be a more valuable contribution than refining the threshold,

and it requires the micro-foundation this section declines to supply.

Cross-sectional use of this ordering is limited in two ways. The buffer is not exogenous to the decision being explained, since a sovereign that buys cover should optimally hold fewer reserves; where a high buffer is driven by something outside the catastrophe problem, as hydrocarbon revenue drives it in the Gulf, the partial reading is defensible, and where it is not, it is not. And the aggregation of three distinct characteristics into σ is not specified, so σ orders sovereigns only where the three move together.

3.4 Why cover might not be bought

Where Proposition 1 holds and protection is valuable on the buffer evidence, a sovereign that buys none requires explanation. Several accounts compete, and they are not easily separated.

A sovereign under comprehensive sanctions, in default, in conflict, or without a contracting authority recognised by counterparties cannot place risk at any price. Religious objection to conventional insurance and conventional securitisation is a standing issue in several markets in the region, where Takaful and Sukuk-structured alternatives remain thin at the sovereign catastrophe layer. Expected post-disaster grants and concessional lending reduce the effective shadow value of the sovereign's own foreign currency and so reduce optimal cover, which is the charity-hazard argument. A premium is an appropriated, annually voted line item while the carry cost of reserves sits on the central bank's balance sheet and appears in no budget document, so the two reach the decision-maker with different salience. Pre-arranged cover may weaken incentives for ex ante risk reduction, though parametric structures mitigate this since the payout is independent of realised loss and the sovereign keeps the full benefit of reducing damage.

The sixth is a fixed cost of market access. Entering requires a probabilistic catastrophe model, a legal framework defining perils and triggers, a public entity with standing to contract, a payout mechanism, and market relationships. Write F for the one-time cost. Entry requires the capitalised gain from substitution to exceed it:

$$q\left[\delta - p\left(M(pq) - 1\right)\right] KA \;\geq\; F,$$

with K the notional and A an annuity factor over the horizon the capability persists. The factor q appears because reserves delivering the same expected foreign currency as notional K need a stock of only qK; omitting it overstates the capitalised gain by $1/q$.

Proposition 2. *Support directed at F lowers the entry threshold permanently, so a sovereign brought across it remains across when support ends; support directed at the premium raises the left-hand side only while paid, so a sovereign brought across by premium subsidy exits when the subsidy lapses. Where $\delta < p\left[M(pq) - 1\right]$ the bracket is negative and no fixed-cost support of any size induces entry, since the condition then fails even at $F = 0$. Premium subsidy is not symmetric in this respect: a subsidy covering a share s of the premium replaces π with $(1 - s)\pi$ in the sovereign's own comparison, and for s large enough the sovereign buys whatever the unsubsidised bracket would have said.*

The asymmetry in the final clause is the operative result and it runs against the instrument it might appear to favour. Where substitution does not pay at market prices, fixed-cost support correctly induces nothing, because the transaction is not worth doing. Premium subsidy induces it anyway. A sovereign for which $p\left[M(pq) - 1\right] > \delta$ that buys cover under premium subsidy has been moved to a transaction that is value-destroying at market prices and becomes so again the moment the subsidy lapses. Section 4.3 finds the substitution condition failing for the four Gulf states with ample buffers, so the bracket is negative there; for the five economies that satisfied the condition the bracket was positive and nothing was bought, which is the finding Section 4.3 sets out and which the entry condition does not by itself explain. Section 7.4 takes up what follows for a donor practice built largely on premium support.

4. Threshold spreads

4.1 The inversion

Proposition 1 contains one quantity nobody publishes for a sovereign, the annual probability p of a catastrophe large enough to require external financing, and one the market quotes, the multiple M. Inverting the condition turns it into a threshold spread: cover is the cheaper instrument when the sovereign spread exceeds $p\,[M(pq) - 1]$, the hazard frequency multiplied by the market's loading over expected loss.

The inversion converts the comparison into a question a reader can put to their own sovereign. A central bank knows its own spread. It can form a view on the return period of an event that would force it to the market. The multiple is quoted at each renewal. What the inversion does not do, and what it is often assumed to do, is dispense with p. The threshold is proportional to p, so a reader must either assume a frequency or read the threshold as a range across plausible ones, and every table below is tabulated accordingly.

The threshold gives one edge of a region, not the region itself, and the second edge has to be stated because the rest of the paper turns on it. Proposition 1 is a price comparison and it is monotone in δ: the higher the spread, the stronger the case for cover. Taken alone it would imply that the sovereign with the worst borrowing terms has the strongest case, and in the limit that a sovereign shut out of the market has an unanswerable one. Section 3.2 explains why that limit does not follow. But a sovereign also transfers risk only if a counterparty will contract with it, which is a condition on standing rather than on price.

Definition (the band). *A sovereign is in the band in a given period when the comparison applies to it and it favours cover. Three conditions:*

(i) Accessibility. The sovereign can obtain foreign currency at the margin, so that both instruments are available and the price comparison is defined. This fails where the sovereign is quantity-rationed in foreign currency rather than merely paying a high price for it, which we take to be the case where it has no access to international capital markets.

(ii) Standing to contract. An international counterparty will write the risk. This fails where the sovereign is under sanctions bearing on financial contracting, lacks a contracting authority recognised by counterparties, or has arrears whose settlement conditions bar new commitments. The grounds are applied economy by economy in Section 5.3 and recorded in the final column of Table 5.

(iii) Price. $\delta > p\,[M(pq) - 1]$, the condition of Proposition 1: the carry cost of reserves exceeds the hazard frequency multiplied by the market's loading over expected loss.

The order is deliberate: condition (i) restricts the domain of condition (iii) rather than adding a hurdle beyond it. Section 3.2 gives the reason — a price comparison between two instruments is meaningful only where either can be acquired at the margin, and where neither can, the comparison is undefined rather than resolved in favour of cover. Putting accessibility first also forecloses a circularity. Were price evaluated first and allowed to rise without bound under rationing, the same scarcity would satisfy one condition and defeat the other by construction, and the finding that they never hold together would be definitional. Treating accessibility as a domain restriction means the sovereigns it excludes are not ones for which cover was cheap and unaffordable, but ones about which this framework is silent.

Two candidate limbs are deliberately excluded from condition (i). A reserve floor under a Fund programme is a constraint on drawing reserves down, not on acquiring foreign currency, so it is not rationing in the sense Section 3.2 defines; it belongs in condition (iii) as a fungibility correction to δ, which is where Section 3.2 puts it, and counting it twice was the residual form of the circularity this Definition was rewritten to remove. And whether a sovereign's premium was in fact met by a third party is the outcome being explained, not a state variable, so conditioning the domain on it was circular in a more elementary way.

One thing the Definition deliberately omits is the participation condition of Section 3.3. A

sovereign can be in the band as defined here and still rationally hold no cover, because v/λ may fall below M, in which case cover is worth buying at no spread whatever. We leave participation out because we cannot evaluate it for any sovereign in the region, and including an unevaluable condition would make the band unfalsifiable. The consequence is that band membership is necessary and not sufficient for cover to be worth buying, and the five economies Table 4 records as satisfying all three conditions are candidates for the institutional explanations of Section 3.4 only on the assumption that participation held for them, which we cannot check.

Three features of the Definition should be noted. None of the three conditions is a spread level, so the band is not an interval on the real line and the word is a convenience. Conditions (i) and (ii) concern insurance rather than borrowing: paying a premium is not issuing debt, so a sovereign locked out of the bond market may retain the standing to buy cover even where condition (i) fails, as Djibouti does. And both are coarse, a yes or no rather than a price of access.

Table 1. Threshold sovereign spread above which parametric cover is the cheaper instrument (basis points)

Panel A. From the fitted pricing function, soft market

Return period	p	$q = 1.0$	$q = 0.8$	$q = 0.6$
1-in-200	0.005	147*	167*	195*
1-in-100	0.010	193	222	264*
1-in-50	0.020	236	280	342
1-in-30	0.0333	251	310	394
1-in-20	0.050	238	312	417

Panel B. From the fitted pricing function, hard market

Return period	p	$q = 1.0$	$q = 0.8$	$q = 0.6$
1-in-200	0.005	256*	287*	331*
1-in-100	0.010	354	400*	466*
1-in-50	0.020	475	543	640
1-in-30	0.0333	570	661	791
1-in-20	0.050	638	753	917

Panel C. Implied by the eight observed remote-layer risk margins, at a one-in-a-hundred-year frequency

Trigger reliability q	Soft issuance, 2017–2021	Hard issuance, 2023–2026
1.0	150–340	300–600
0.9	178–389	344–678
0.8	213–450	400–775
0.7	257–529	471–900
0.6	317–633	567–1,067
0.5	400–780	700–1,300

Panel D. Cluster bootstrap of the fitted threshold, fifth to ninety-fifth percentile (basis points)

Return period	Soft, $q = 1.0$	Soft, $q = 0.6$	Hard, $q = 1.0$	Hard, $q = 0.6$
1-in-200	123–174	161–235	221–293	283–386
1-in-100	162–224	221–311	308–397	402–530
1-in-50	196–271	287–395	410–527	556–716
1-in-30	201–294	327–452	468–643	684–879

Return period	Soft, $q = 1.0$	Soft, $q = 0.6$	Hard, $q = 1.0$	Hard, $q = 0.6$
1-in-20	169–291	336–485	500–731	761–1,032

Notes. Panels A to C give the threshold spread from Proposition 1 as corrected, $p\,[M(pq)-1]$: the market's loading over expected loss, not the gross price. Panels A and B compute it from the pricing function $M(\cdot)$ fitted to Table 2 by regime. Cells marked with an asterisk lie outside the range of expected losses in the fitting sample, 0.79 to 6.54 per cent soft and 0.90 to 5.84 hard, and are extrapolations rather than combinations the observed market has priced; the remaining twenty-one cells are within it.

Panel A is not monotone in the hazard frequency, and the turning point is a property of the estimated pricing function rather than an artefact. The threshold is p times the loading, and the loading falls steeply enough in expected loss that in the $q = 1$ column of Panel A the product peaks near a one-in-thirty-year frequency and declines thereafter, while at lower trigger reliabilities the turning point lies below the tabulated grid and those columns rise throughout it. Cover for frequent events is priced at a low multiple, so the loading the sovereign pays above expected loss is small relative to the carry it avoids.

Panel D attaches uncertainty to Panels A and B, which are otherwise point thresholds from an estimated function. We resample the thirteen transactions with replacement two thousand times, refit $\log M = a + b \log \mathrm{EL} + \gamma \mathbb{1}[\text{hard}]$ on each draw, and recompute the threshold; the panel reports the fifth and ninety-fifth percentiles. The bands are informative rather than wide: at a one-in-a-hundred-year frequency and a perfect trigger the soft threshold lies between 162 and 224 basis points and the hard between 308 and 397, which do not overlap. This converts Section 4.3's classification rule into a probability. At those parameters every one of the eleven qualifying observations exceeds the threshold in all two thousand draws, the lowest of them, Egypt's 360 basis points in 2001, clearing the ninety-fifth percentile of 224 with room to spare. The five-economy result is not an artefact of where the point estimate happens to fall.

Panel C is an identity rather than a fitted relation, but it is no longer free of p. For an observed contract the risk margin is $\pi = pq\,M(pq)$, so the threshold is $\pi/q - p$. The subtraction is material: 50 basis points at a one-in-two-hundred-year frequency, 100 at one-in-a-hundred, 200 at one-in-fifty. Panel C is tabulated at one-in-a-hundred, and at one-in-two-hundred every figure rises by 50 while at one-in-fifty every figure falls by 100. The inversion is sometimes presented as a device that avoids estimating p; it does not, and the frequency has to be assumed or stated as a range wherever the threshold is used.

The two trigger channels must not be added. Panel C scales the observed margin by $1/q$ because the premium is held at its observed level; Panels A and B scale the fitted multiple by q^b, where b is the within-regime fitted elasticity of about -0.43, because degrading the trigger lowers the expected loss and the market reprices it. Working in the other direction, the depletion term in Section 3.2 lowers each threshold by $p(\delta' - \delta)T$: at a post-disaster spread widening of 300 basis points sustained over three years that subtracts 9 basis points at a one-in-a-hundred-year frequency and 45 at one-in-twenty, an order of magnitude smaller than either trigger correction.

4.2 Calibrating to observed placements

The threshold can be read almost directly off any transaction that publishes both a modelled expected loss and a price, because the risk margin a sovereign pays equals pqM, so the threshold $p[M(pq) - 1]$ is that margin divided by the trigger probability, less the hazard frequency. The World Bank's capital-at-risk notes publish both quantities at issuance, which makes the sovereign catastrophe bond population the natural calibration sample.

Table 2. Sovereign catastrophe bond pricing, 2017–2026

Sponsor and year	Peril	Modelled expected loss	Risk margin	Multiple
Mexico 2018	Earthquake, remote layer	0.79%	2.50%	3.16
Chile 2018	Earthquake	0.86%	2.50%	2.91
Mexico 2020	Earthquake, remote layer	0.90%	3.50%	3.89
Mexico 2024	Earthquake, remote layer	0.90%	4.00%	4.44
Chile 2023	Earthquake and tsunami	1.00%	4.75%	4.75
Jamaica 2024	Atlantic named storm	1.50%	7.00%	4.67
Jamaica 2021	Atlantic named storm	1.52%	4.40%	2.89
Colombia 2018	Earthquake	1.56%	3.00%	1.92
Kyrgyz Republic 2026	Earthquake and precipitation	2.30%	6.00%	2.61
Jamaica 2026	Atlantic named storm	2.48%	6.75%	2.72
Philippines 2019	Earthquake	3.00%	5.50%	1.83
Philippines 2019	Tropical cyclone	3.00%	5.65%	1.88
Mexico 2017	Earthquake	3.43%	4.50%	1.31
Mexico 2017	Pacific named storm	3.96%	5.90%	1.49
Mexico 2020	Pacific named storm	4.06%	6.50%	1.60
Mexico 2024	Pacific named storm	4.09%	12.00%	2.93
Peru 2018	Earthquake	5.00%	6.00%	1.20
Mexico 2017	Atlantic named storm	5.56%	9.30%	1.67
Mexico 2020	Earthquake, working layer	5.78%	9.00%	1.56
Mexico 2020	Atlantic named storm	5.61%	10.00%	1.78
Mexico 2024	Earthquake, working layer	5.84%	11.00%	1.88
Mexico 2024	Atlantic named storm	5.69%	13.50%	2.37
Mexico 2018	Earthquake, working layer	6.54%	8.25%	1.26

Notes. Twenty-three tranches, ordered by expected loss within the remote and intermediate groups; the working-layer group at the foot of the table is not sorted. The risk margin is the investor spread over the collateral return, which is the price quantity these transactions publish; it is not a total coupon. Multiple is risk margin divided by modelled expected loss at issuance. All are parametric except the two Philippine tranches, which use a modelled-loss trigger. Issuer is the IBRD except the Kyrgyz Republic, where it is the Asian Development Bank. That row prices Kyrgyz catastrophe risk but carries Asian Development Bank credit rather than sovereign credit, which is why Section 4.4 includes it in the pricing sample and excludes it from the spread comparison; the identically priced Tajikistan tranche of the same issuance is omitted because including a second tranche at the same price and expected loss would double-weight one transaction in the fit. Sources in Appendix A.

The distribution has three features that bear on the threshold.

Across all twenty-three tranches the mean multiple is 2.47 and the median 1.92; excluding the two Philippine tranches, whose modelled-loss trigger makes their published expected loss a loss fraction rather than a payout probability and so not strictly commensurable with the rest, the mean is 2.52 and the median 2.37.

Comparison with the broad market requires care, because the standard benchmark is a

no-intercept regression slope rather than a mean of multiples, and the slope weights by the square of expected loss and so is pulled toward the working layers. Computed the same way, the sovereign sample gives 1.81 over all twenty-three tranches, 1.55 for 2017 to 2021 and 2.43 for 2023 onward, against the 2.69 that Lane and Mahul (2008) report for 1997 to the first quarter of 2008 and the 2.31 to 2.37 of their cycle-adjusted specifications. Appendix A.9 records which specification each of their published figures comes from, which their own paper does not do and without which the comparison cannot be made like for like. Sovereign placements are therefore priced below the broad market of that era on the like-for-like construction, and only the hard-market tranches reach the cycle-adjusted figure. The comparison is period-confounded and we do not lean on it: the sovereign sample runs from 2017 to 2026 and the benchmark from 1997 to 2008, and the paper's own broad-market series moves from 5.15 to 2.40 across 2023 to 2026, so the level of either estimate is a statement about its own cycle. What the comparison supports is only the weak claim that nothing in these prices suggests a surcharge for sovereign counterparty status.

The period split is sharp and it matters more than the level. Tranches issued between 2017 and 2021 average a multiple of 2.02 and those issued from 2023 average 3.30. On the eight remote layers the separation is cleaner still: the five soft-market tranches price at multiples of 1.92 to 3.89, and the three hard-market tranches at 4.44, 4.67 and 4.75, with no overlap. The direction matches the broad primary market, where the comparable nine-month average multiple was 5.15 in 2023, 4.14 in 2024, 3.26 in 2025 and 2.40 in 2026; but the sovereign sample has only three post-2022 remote observations, Chile's of 2023 and Mexico's and Jamaica's of 2024, which is too few to say whether it tracks that decline or merely sits inside it.

The multiple falls steeply in expected loss, and this sample is large enough to put a number on the relation that Section 3.2 relies on. Remote layers with expected losses at or below 1.6 per cent price at multiples between 1.92 and 4.75, while working layers at or above 5 per cent price between 1.20 and 2.37. Fitting $\log M = a + b \log(\mathrm{EL})$ across all twenty-three tranches gives an elasticity of $b = -0.456$, with a standard error of 0.078 and an R^2 of 0.62; excluding the two modelled-loss tranches changes it to -0.453. This is the first estimate of M' we are aware of for the sovereign segment, and it is what makes Proposition 1's restatement operational rather than qualitative. Two cautions attach, and one of them turns out to cut the other way. The tranches are not independent draws: thirteen are Mexican, four of the eight hard-market tranches belong to Mexico's 2024 transaction, and tranches of one transaction share a modelling vendor, a placement date and a book. Re-estimating with the regime shift in the specification and standard errors clustered by transaction, which gives thirteen clusters, the elasticity is -0.427 with a cluster-robust standard error of 0.026, a t of -16.4, and the regime shift is 0.436 in logs, a level ratio of 1.55, with a t of 6.4. Allowing for the correlation structure therefore tightens both estimates rather than loosening them, and the 55 per cent level shift and the negative elasticity both survive it. The second caution is not resolvable here: the expected loss appears in the denominator of the dependent variable, since M is the margin divided by it, so measurement error in the published expected loss biases the elasticity mechanically. Regressing the log margin on the log expected loss is not a remedy, because $\log M = \log \pi - \log \mathrm{EL}$ makes the two slopes differ by exactly one and the standard errors identical; that regression is sometimes reported as a robustness check, which it is not. Neither specification is identified against measurement error in the expected loss, and the sign rests on the raw price gap between remote and working layers, which is large: multiples of 1.92 to 4.75 at expected losses at or below 1.6 per cent against 1.20 to 2.37 at or above 5 per cent. The elasticity should be read as a well-determined sign and order of magnitude rather than a precise coefficient.

The regime split runs through the intercept and not the slope, which is the useful part. Fitted within regime, the elasticity is -0.4262 for 2017 to 2021 and -0.4292 for 2023 onward, statistically indistinguishable, while the level term rises from 2.9305 to 4.5427; the within-regime fits are tighter than the pooled one, at R^2 of 0.79 and 0.82. The reinsurance cycle therefore

shifts the price of catastrophe risk roughly proportionally across the loss distribution rather than steepening it, so a sovereign's position relative to the threshold moves with the cycle by a factor it does not control, but the trade-off it faces between trigger reliability and price does not change with the cycle. Panels A and B of Table 1 are computed from these two fitted functions, which removes the earlier practice of gridding over p and M independently and so producing cells, at high frequency and high multiple together, that the pricing function excludes. Nine of the thirty cells still lie outside the fitted range of expected losses and are marked as extrapolations in the table.

A sovereign buying catastrophe protection is buying the remote layer, so the remote-layer figures are the relevant ones.

Risk margins on the eight remote sovereign layers are 250, 250, 300, 350, 440 in the soft market and 400, 475, 700 in the hard market. Since the threshold is the risk margin divided by the trigger probability, less the hazard frequency, substitution at a perfect trigger and a one-in-a-hundred-year frequency requires a sovereign spread of **at least 150 to 340 basis points at soft-market prices and 300 to 600 at hard-market prices**, with both ranges rising as the trigger worsens and falling as the assumed frequency rises, as Table 1 Panel C sets out. We report the two regimes separately rather than pooling them, and we do not treat Jamaica's 700 as an outlier: it is the top of a three-observation hard-market range, and on a pooled reading it would be discarded for no reason other than that it is inconvenient. Working-layer cover, which a sovereign would buy only to smooth frequent events, carries risk margins of 600 to 1,350 basis points and so implies thresholds of 500 to 1,250 at the reference frequency and is not a reserve substitute on any reading.

The threshold is not a number but a range whose level is set by the reinsurance cycle. The fitted level term rises by about 55 per cent between the two regimes, which at a one-in-a-hundred-year frequency and a perfect trigger moves the threshold from 193 to 354 basis points, a shift of 161 basis points. That is more than the entire spread Saudi Arabia paid at any observation in Table 3 and more than half the spread Morocco paid at its 2022 peak, so a finance ministry that satisfied the condition at a 2019 renewal could fail it at a 2024 one with nothing changed on its own balance sheet. It also follows that the figure to use for a prospective placement is the hard-market one, since that is where the market has been since the January 2022 renewal, and we date the regime from 2022 throughout.

The only historical catastrophe pricing series for this geography, the early reinsurance programme of the Turkish pool, has rates on line of 420 to 613 basis points, implying thresholds of 320 to 513 at the reference frequency, which we report in Appendix A for comparison. That range sits inside the corrected hard-market range of 300 to 600 and overlaps the soft-market range of 150 to 340 at its lower end, so the two calibrations agree rather than one correcting the other. This is reassuring but not probative: that programme covered a residential pool's excess-of-loss reinsurance rather than a sovereign parametric layer, it predates the modern insurance-linked securities market, and Türkiye is outside the sample region.

4.3 Who clears the threshold

A threshold of 150 to 340 basis points at soft-market prices, and 300 to 600 at hard-market prices, on the reference parameters of a perfect trigger and a one-in-a-hundred-year frequency, invites the obvious question of which sovereigns in this region have ever cleared it. Published data answer it.

No continuous sovereign spread series spanning 2000 to 2026 exists in the public record for these economies. Three partial sources do. The IMF's Global Financial Stability Report statistical appendix published EMBI Global yield spreads by country from 2000 until the series was discontinued after the 2015 vintage. Dated cross-sections of five-year US dollar credit default swap spreads are published in Gulf market commentary from 2021. And individual IMF staff reports give point observations. Table 3 assembles what these yield. The three are different

instruments and are reported in separate panels; they should not be concatenated into one series.

Table 3. Sovereign spreads, selected MENAAP economies

Panel A. EMBI Global yield spreads, basis points, end of period

Economy	2001	2003	2005	2008 Q4	2011	2013	2014
Egypt	360	131	58	385	607	443	351
Pakistan	1,115	289	198	2,112	1,274	606	525
Iraq	—	—	—	1,282	603	511	579
Jordan	—	—	—	—	500	290	—
Morocco	518	160	75	—	—	270	224
Tunisia	—	146	81	464	—	—	—

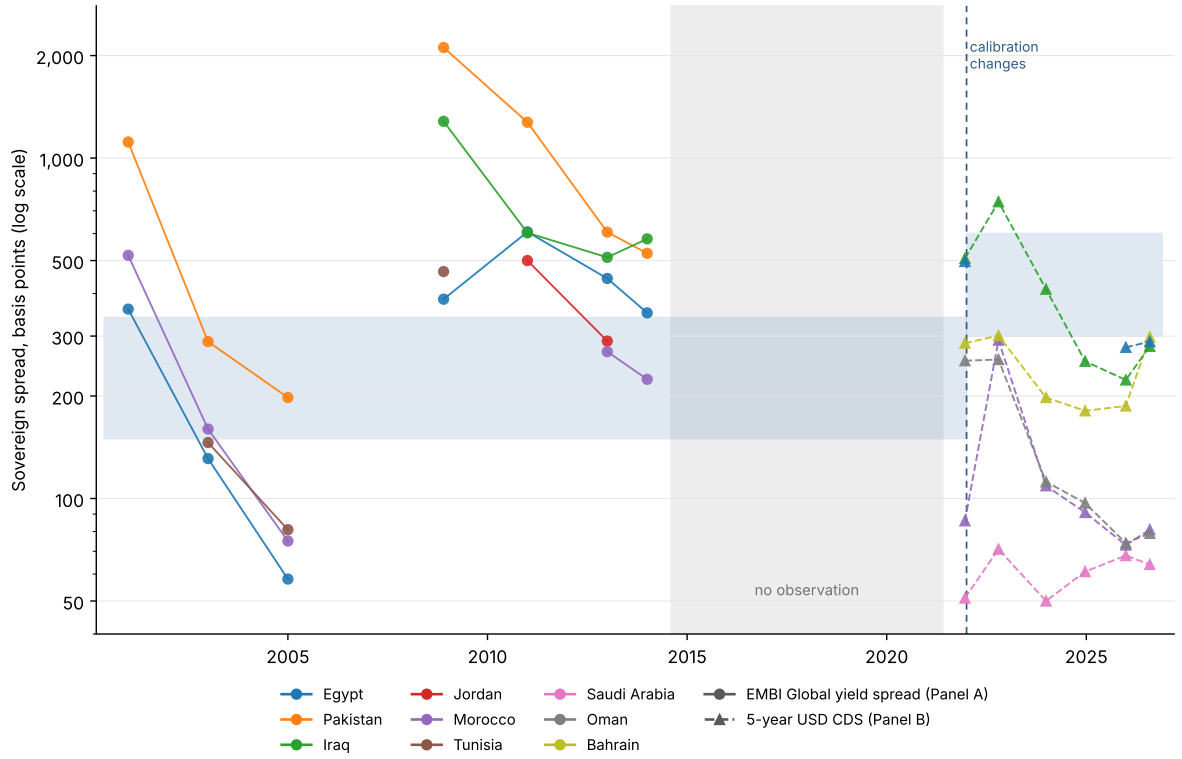
Panel B. Five-year US dollar CDS spreads, basis points, selected dates

Economy	31 Dec 2021	19 Oct 2022	31 Dec 2023	19 Dec 2024	31 Dec 2025	4 Aug 2026
Saudi Arabia	51	71	50	61	68	64
Morocco	86	292	109	91	73	81
Oman	254	256	112	97	74	79
Bahrain	286	301	198	181	187	298
Iraq	507	744	412	253	223	280
Egypt	497	—	—	—	278	289
Türkiye	555	731	279	250	207	235

Notes. Panel A is from the IMF Global Financial Stability Report statistical appendices, source JPMorgan, on the EMBI Global definition. Dashes denote an economy absent from the index in that year, or a value not reported in the vintage used. Egypt's index membership changed over the period: trade coverage records a removal from the EMBI in June 2011 and a readmission in January 2022, which we cannot reconcile with the 2013 and 2014 Egyptian readings printed in the Fund's own statistical appendix and reproduced here, and we report the appendix as published rather than suppress figures we cannot explain. The Gulf states were admitted to the EMBI benchmarks only in 2019. Panel B is CDS, a different instrument from a bond index spread, and the two panels are therefore not continuous; each Panel B column is taken from the one daily report that prints that date, either as its own date or as its year-end reference column, and figures are rounded from two decimal places. Dashes in Panel B denote a quotation reported as not available. Jordan, Tunisia, Pakistan and Djibouti are not quoted in the CDS source. The two panels leave 2015 to 2020 unobserved for every economy in the sample, Panel A ending in 2014 and Panel B beginning at the end of 2021. That gap contains Morocco's 2020 placement and the Pakistani window discussed in Section 7.1, and we record it as missing data rather than as evidence of anything. Türkiye is shown as the comparator used in Section 6.3 and is not part of the region. Sources and vintages are in Appendix A.

Note. The shaded horizontal band is the threshold at a one-in-a-hundred-year frequency and a perfect trigger: 150 to 340 basis points on soft-market pricing to the end of 2021, 300 to 600 from 2022. The dashed vertical rule marks where the dated calibration changes; readings at 31 December 2021 fall before it and are judged against the soft range, as Section 4.3's rule requires. Points above the applicable band met the price condition at that date. The grey vertical panel marks the window between the last Panel A reading and the first Panel B reading, in which no spread observation exists for any economy in the sample. Every economy in Table 3 is plotted except Türkiye, which that table carries as a comparator outside the region. The two instruments are drawn separately and are not one series: circles joined by solid lines are Panel A EMBI Global yield spreads, triangles joined by dashed lines are Panel B five-year US dollar credit default swaps, and no line joins a point in one panel to a point in the other. Lines break wherever consecutive observations are more than about two and a half years apart, so

Figure 1. Sovereign spreads against the substitution band



no segment interpolates across a gap. The option-adjusted Egyptian readings of 2023 and 2024 and the staff-report point observations for Jordan and Oman are on further bases again and are omitted here; they appear in Table 4 and Appendix A.5. Morocco and Iraq and Egypt appear in both panels. The figure is an aid to reading Table 3 against the band, and no claim in the text rests on comparing a circle with a triangle.

Three things follow from the panel. The first two bear on claims commonly made about this region on the basis of credit ratings rather than prices, and the prices do not support them.

First, sovereigns in this region did spend extended periods in the band. Egypt stood at 607 basis points at the end of 2011 and, on option-adjusted measures reported by JPMorgan, between roughly 990 and 1,434 basis points through 2023 and January 2024. Pakistan was above 600 basis points in 2001, 2008 and 2013, reaching 2,112 in the fourth quarter of 2008 and 1,274 in 2011, and stood at 525 in 2014. Iraq is at 1,282 in the fourth quarter of 2008 and between 511 and 603 at each of 2011, 2013 and 2014, and its credit default swaps were at 507 and 744 basis points at the two 2021 and 2022 observations in Panel B. The record therefore contains sovereigns well above the threshold for extended runs, which is worth establishing because the absence of cover in this region is often attributed to spreads being too low for substitution to be worth considering.

Second, the Gulf was not uniformly investment grade. Bahrain fell below investment grade at Standard and Poor's in February 2016, at Moody's in March 2016 and at Fitch in August 2016, and has remained there continuously, with a further downgrade to B at Standard and Poor's in November 2025 and at Fitch in February 2026, these two per-notch dates resting on contemporaneous press reporting rather than on the agency releases, as Appendix A.6 records. Oman fell below at Standard and Poor's in May 2017, at Fitch in December 2018 and at Moody's in March 2019, and regained investment grade at all three only between September 2024 and December 2025, with the IMF reporting its spreads falling from an average of 310 basis points in 2022 to 129 in the first ten months of 2025. Saudi Arabia, the United Arab Emirates, Qatar and Kuwait carry investment-grade ratings and we located no fall below investment grade for

any of them in the period and Saudi Arabia traded between 50 and 71 basis points across the 2021 to 2026 window. Of the six Gulf states, then, four carry investment-grade ratings with no fall below located in the period and two fell below and are recorded above, and the two are those whose fiscal positions are weakest. Bahrain's credit default swaps stood at 298 basis points on 4 August 2026, which the dated rule judges against the hard calibration, where it falls just below the 300 basis point floor and so does not meet the condition. That is not reconcilable with treating Bahrain as an ample-buffer sovereign, and Section 5.3 does not: on the Fund's published indicators Bahrain holds 2.4 months of prospective nonhydrocarbon import cover, the thinnest in the Gulf, with no published adequacy ratio and gross government debt at 134 per cent of GDP. Bahrain's 2025 Article IV staff report was withheld from publication by the authorities, so these figures come from the Executive Board press release.

Third, applying the Definition mechanically places every economy, and the placement contradicts what the earlier sections of this paper led us to expect. Table 4 does this. The rule is stated once, with its accessibility limb included, and applied to every observation without exception. Each dated observation in Table 3, and each point observation its notes identify, is judged against the calibration for its own regime: soft before 2022 and hard from 2022, the sample containing no 2022 issuance and the hardening being conventionally dated to the January 2022 renewal. The price condition counts as met where the spread reaches the top of the applicable range, indeterminate where it falls inside it, and not met below. Accessibility is read off the same data: a quoted spread is evidence that the sovereign was being priced in international markets, so condition (i) is satisfied at every date for which Table 3 reports a reading, and fails where no reading exists and the sovereign is separately recorded as without market access. Standing is as recorded in Table 5.

Table 4. Applying the Definition: every date at which the price condition is met

Economy	Dates at which price is met, with the reading in basis points	Accessible	Standing	All three met
Egypt	2001 (360), 2008 Q4 (385), 31 Dec 2021 (497)	Yes, spread quoted	Yes	**Yes**
Egypt, reported but outside the headline	2011 (607), 2013 (443), 2014 (351)	Yes, but see note on index membership	Yes	Not counted
Egypt, reported but outside the headline	Apr 2023 (1,434), Oct 2023 (1,277), Jan 2024 (990)	No — foreign currency rationed	Yes	No, on accessibility
Pakistan	2001 (1,115), 2008 Q4 (2,112), 2011 (1,274), 2013 (606), 2014 (525)	Yes	Yes	**Yes**
Jordan	2011 (500); indeterminate at late Jun 2025 (above 300)	Yes	Yes	**Yes**
Tunisia	2008 Q4 (464)	Yes	Yes	**Yes**
Morocco	2001 (518)	Yes	Yes	**Yes**
Iraq	2008 Q4 (1,282), 2011 (603), 2013 (511), 2014 (579), 31 Dec 2021 (507), 19 Oct 2022 (744)	Yes	Marginal, legacy external arrears	No, on standing
Oman	None; indeterminate at 31 Dec 2021 (254) and 2022 average (310)	Yes	Yes	No, on price

Economy	Dates at which price is met, with the reading in basis points	Accessible	Standing	All three met
Bahrain	None; indeterminate at 31 Dec 2021 (286) and 19 Oct 2022 (301)	Yes	Yes	No, on price
Saudi Arabia	None	Yes	Yes	No, on price
United Arab Emirates	No reading; rated investment grade	Yes	Yes	No, on price, by inference
Kuwait	No reading; rated investment grade	Yes	Yes	No, on price, by inference
Qatar	No reading; rated investment grade	Yes	Yes	No, on price, by inference
Djibouti	No reading; no market access	No	Yes	No, on accessibility
Algeria	No reading in either source	Not established	Yes	Not established

Notes. The thirteen economies Table 5 records as able to transact, plus Iraq, which Section 5.3 classifies as marginal; the six Table 5 records as unable to transact are omitted, their absence of cover carrying no information about price. Readings are from Table 3 and from the staff-report point observations recorded in this table and in Appendix A.5. The classification uses the corrected threshold $\pi/q - p$ of Table 1 Panel C at the reference parameters $q = 1$ and $p = 1/100$, which puts the applicable range at 150 to 340 basis points for observations before 2022 and 300 to 600 from 2022; the price condition counts as met at or above the top of the range, indeterminate inside it, and not met below. Accessibility is read off the same data: a quoted spread is evidence that the sovereign was being priced in international markets, so condition (i) is satisfied at every date for which Table 3 reports a reading, and fails where no reading exists and the sovereign is separately recorded as without market access. Standing is as recorded in Table 5. Egypt appears in three rows because its nine qualifying readings are not of one kind, and only the first row enters the headline count. The 2023 and 2024 readings are option-adjusted spreads on a different basis from either panel, as Appendix A.5 records, and they fall in the Egyptian foreign exchange crisis of import backlogs, a parallel-market premium and the March 2024 devaluation; a sovereign in that position is quantity-rationed in foreign currency, which is what condition (i) excludes, so reading accessibility off the existence of a bond-market quotation would be the very conflation Section 3.2 rejects. The 2011, 2013 and 2014 readings post-date the June 2011 removal from the EMBI that the note to Table 3 reports and cannot reconcile with the Fund's own statistical appendix; we print them and do not count them. Excluding all six leaves eleven observations across the same five economies, and that is the number the headline uses. Ten of the eleven are Panel A EMBI Global yield spreads and the eleventh, Egypt's reading of 497 at 31 December 2021, is a Panel B credit default swap. The two are different instruments and the count mixes them; we state this rather than suppress it. Dropping the credit default swap reading leaves ten observations across the same five economies and changes no conclusion, since Egypt qualifies on its 2001 and 2008 EMBI readings in any case. The United Arab Emirates, Kuwait and Qatar appear in neither spread source and are recorded as not meeting the price condition on their investment-grade standing throughout and on the Gulf credit default swap levels in Table 3 Panel B, which is an inference rather than an observation.

Because the threshold depends on both the hazard frequency and the trigger reliability, and neither is published for any economy in the region, the classification has to be reported as a surface rather than a point. Table 4a does that.

Table 4a. Sensitivity of the classification: observations and economies meeting

the price condition

Trigger reliability q	$p = 1/200$	$p = 1/100$	$p = 1/50$	$p = 1/30$
1.0	14 obs, 5 economies	17 obs, 5 economies	22 obs, 7 economies	27 obs, 7 economies
0.9	14 obs, 5 economies	14 obs, 5 economies	19 obs, 5 economies	25 obs, 7 economies
0.8	11 obs, 4 economies	13 obs, 5 economies	17 obs, 5 economies	23 obs, 7 economies
0.7	8 obs, 2 economies	8 obs, 2 economies	14 obs, 5 economies	17 obs, 5 economies

Notes. Counts are on the wide screen, that is dates meeting the price condition alone; on the strict screen the reference cell is eleven observations and the five-economy result is unchanged. Iraq is excluded throughout on standing. The five economies are Egypt, Pakistan, Jordan, Tunisia and Morocco; the seven add Bahrain and Oman, which cross from indeterminate into met at higher hazard frequencies; the four drop Tunisia and the two are Egypt and Pakistan alone. The count rises in p because the threshold is $\pi/q - p$, and falls in declining q because the observed margin is divided by a smaller number. Morocco's 2023 renewal fails the price condition in every cell of this table; its 2020 placement falls in the unobserved window and is not classified in any cell.

Three things follow, and the first cuts against the pricing argument the preceding sections build.

The band was occupied, repeatedly, by sovereigns that could have acted on it. Egypt, Pakistan, Jordan, Tunisia and Morocco each satisfied all three conditions at one or more dated observations between 2001 and the end of 2021, and none bought cover at any of them. Two counts are in play and they measure different things, so we keep them apart. Seventeen is the number of dated readings at which the **price condition alone** is met, which is what the second column of Table 4 records. Eleven is the number at which **all three conditions** are met: of the other six, three fail condition (i) outright, being Egypt's crisis-period option-adjusted spreads, and three rest on index readings whose provenance we cannot reconcile. Eleven is the figure the finding rests on, five of them Pakistan's and three Egypt's, and the five-economy result is the same on either count. The dates matter because they rule out the explanation that first suggests itself. Egypt's readings at 2001, 2008 and end-2021, Pakistan's at 2001, 2008, 2011, 2013 and 2014, Jordan's 2011, Tunisia's 2008 and Morocco's 2001 are ordinary market observations at which the sovereign was being quoted internationally with its contracting standing intact; they are not crisis readings whose spreads record the circumstances foreclosing a purchase. The price condition therefore does not explain this region's absence of sovereign catastrophe cover. Where the condition could be evaluated and was satisfied, nothing was bought, and the explanation has to be sought in Section 3.4's institutional and behavioural candidates.

The finding is robust across the parameter surface but it is not parameter-free, and Table 4a is where a sceptical reader should look first. The five-economy result holds in thirteen of the sixteen cells and fails in three. At a one-in-two-hundred-year frequency with $q = 0.8$ Tunisia drops out and four economies remain. At $q = 0.7$ with frequencies of one-in-a-hundred or rarer, which is two cells, the result collapses to Egypt and Pakistan alone, eight observations, and those are the two economies for which a reader might most readily argue that high spreads and impaired capacity to transact were the same episode. Seven economies appear in four cells, all of them at a one-in-fifty-year frequency or more frequent, where Bahrain and Oman cross in. We therefore state the finding as holding for five economies at the reference parameters and for no fewer than two anywhere on the surface, and we flag the $q = 0.7$ corner as the one that would most weaken it. What no cell of Table 4a changes is Morocco's 2023 renewal, which fails the price condition throughout.

Three qualifications attach to the result above. The thresholds are derived from Mexican, Chilean, Colombian, Jamaican and Peruvian contracts, so using them for MENAAP sovereigns transplants the hazard frequencies implicit in those placements; we checked the sensitivity and the five-economy classification is unchanged between one-in-five-hundred and one-in-a-hundred,

while at one-in-fifty and more frequent Bahrain and Oman cross in and the count rises to seven, as Table 4a records; the transplant should be stated. Section 4.4 screens the sponsors of those contracts against the same calibration, which removes the transplanted margin but not the transplanted frequency, and reports why that exercise does not bear on this finding either way. The observations are end-period index readings on an eight-date grid, not decision opportunities: sovereign cover is placed at renewals after years of preparation, fifteen years in Morocco's own chronology, so the claim is that the condition was satisfied at every date on which we can observe it and no cover was placed in the surrounding years, which is weaker than a claim about decisions. The sovereign-year panel Section 4.5 calls for is the object this finding properly requires, and we do not have it. And none of this establishes that cover was worth buying at all, since Section 3.3's participation condition is unevaluable.

Bahrain and Oman are not anomalies, though on a careless reading Bahrain looks like one. Neither reaches the met classification at the reference parameters: Bahrain's 286 and 301 basis points and Oman's 254 and 310 fall inside the indeterminate band, and their remaining observations fall below it. They cross into met only at a one-in-fifty-year frequency or more frequent, which is the seven-economy column of Table 4a. The error to avoid is judging a 2026 Bahraini reading against the 2017 to 2021 calibration; the dated rule removes that, and the dated rule leaves both short of the condition at the reference frequency, which is the calibration the headline uses, though not at the more frequent return periods Table 1 also tabulates.

The region's two buyers then sit as follows.

Morocco met all three conditions in 2001, bought in 2020 at a date for which no spread observation exists, and renewed in 2023 at a date that fails the price condition. Its EMBI Global spread stood at 518 basis points in 2001, above the soft-market range, with a quoted spread and unimpaired standing. It placed its parametric contract nineteen years later, in a year that falls inside the 2015 to 2021 gap in which no spread observation exists for any economy in this sample, so the 2020 placement cannot be judged against the threshold at all and we do not judge it. What can be judged is the 2023 renewal, at a credit default swap reading of 109 basis points, far below the hard-market floor of 300. The surrounding Panel B readings, which run from 73 to 292 basis points between the end of 2021 and August 2026, are all below the applicable range, the October 2022 peak of 292 included. Two qualifications attach and both matter. These are credit default swaps and the 2001 reading is an EMBI Global spread, so the comparison across the gap is a comparison across instruments. And the renewal is one dated observation, not a decision window. So what the record supports is narrower than it first appears: the one economy in the region that transacted at macroeconomic scale renewed at a date when the pricing case was weak, and passed over the date when it was strongest, with the original placement unobservable. That is not a point about Moroccan policy, which had no probabilistic hazard model in 2001 and no legal vehicle to contract through; it is a point about the pricing condition's power to explain observed behaviour, which on this evidence is low.

Djibouti falls outside the domain of the comparison. It has no spread quotation, no access to international capital markets, debt distress on the Fund and Bank's own assessment, and the market-financing stress test in its debt sustainability analysis marked not applicable. Condition (i) therefore fails and Proposition 1 does not apply to it. It has the standing to contract, which the purchase itself demonstrates, and it bought: annual limits of USD 2.2 million, five hundredths of one per cent of GDP, against perils capable of costing multiples of that, with a donor commitment of USD 2 million across the five-year term meeting the cost. The premium itself is not disclosed, so the subsidised share cannot be computed; what is on the record is who paid.

The case does not show that the sovereign for which cover is cheapest cannot afford it, because the framework does not establish that cover was cheapest for Djibouti; the comparison is undefined where neither instrument is accessible at the margin. It shows something narrower: a sovereign outside the domain entered the market because a third party paid, and the limit cannot

move a balance of payments. A price-only reading reaches the opposite conclusion by letting an unobservable spread stand in for an unbounded one, which Section 3.2 rejects. Section 7.1 draws the consequence for donor practice, which bears on Proposition 2 rather than Proposition 1.

Neither purchase, then, is explained by the substitution argument, and for different reasons: Morocco's because the price condition failed at the dates it transacted, Djibouti's because the price condition does not apply to it. A sovereign at 100 basis points that buys parametric cover is buying something other than cheaper reserve carry: a payout independent of loss adjustment, a pre-committed fiscal response that does not require a supplementary budget, a signal to domestic constituencies, or, in Morocco's case, the sovereign leg of a structure whose main purpose is compensating uninsured households through the solidarity fund. Those motives are real and this framework does not price them. It establishes only that they are not reserve substitution.

Iraq requires a separate note, because it met the price condition at six dated observations, more than any economy counted in the headline row, and is excluded on standing alone. Egypt has nine qualifying readings in all, but they are of three kinds and only the first row of three enters the count. Legacy external arrears are the ground: the Fund's 2025 Article IV records that Iraq's debt stock includes legacy arrears to non-Paris Club creditors and that comparable treatment has been requested and not obtained. The report gives no origination date, no amount and no named creditor, so we date the arrears no further than the word legacy supports. Section 5.3 records the classification as marginal rather than negative. Iraq is therefore neither available as a clean counter-example nor safely set aside, and a reader who thinks the arrears should not disqualify it would add a sixth economy to the five above and strengthen the finding rather than weaken it.

Tunisia is the thinnest of the five cases and should be read as such. It met the price condition at one observation, its 2008 fourth-quarter reading of 464 basis points, with a quoted spread and unimpaired standing, and it holds a proposed compulsory extension and no sovereign cover. Its inability to issue foreign-currency debt since 2019 does not add to that: on the argument of Section 3.2 a sovereign without market access falls outside the domain of the comparison, so the post-2019 period is silent rather than supportive. Tunisia is also the row that drops out first in Table 4a. The four other economies rest on more observations, and a reader who discounts Tunisia entirely is left with Egypt, Pakistan, Jordan and Morocco.

The resulting picture is not that the band is empty. It is that in this region the band is occupied by the sovereigns least able to act on it, and that the two sovereigns which did act were one below the price condition and one inside it but dependent on donors to transact at all.

4.4 The sovereigns that did buy, and why they do not settle the question

Section 4.5 calls for a panel of sovereign spreads against contemporaneous rates on line, and the obvious partial version of it can be run now: the sovereigns that placed the contracts in Table 2 published their prices, and their spreads at the placement dates are recoverable from official sources. We ran it. It does not settle the question in either direction, and the reasons it does not are worth more than the counts.

One route is closed at the outset and we record the decision, because it is the route a reader would expect us to take. Each tranche publishes a modelled expected loss, and since the unconditional payout probability is pq, setting $q = 1$ would make p equal to that published figure and give each buyer a threshold computed entirely from its own contract. We do not do this. The p of Proposition 1 is the frequency of an event large enough to force the sovereign to the external market, which Section 4.5 sets out at length and which is not the frequency of a damaging catastrophe; catastrophe risk modelling produces modelled losses and not that quantity. Reading p off a published expected loss would substitute the one for the other in the single place where the result is load-bearing. Table 4b therefore uses the same reference calibration as Table 4, which keeps the two screens comparable at the cost of leaving the frequency transplanted. What the exercise removes is the transplanted *margin*, since each threshold below is computed from

the price that sovereign actually paid, and not the transplanted frequency.

Table 4b. Remote-layer placements against the reference threshold

Placement	Date	Risk margin	Expected loss	Threshold	Spread at nearest month-end	Verdict
Mexico 2020	9 Mar 2020	350	90	250	372	Clears
Mexico 2024	17 Apr 2024	400	90	300	309	Indeterminate
Mexico 2018	7 Feb 2018	250	79	150	223	Clears
Colombia 2018	7 Feb 2018	300	156	200	178	Does not clear
Chile 2023	17 Mar 2023	475	100	375	153	Does not clear
Jamaica 2024	25 Apr 2024	700	150	600	150	Does not clear
Chile 2018	7 Feb 2018	250	86	150	106	Does not clear
Jamaica 2021	19 Jul 2021	440	152	340	Not established	—

Notes. All figures in basis points, ordered by spread. The eight remote layers of Table 2, that is the tranches with expected losses at or below 1.6 per cent; the fifteen intermediate and working layers are omitted on the ground given in Section 4.2, which is a statement about what a sovereign buying macroeconomic protection is buying and not a claim that sovereigns buy nothing else — the note below records that they do. The threshold is $\pi/q - p$ at the reference parameters $q = 1$ and $p = 1/100$, as in Table 4, and is a floor on both counts: the observed margin is divided by q, so any trigger short of perfect raises it. Spread readings are J.P. Morgan EMBI Global sovereign stripped spreads. They are month-end composites compiled by the ECLAC Washington Office, cross-checked against the monthly averages of daily differentials published in the Banco Central de Reserva del Perú series where that series carries the sovereign; the two aggregations agree within 11 basis points wherever both exist, and where a placement falls between month-ends the nearer reading is printed. The three placements of 7 February 2018 were priced out of one book on one day, each on its own sovereign's modelled loss; Section 4.2 clusters them as three transactions, which is what gives Table 2 thirteen clusters over twenty-three tranches, and we keep that convention here. They are nonetheless not three independent draws, sharing a book, a date and a market, and the eight rows below rest on four sovereigns in all. Mexico's 2024 transaction placed its remote earthquake layer on 17 April and its Pacific storm layer, a working layer, on 15 May; only the former is in scope. Jamaica's 2024 date is the World Bank pricing announcement of 25 April; the notes settled in May, and the verdict is unchanged on either month-end, the readings being 150 basis points at 30 April and 150 again at 31 December 2024. A placement counts as clearing where its spread exceeds the threshold by more than 15 basis points, as not clearing where it falls short by more than 15, and as indeterminate inside that interval, which is roughly the resolution of a month-end index reading against an intra-month placement; Mexico's 2024 placement, at 309 against a threshold of 300, is the one row the band decides. The 9 March 2020 Mexican placement falls at the start of the pandemic repricing, between a February month-end of 372 and a March month-end of 653; the February reading is printed because it is the conservative one at $q = 1$, and the sensitivity note below records where that choice becomes decisive. Jamaica's 2021 reading is not established, for the reasons given in Appendix A.10, which also records the vintages the other readings are drawn from and, for the two placements falling in years no listed vintage covers, says so. Cat bond coupons are fixed at price guidance some weeks before settlement, so no reading here is a price on the day of the decision.

Four objections apply to any reading of this table, and three of them are ours rather than a referee's.

The first is decisive and it is about domain, not data. Mexico, Chile and Colombia are inflation-targeting floaters whose public debt is predominantly local-currency and whose reconstruction is largely priced in their own currencies. Section 3.1's premise is a sovereign whose reconstruction is priced in a currency it does not issue, and Assumption 1 makes v approach unity as monetary sovereignty grows, which drives the participation condition $v/\lambda > \min\{M, 1+\delta/p\}$ to fail before the price condition is reached at all. For those three the comparison is not operative, and testing their price condition is the error this paper rejects for Djibouti, run in reverse: there we decline to place a sovereign deep inside the band when it lies outside the domain, and here we would be declaring sovereigns outside the band when the band does not apply to them. Jamaica is the one sponsor in the table for which the premise plainly holds, and it has one evaluable placement, which does not clear. One observation is not a generalisation.

The second is that the ordering in the table runs in the direction Proposition 1 implies, not against it. Ranked by spread the three highest observations are a clear, an indeterminate and a clear, and the four lowest all fail; the sovereign with the highest spread in the single Pacific Alliance book cleared and the two with lower spreads did not. Spread and threshold are not positively correlated across these placements, so the pattern is not an artefact of which layer each sovereign happened to buy. Seven observations separating almost monotonically is weak evidence, but it is weak evidence for the condition's ordering, and we record it as such rather than against it.

The third is the drawdown probability. Section 3.2 notes that if the marginal reserve unit is drawn with probability $z < 1$ in the disaster state the reserve arm becomes $1 + \delta/(pz)$, which favours cover, and sets z to one throughout for want of any basis for estimating it. Solving $\delta = pz[M(pq) - 1]$ for the value at which each verdict reverses, with the multiple the reference calibration implies rather than the contract's own, gives z of about 0.71 for Chile's 2018 placement, 0.41 for Chile's of 2023 and 0.25 for Jamaica's of 2024. Mexico's 2024 indeterminacy requires z above one, which is not a probability, so it does not resolve to a clear at any admissible value. The three negative verdicts are therefore reversible, but not uniformly and not cheaply: Chile's 2018 placement turns on a drawdown probability a little over two thirds, which is unexceptional, while Jamaica's 2024 placement requires the marginal reserve unit to be drawn only a quarter of the time. What the exercise shows is that the verdicts are not robust to a parameter the paper concedes it cannot estimate, and that the concession runs further for some placements than for others.

The fourth is coverage. Table 2 is the sovereign catastrophe bond record, not the sovereign risk transfer record. The regional pools publish portfolio rates on line — Appendix A.9 records the Caribbean facility's 2024–25 portfolio at 4.91 per cent against limit, with its earthquake and cyclone segment at 3.60 — and Appendix A.8 shows that a threshold can be computed from a rate on line alone, as it is for the Turkish pool. Those memberships are dozens of sovereign-years of published prices and we have not screened them. Until that is done, "the placements whose prices are public" means the bond record and not the universe.

What survives is narrower than the test was meant to deliver and it is not nothing. On the evidence available, the price condition is not visibly what moved these sponsors either: the sovereign that bought most often is the one that most often cleared, which is consistent with the condition, while the two placements furthest from their thresholds were made by sovereigns whose spreads were among the lowest in the sample. The honest statement is that the buyers neither confirm nor refute Section 4.3's finding, that three of the four lie outside the comparison's domain, and that the object the finding properly requires remains the sovereign-year panel of Section 4.5 rather than six placement dates. We report the attempt because the question is the first one a reader will ask of Section 4.3, and because a negative answer about what the available data can identify is more useful than a count presented as a result.

One placement in the 2026 record cannot be tested at all, and the reason is instructive. The Asian Development Bank's inaugural disaster relief bonds of May 2026, 80 million US dollars

each for the Kyrgyz Republic and Tajikistan at a 600 basis point risk margin, carry the Bank as obligor, the sovereigns being beneficiaries that receive payouts through national social protection systems. No sovereign credit enters the instrument and there is no spread against which to compare the margin. Neither sovereign has a quotation in any case: the Fund's 2025 Article IV for the Kyrgyz Republic records no Eurobond outstanding and uses an estimated yield in its debt sustainability analysis. The structure matters here because it detaches a placement from the accessibility condition of Section 4.1, which is the condition Djibouti's position turns on and the one no premium subsidy addresses. Section 7.4 returns to it.

4.5 Measurements that would locate a sovereign in the band

Converting the band from a qualitative claim into a mapping would require measurements that do not currently exist for this region.

A panel of sovereign spreads against contemporaneous rates on line would show directly how many sovereign-years in any region fall inside the band. Section 4.4 attempts the narrower question this section previously declined, namely whether the sovereigns that bought cover were inside the band when they bought, and reports that the available placements cannot settle it. What it cannot do is establish how often a sovereign-year falls inside the band without a purchase following, which requires the panel and not the six placement dates; the panel would also carry the dates between renewals, at which no transaction was available, and distinguishing those is itself unresolved.

Estimates of p, the frequency of an event large enough to force a sovereign to the market, would locate each sovereign on Table 1. This is not the frequency of a damaging catastrophe. It depends on the share of loss falling on the public balance sheet, on available fiscal space and on the external position when the event arrives, so a country with frequent hazard and deep buffers may have a low value of it. Catastrophe risk modelling as practised produces modelled losses, not this quantity, and nobody publishes it for any country in this region.

Estimates of F in money rather than in elapsed time would make Proposition 2 testable where it applies. Section 6.2 reports the fifteen years Morocco took and declines to infer a large fixed cost from it, but elapsed time is equally consistent with low political priority, competing fiscal demands or the length of a project cycle, and the one cost figure available, a grant facility of up to USD 5 million for Morocco's modelling, is small relative to the gains posited.

5. The MENAAP catastrophe cover inventory

5.1 Construction

No dataset records which sovereigns hold pre-arranged catastrophe cover, so we assembled one. The population is the twenty economies the World Bank assigns to its Middle East, North Africa, Afghanistan and Pakistan region and for which any institutional record exists: Morocco, Algeria, Tunisia, Libya, Egypt, Jordan, Lebanon, Iraq, Yemen, West Bank and Gaza, Saudi Arabia, the United Arab Emirates, Oman, Qatar, Kuwait, Bahrain, Iran, Pakistan, Afghanistan and Djibouti. Syria is excluded: it has been without accessible insurance-sector or fiscal reporting for the whole period, and including it would add a row that no source could populate. Türkiye is not in the region on the World Bank's classification and enters this paper only as an out-of-region comparator in Section 6.3.

For each economy we established the presence or absence of five instrument classes at any point between 2000 and 2026: a compulsory catastrophe extension to property or motor insurance; a national catastrophe pool or solidarity fund; a sovereign parametric or reinsurance contract; a catastrophe bond; and a World Bank Catastrophe Deferred Drawdown Option or comparable contingent credit line. Membership of a regional risk pool was checked for every economy and is reported in the text instead of as a column, since there is one instance.

An instrument counts as in force only if it meets three criteria. It must be pre-arranged, so a budget contingency reserve funded by annual appropriation does not qualify however large, because it is a claim on the sovereign's own resources, not on a counterparty. It must be in force rather than proposed. And it must operate at the sovereign layer or through a public entity, so private market penetration is reported separately, bearing as it does on the size of the uninsured loss rather than on the state's position.

Sources run in a fixed order of preference: national legislation and official gazettes; insurance regulators; World Bank project appraisal, restructuring and implementation completion documents; IMF Article IV and programme documents; and the Artemis deal directory for catastrophe bond issuance, whose record begins in December 1996 and which was searched by sponsor, by sponsor country and by peril for each of the twenty economies up to 4 November 2026, the accession date for the whole inventory. One sponsor-filtered page of that directory, for Türkiye after 2015, truncated before its final rows, and Appendix A.3 records that the corresponding negative finding is a not-found rather than a confirmed absence. Industry press was used only where no institutional source carried the fact. Every cell of Table 5 is sourced in Appendix A, together with the conflicts between sources and how each was resolved.

Negative findings required separate treatment, since most of the table consists of them. For the six Gulf states we checked four records: the complete catastrophe bond issuance history, by sponsor and by peril; membership of the three regional risk pools that exist, whose members lie in the Caribbean and Central America, Africa and the Pacific; eligibility for a Catastrophe Deferred Drawdown Option; and the most recent Article IV staff report for each, which was available for five of the six, Bahrain's 2025 report not having been published because the authorities withheld consent. These are not four independent confirmations. No regional pool is domiciled in the region, and joining one of the three that exist elsewhere requires a geographic fit that only Djibouti has, so non-membership carries little information; ineligibility for a Catastrophe Deferred Drawdown Option follows from high-income status, not from any choice, and is evidence about supply rather than demand; and a private reinsurance placement would not ordinarily appear in a surveillance document. The issuance record is the one strong negative. We report the absence as established for catastrophe bonds and as well supported but not proven for sovereign reinsurance placements.

5.2 The inventory

Table 5 reports the result. Across twenty economies and twenty-seven years, two sovereign risk transfer contracts were in force: Morocco's USD 275 million parametric contract and Djibouti's USD 2.2 million multi-peril policy with the African Risk Capacity. No MENAAP sovereign has sponsored a catastrophe bond, and no regional risk pool is domiciled in the region, and the single instance of a MENAAP sovereign joining one of the three that exist elsewhere is Djibouti's membership of the African Risk Capacity. Compulsory catastrophe extensions exist in two economies, Morocco and Algeria, and a funded national pool in one. One contingent credit line has existed, Morocco's, and it was drawn for the pandemic rather than for the 2023 earthquake; six of the Gulf states are ineligible for the instrument on income grounds. Six economies could not have contracted with an international counterparty for most or all of the period, and a seventh, Iraq, only marginally.

The density of the table is the first result. Of the one hundred instrument cells, excluding the transaction-capacity column, eighty-eight record an absence and one is unverified. Of the remaining eleven, three are proposals and two are appropriation-funded budget lines that do not meet the inclusion criteria of Section 5.1, leaving six cells in which an instrument was actually in force or drawn, and four of those six belong to Morocco. The question the rest of this section addresses is how much of that absence the pricing condition already accounts for.

Table 5. Sovereign catastrophe risk financing instruments, MENAAP economies, 2000–2026

Economy	Compulsory cat extension	National pool or fund	Sovereign risk transfer	Cat bond	Contingent credit	Able to transact
Morocco	In force (Law 110-14, 2020)	In force (FSEC, 2019)	In force to 2023 renewal (USD 275m parametric)	Proposed	Drawn 2020 (COVID)	Yes
Algeria	In force (Ord. 03-12, 2004)	Budget funds only	Not verified	None	None	Yes
Tunisia	Proposed	None	None	None	None	Yes
Egypt	None	None	None	None	None	Yes
Jordan	None	None	None	None	None	Yes
Saudi Arabia	None	None	None	None	Ineligible	Yes
United Arab Emirates	None	None	None	None	Ineligible	Yes
Oman	None	Proposed 2024	None	None	Ineligible	Yes
Qatar	None	None	None	None	Ineligible	Yes
Kuwait	None	None	None	None	Ineligible	Yes
Bahrain	None	None	None	None	Ineligible	Yes
Pakistan	None	Budget funds only	None	None	None	Yes
Djibouti	None	None	In force (ARC parametric, USD 2.2m)	None	None	Yes
Iran	None	None	None	None	None	No (inferred)
Lebanon	None	None	None	None	None	No (financial collapse)
Libya	None	None	None	None	None	No (contracting authority)
Yemen	None	None	None	None	None	No (conflict)
Iraq	None	None	None	None	None	Marginal
Afghanistan	None	None	None	None	None	No (sanctions, recognition)
West Bank and Gaza	None	None	None	None	None	No (no sovereign counter-party)

Notes. FSEC is the Fonds de Solidarité contre les Événements Catastrophiques, Morocco's public catastrophe fund. "Ineligible" denotes a high-income sovereign outside the IBRD lending categories, for which the Catastrophe Deferred Drawdown Option is structurally unavailable. "Budget funds only" denotes an appropriation-funded contingency vehicle, which does not meet the inclusion criteria in Section 5.1 and is shown for completeness and not counted as cover. Algeria's compulsory regime rests on Ordinance 03-12, enacted after the 2003 Boumerdès earthquake; the associated state guarantee to national reinsurers could not be confirmed on a primary source. Oman's national catastrophe pool was submitted to the regulator in March 2024; an assessment of observance of insurance core principles published in June 2025, following an on-site review in late 2024, does not record it as established. Morocco's parametric contract is recorded as in force to its 2023 renewal; its status for 2026 could not be verified from any source opened, as Appendix A.1 notes. Morocco's Catastrophe Deferred Drawdown Option was restructured in March 2020 to admit public health emergencies and drawn on 3 April 2020 for the COVID-19 response; it is a separate instrument from the parametric contract, which coincidentally carries the same limit, and it was not drawn for the 2023 earthquake. Djibouti signed a five-year multi-peril parametric policy with the African Risk Capacity on 19 January

2023, carrying annual limits of USD 1.4 million for extreme rainfall in Djibouti city and USD 0.8 million for nationwide drought, with USD 2 million of premium subsidy from the World Bank including USD 1 million through the Global Shield Financing Facility. It is recorded here as in force. Its combined limit of USD 2.2 million is 0.05 per cent of 2024 GDP, which is immaterial at the macroeconomic scale this paper is concerned with, and Section 5.3 treats it accordingly instead of excluding it. The final column records whether the sovereign could place risk with international counterparties during the period, on the grounds set out in Section 5.3. Sources for every cell are in Appendix A.

Counted against the criteria in Section 5.1, which require an instrument to operate at the sovereign layer or through a public entity, two economies qualify: Morocco, through its solidarity fund and parametric contract, and Djibouti, through its parametric policy. Algeria's compulsory extension and Morocco's are statutory mandates on private policies, not sovereign instruments, and are tabulated for completeness rather than counted here. Across the region there are two placed sovereign risk transfer contracts, but only one of macroeconomic scale: Morocco's limit of USD 275 million against Djibouti's USD 2.2 million, a ratio of 125 to 1. Djibouti's limit is five hundredths of one per cent of its GDP and cannot move a balance of payments, so the finding is better stated as one placed contract of any macroeconomic consequence in twenty economies. No MENAAP sovereign has ever sponsored a catastrophe bond. The only sponsor in the adjacent geography identified in the issuance record is the Turkish pool, and Türkiye falls outside MENAAP. No regional risk pool covers the region itself, and Djibouti's membership of the African Risk Capacity is the single instance of a MENAAP sovereign joining one of the three pools that exist elsewhere.

5.3 Partitioning the absence

A row of zeros carries no information until the zeros are sorted, and the framework in Section 3 sorts them.

Sovereigns that could not transact. Six of the twenty could not have placed risk with international counterparties for reasons the framework does not contain, and a seventh, Iraq, is marginal on the same grounds for part of the period. We write of seven below where the grouping is what matters and distinguish Iraq where its classification bears on the argument. Iran has been subject to sanctions regimes that, on the evidence of its complete absence from the issuance record and from every risk-financing diagnostic covering the region, have foreclosed placement with international reinsurers and issuance into the 144A market. We did not examine the underlying legal instruments and record this as a strong inference, not a legal finding. Lebanon has been through a currency collapse and has had no approved IMF arrangement since a staff-level agreement in April 2022, and its banking and public finances have not supported external contracting on commercial terms. Libya, Yemen, Afghanistan and West Bank and Gaza have lacked, for most or all of the period, a contracting authority recognised by counterparties, an underwritable risk, or a sovereign counterparty at all; Iraq is marginal on the same grounds for part of the period. Their absence of cover carries no information about premiums, spreads, buffers or fixed costs, and we exclude the six from everything that follows. Iraq is carried forward, because it is quoted in both spread sources and meets the price condition at six observations, and Table 4 records it as excluded on standing rather than silently dropped. This materially qualifies the headline: the finding is near-universal absence among thirteen transactable sovereigns, not among twenty.

Sovereigns for which cover is worth little. Table 6 reports the monetary position of the economies for which the question is live.

Table 6. Monetary position and buffers

Economy	Exchange rate arrangement	Reserves, months of imports	Reserves, % of IMF adequacy metric	Other liquid external assets	Cover
Saudi Arabia	Conventional peg since 1986	14.9 (2024)[a]	187 (2024)	PIF, SAR 3.42tn (2024)	None
United Arab Emirates	Conventional peg since 1997	7.6 (2024)[b]	109 (2024)	Not disclosed	None
Kuwait	Conventional peg to a composite	7.0 (2025)[c]	104 (2025)	KIA, USD 1,002bn, 640 per cent of GDP (2025)	None
Qatar	Conventional peg	8.0 (2024)[c]	62 (2024)	QIA, about 220 per cent of GDP (2023); QCB and QIA together about 600 per cent of the metric	None
Oman	Conventional peg since 1986	4.2 (2024)[c]	68 (2024)	OIA USD 53bn; liquid external buffers USD 28.4bn	None
Bahrain	Conventional peg	2.4 (2024)[d]	Not reported	Not disclosed	None
Morocco	Band of ±5 per cent	5.6 (2023)[e]	94.5 unadjusted, 122.1 adjusted (2023)	None material	Parametric, USD 275m
Pakistan	Managed float under programme	0.8 (FY2023), 2.3 (FY2025)[a]	17.9 (FY2023), 54.3 (FY2025)	None material	None

Notes. Türkiye is excluded from this table, as from the population in Section 5.1, and its figures are given where the comparator case is developed in Section 6.3. Import cover is taken from the relevant IMF staff report in every row, but the staff reports do not use a common denominator and the column should not be read as a single series. The superscripts record the basis: [a] next year's imports of goods and services; [b] next year's imports of goods and services net of re-exports, an adjustment specific to the United Arab Emirates; [c] prospective imports, coverage not specified in the source; [d] prospective imports of nonhydrocarbon goods and services, a narrower denominator than any other row; [e] prospective imports of goods and non-factor services. World Bank reserve series use a different denominator again and give materially lower figures for the Gulf; they are not spliced with these. Qatar is reported on two bases because the Fund assesses it on both: central bank reserves alone are 62 per cent of the metric, below the 100 per cent benchmark, and the Fund's adequacy judgement rests on central bank reserves together with the sovereign fund's foreign assets. Bahrain has no published adequacy ratio and no published sovereign asset figure, and its only published IMF indicators come from a press release, the 2025 staff report having been withheld from publication; it is treated separately below for that reason. Adequacy ratios are revised across report vintages; the Pakistani FY2025 figure is from the December 2025 review and the FY2023 trough from the 2024 Article IV, as Appendix A.4 records, and Turkish figures are from the 2026 Article IV, which revised the 2023 reading downward from above 100 per cent as reported in 2024. Morocco is reported on two ARA bases because the IMF publishes both and they differ materially: on the unadjusted metric Morocco stood below 100 per cent in 2023 and reached 102.8 per cent only at end-2025, while the metric adjusted for capital controls, which is the one the Fund uses to assess adequacy, places it comfortably above. Which variant is used changes the characterisation of the case, and both are used below. Sources and vintages are in Appendix A.

Saudi Arabia holds reserves at nearly twice the adequacy metric and close to fifteen months of

imports, alongside a sovereign fund reported at SAR 3.42 trillion at end-2024. On Assumption 1 the shadow value of disaster-state foreign currency for such a sovereign is close to unity whatever its exchange rate arrangement, the marginal value of protection is correspondingly low, and its absence of cover requires no further explanation. The United Arab Emirates and Kuwait clear the adequacy benchmark on central bank reserves alone, at 109 and 104 per cent of the metric, though both sit near the bottom of the 100 to 150 per cent range the Fund calls adequate rather than at Saudi Arabia's 187. For those two the ample-buffer reading rests less on the reserve ratio than on what stands behind it: Kuwait's investment authority is recorded by the Fund at 640 per cent of GDP and its net foreign asset position at 803 per cent, and the Emirati net international investment position at 178 per cent. The buffer classification is therefore a judgement applied case by case rather than a rule, and it reads sovereign fund assets as liquid for Kuwait and Qatar while discounting them for Oman, on the liquidity grounds set out in the Oman discussion below. An analysis conditioning on the exchange rate regime alone would have treated these hard pegs as the framework's most striking failures, and the buffer argument in Section 3.3 is what prevents that.

Qatar belongs in this group but not on the reserves-only measure, and the distinction matters because it recurs. Qatari central bank reserves are 62 per cent of the adequacy metric, which on a reserves-only reading would place Qatar below Oman. The Fund's own adequacy judgement is made on central bank reserves together with the investment authority's foreign assets, which it puts at about 220 per cent of GDP and the combination at roughly 600 per cent of the metric. We follow the Fund. The point to note is that for a Gulf sovereign with a large fund the reserves-only adequacy ratio is not the buffer that bears on Assumption 1, and the same ambiguity drives the Oman case below.

Bahrain. Bahrain does not belong in this group, though income and peg alone would put it there. The Fund's published indicators put Bahrain at 2.4 months of prospective nonhydrocarbon imports in 2024, the thinnest cover of any Gulf state and thinner than Morocco's; no adequacy ratio is published; no sovereign asset figure is published; and gross government debt stood at 134 per cent of GDP in 2024, projected to approach 150 per cent by 2026. Bahrain therefore belongs with the remainder below rather than with the ample-buffer group. It is not a case the price condition fails to explain, Table 4 recording it as indeterminate at two observations and met at none; what it does establish is that income and peg are poor proxies for a buffer, since at 2.4 months of cover Bahrain's external position is closer to Pakistan's than to Qatar's.

One qualification is owed. The 2025 Article IV staff report for Bahrain was not published, the authorities having withheld consent, so every figure here comes from the Executive Board press release rather than from a staff report, and the published series has been revised across vintages in a way that bears on any time-series use. The 298 basis point reading is also a single date in a series that stood at 181 and 187 at the two preceding observations.

Oman. Oman is the case the two readings disagree about.

On official reserves Oman is a thin-buffer sovereign: 4.2 months of import cover and 68 per cent of the adequacy metric at end-2024, a monetary position materially different from Saudi Arabia's despite the identical peg. On that reading its marginal value of protection is high, it holds none, and affordability cannot be the explanation: the Omani Investment Authority reported some USD 53 billion of assets, while remote-layer sovereign cover at the eight risk margins in Table 2 costs between USD 2.5 million and USD 7.0 million a year for each USD 100 million of limit, so a USD 500 million programme would run to between USD 12.5 million and USD 35 million annually; on the three hard-market margins alone the range is USD 20 million to USD 35 million. No premium for either instrument in Table 5 is disclosed, so this is an inference from the sovereign placements in Table 2, not a quotation for Oman. Oman would then be evidence for the fixed-cost account, and the chronology supports it: a pool submitted to the regulator in March 2024 and still not established by the regulatory assessment of June 2025. Section 4.3 qualifies this in both directions. The IMF reports Omani spreads averaging

310 basis points in 2022 and 129 over January to October 2025, and the credit default swap series puts Oman at 254 and 256 basis points at the end of 2021 and in October 2022. The two are different instruments and the first is a calendar-year average judged against a rule stated for dated observations, so neither is a clean reading. On the rule as stated, which counts the price condition as met only at or above the top of the applicable range, none of these readings meets it and all but the 2025 figure fall inside the indeterminate interval. Oman is therefore not in the band in the sense of the Definition at any observation; what the readings support is that it came closer in 2021 and 2022 than at any time since. Whether Proposition 1 held for Oman depends on which year the purchase is imagined in, and the paper cannot resolve it on a threshold stated as a range.

The difficulty is that the same USD 53 billion appears on both sides of that argument. If sovereign fund assets are part of the state's disaster-state foreign currency capacity, then Oman's effective buffer is large, the shadow value is close to unity, and Oman is simply the Saudi case: its absence of cover is efficient and evidences nothing about fixed costs. The affordability rebuttal and the thin-buffer premise cannot both draw on the same assets, and the argument above does.

A third figure, which the IMF reports and which neither reading has used, narrows the gap. Oman's total liquid external buffers, including the investment authority's liquid foreign assets, stood at USD 28.4 billion at end-2024 against official reserves of USD 18.4 billion. The relevant buffer is therefore neither the USD 18.4 billion of reserves nor the USD 53 billion of fund assets but something close to USD 28.4 billion, which is roughly 154 per cent of the official reserve stock and implies an effective adequacy ratio above 100 per cent instead of the 68 per cent the reserves-only measure gives.

On that figure Oman is closer to the ample-buffer group than to the thin-buffer one, and the case for reading its absence of cover as a fixed-cost failure correspondingly weakens. Further considerations bear on the question and do not all point the same way. Sovereign fund assets are not held for liquidity: a substantial share of the Omani Investment Authority's portfolio is in unlisted and strategic holdings whose disposal at short notice would be costly, whereas official reserves are liquid by construction. Drawing on the fund is a fiscal decision requiring government authorisation, while drawing on reserves is a monetary operation, so the two are not equivalent in the weeks after an event. Against that, a sovereign with USD 53 billion of external assets of any liquidity is plainly not foreign-currency constrained in the sense Pakistan was in 2023, and treating official reserves as the whole buffer overstates its exposure.

Our reading is that Oman sits between the two groups and that the evidence does not place it cleanly in either. We note in addition that on plausible parameters Proposition 1 fails for Oman, so that even on the thin-buffer reading its lack of cover may be correct instrument choice, not exclusion from the market. Oman therefore weakens the fixed-cost account rather than establishing it, and Section 7.4 shows that the group below does not rescue it either.

The remainder. The partition leaves five groups, not three, and they account for all twenty economies. Seven sovereigns could not transact. Four hold buffers against which cover is worth little. Morocco holds the region's only substantial sovereign instrument. Algeria is a group of one: it has the standing to contract, holds a compulsory extension that is a statutory mandate on private policies rather than an instrument of the state, and appears in neither spread source, so no observation places it with respect to the price condition and Table 4 records it as unestablished throughout. Morocco's buffer position is itself instructive: at 5.6 months of prospective imports and 94.5 per cent of the unadjusted IMF adequacy metric in 2023, it was not an ample-buffer sovereign on the measure the Fund publishes first, and reached 102.8 per cent unadjusted only at the end of 2025. That is the thinnest buffer of any economy here holding cover, so protection plausibly had real value to it — a judgement we rest on the buffer figures rather than on Assumption 1, whose ordering Section 3.3 confines to the ample-buffer tail. Section 4.3 shows the price condition nonetheless failed at every observation in the placement period.

The remainder, for which the question is live and the buffer is thin, are Pakistan, Egypt, Jordan, Tunisia, Bahrain, Djibouti and, on one reading, Oman. For this group we do not appeal to Assumption 1, for the reason given in Section 3.3: the ordering it generates is in v, the operative quantity is v/λ, and at the thin-buffer end nothing licenses the assumption that the ratio moves with the buffer. The partition here rests on the observable buffers themselves, which is weaker and sufficient. Section 4.3 shows the group divides again on price. How that group divides on price is set out in Section 4.3 and not repeated here: four of its members met all three conditions at one or more dates and none bought, Djibouti falls outside the domain of the comparison for want of market access, and Bahrain and Oman reach only the indeterminate classification. What this section adds is the buffer evidence behind those placements, and the point that the thin-buffer group is where the price condition most often held.

Of the seven, only Djibouti holds any cover, and its limit of USD 2.2 million against a 2024 GDP of USD 4.3 billion is too small to alter a balance of payments. Djibouti is also, as Section 4.3 sets out, outside the domain of the price comparison rather than deep inside the band, since it has no market access and neither instrument is available to it at the margin; what remains true is that a donor paid and that the scale of what it bought is the point rather than an aside. Djibouti is instructive for Proposition 2 in a direction that complicates it, since entry there was achieved with donor premium support, the instrument Section 7.4 deprecates, and it worked, though it bought no scale.

5.4 Limits of the cross-section

The event record is thin and unevenly reported. For several economies, Egypt and Jordan among them, no post-2000 event in these economies was found with a reported direct loss above half a per cent of GDP, a screen we applied by searching for a reported loss figure economy by economy in the sources listed in Appendix A rather than by querying a disaster-loss database, which this project could not reach; the screen is therefore a not-found rather than a measured absence, and the loss databases themselves note that damage figures depend on country self-reporting, so absence of a recorded loss is not evidence of absence of loss; we did not find a study quantifying that bias for this region. We therefore do not use the event count as a measure of exposure, and present the inventory as a cross-section of instruments and not as a matched sample of events and responses.

Loss measurement is not stable either. The 2023 Moroccan earthquake carries three published figures that are not commensurable. The government announced a forward reconstruction and upgrading budget of MAD 120 billion over five years, roughly USD 11.7 billion. An interregional computable general equilibrium exercise estimates the indirect output loss at 0.24 per cent of 2023 GDP, about MAD 3 billion; this is a think-tank modelling result rather than an official or peer-reviewed figure and should be read as such. Private-sector losses covered are reported at USD 40 million, attributed to the Moroccan regulator and the solidarity fund. The three stand in a ratio of roughly 300 to 7 to 1 and measure a forward budget, a modelled flow and an insured stock respectively. Pakistan's 2022 floods are USD 14.9 billion of damage, USD 15.2 billion of loss and USD 16.3 billion of reconstruction need, which summed pairwise produce the widely quoted and definitionally confused thirty-billion figure. Any threshold rule is sensitive to the measure chosen; the measure is named at each figure.

The deepest limitation is that with two placed contracts, only one of macroeconomic consequence, among thirteen transactable sovereigns, the cross-section cannot discriminate between the explanations in Section 3.4. One of the six, the fixed cost of market access, is refuted in Section 7.4 against the entry condition; the cross-section is consistent with each of the remaining five and separates none of them. Section 4.5 sets out what would discriminate.

6. Two cases and a comparator

The cross-section establishes absence. These cases show what absence cost in one instance, what presence delivered in another, and how far the instrument reaches even where it is most developed. Sources for every figure in this section are in Appendix A.

6.1 Pakistan

Pakistan entered the 2022 monsoon season with no sovereign catastrophe risk transfer of any kind. This is documented, not inferred. The World Bank's disaster risk financing diagnostic for Pakistan, prepared in 2020, states that no sovereign risk transfer was available and no contingency credit arrangements were in place, and describes sovereign parametric insurance and catastrophe bonds as options that a 2013 policy had directed should be explored. Household cover was negligible, at under one per cent of residential properties. Public assets are legally required to be insured through a single state provider, but only about thirty per cent of them carry cover, and that cover applies only during construction. The National Disaster Management Fund, established under the 2010 Act, held approximately PKR 1.5 billion, and the same diagnostic describes it as not fully operational.

The floods produced damage of USD 14.9 billion, which the post-disaster needs assessment puts at 4.8 per cent of FY22 GDP on its own GDP convention; losses of USD 15.2 billion; and reconstruction needs of at least USD 16.3 billion. Thirty-three million people were affected. No insured loss figure for the event appears in any primary source we could locate, and Swiss Re's catastrophe review for that year mentions the flooding without attaching a loss estimate to it.

The external sequence followed. By the close of FY2023 gross official reserves were USD 4,455 million, equivalent to 0.8 months of next year's imports of goods and services and 17.9 per cent of the IMF adequacy metric. We give the trough alone, because the Fund's selected-indicators table in the document cited at A.2 reports FY2023 onward on a stated basis and we were unable to source the FY2021 and FY2022 levels on the same basis. The sovereign came close to default and entered a programme, and reserves have since recovered to roughly 54 per cent of the metric by FY2025.

We do not claim that a parametric contract would have prevented this. The near-default had causes well beyond the floods and a counterfactual is not available. A timing point also constrains what the case can show. Cover for the 2022 floods had to be bought before them, at the spreads prevailing in 2020 or 2021, not at the distressed spreads of 2022 and 2023; by the time the external position had deteriorated to the levels above, no counterparty would have written the risk. Whether Proposition 1 was satisfied for Pakistan at the moment the purchase would have had to be made is therefore a question about pre-crisis spreads, and we do not have them. The narrower point stands: at the moment the reconstruction import bill arrived, the sovereign had one instrument for meeting it, that instrument stood at 0.8 months of cover, and the alternative had been identified as missing by its principal development partner two years earlier.

What changed afterwards indicates which margin moves. Pakistan is under a Resilience and Sustainability Facility arrangement approved in May 2025 with access of about USD 1.4 billion. The reform measure table of that arrangement contains one disaster risk financing item, the adoption of an implementation framework under a National Disaster Risk Financing Strategy, which carried a deadline of August 2026. No catastrophe insurance, risk transfer or contingency fund measure appears anywhere in it. Pakistan also lodged the first Global Shield request in Asia in November 2024, and the resulting package, finalised in May 2025, was to be delivered through a microfinance project and a proposed climate risk fund with a parametric backstop targeted for mid-2026. Both of those dates have now passed. We were unable to establish from the sources available whether either was met, and the status of the strategy document that the reform measure presupposes could not be verified. That gap matters, because the duration of

institutional build-out is the main evidence offered for the fixed-cost account, and Section 4.5 explains why duration is weak evidence for a cost parameter in any case.

6.2 Morocco

Morocco holds the region's only sovereign risk transfer contract of macroeconomic scale, Djibouti's being the other of the two placements, and the Al Haouz earthquake of 8 September 2023 tested it. The contract carries a limit of USD 275 million and the World Bank's completion report states that it disbursed that amount for the earthquake, within a total of USD 300 million met by the solidarity fund against USD 40 million met by the private sector. The report calls the instrument sovereign reinsurance; that it is parametric and referenced to peak ground acceleration comes from a World Bank conference presentation and from no document we were able to open. Against a government reconstruction programme of MAD 120 billion over five years the USD 275 million is a modest fraction, though it is roughly seven times what the private market met. The construction of that USD 40 million figure is not defined in the source and it should not be read as a measured insured loss.

A caution applies to the comparison this paper most wants to draw. The timing argument in Section 1 turns on parametric cover settling in weeks instead of the months an indemnity recovery takes, and we were unable to establish from any source we read when in 2023 the USD 275 million actually settled relative to the 8 September event. The case is therefore consistent with the timing claim but does not evidence it, and the settlement-speed premise should accordingly be treated as unverified.

The architecture around the contract is what the case illustrates. Law 110-14, adopted in 2016 and operational from 2020, makes catastrophe cover a compulsory extension to property, motor and bodily-injury liability policies, with per-event and annual caps set by the regulator at MAD 3 billion and MAD 9 billion for natural events. Alongside the private regime sits the solidarity fund, a public institution funded by a one per cent parafiscal charge on non-life policies, which raised over USD 90 million between 2020 and 2023 and compensates the uninsured, for whom residential insurance penetration is below three per cent. It is the only arrangement in the region addressing the household layer and the sovereign layer through one statute.

The chronology is long. The law was adopted in 2016; the implementing decree followed in 2019; the fund became operational at the end of that year; the parametric contract was placed in 2020 and renewed in 2023. The 2004 Al Hoceima earthquake is widely described as the impetus for the regime, and the World Bank project documents in Appendix A.1 date the programme's origins to the mid-2000s, but we found no source that states the causal link, nor one that dates the regime's origins to the earthquake itself, and Section 4.5 notes that elapsed time is weak evidence for a cost parameter, being equally consistent with low political priority, competing fiscal demands or the length of a project cycle; the one figure in money is the grant facility of up to USD 5 million, which Section 7.4 tests against the entry condition. Behind the sequence sit a national probabilistic hazard model, sustained World Bank and donor technical assistance, and a grant facility of up to USD 5 million signed in 2023 to fund the earthquake modelling, flood model update and parametric trigger feasibility work required for the next layer. Roughly fifteen years elapsed, with external assistance throughout. Section 4.5 returns to what that chronology does and does not establish about the size of the fixed cost.

Morocco is also the paper's selection problem. It is the one placement of macroeconomic consequence, and it is simultaneously the economy with fifteen years of sustained external assistance, a prior earthquake that created political demand, a middle-income administrative state and an insurance market able to carry a compulsory extension. Each of those characteristics is a plausible confounder for anything attributed to the cover itself.

Two features of the Moroccan programme are widely misreported. Morocco's Catastrophe Deferred Drawdown Option, approved in December 2019 at USD 275 million, was restructured in March 2020 to cover public health emergencies and drawn on 3 April 2020 for the COVID-19

response, disbursing USD 271,981,770 of the USD 275 million facility. It was not drawn for the earthquake, which was financed through the solidarity fund, and it is a distinct instrument from the parametric contract despite carrying the same limit. Second, Morocco has no catastrophe bond. One has been under discussion since 2019 and is funded for preparation, but none had been issued as of October 2026, and the country's FY2027 finance bill proposes preparing the regulatory framework for insurance securitisation.

6.3 Türkiye as comparator

Türkiye sits outside MENAAP on the World Bank's classification and enters as the nearest high-capacity comparator. Its pool, established by decree in December 1999 after the Marmara earthquake and governed since 2012 by the Catastrophe Insurance Law, provides compulsory first-loss earthquake cover to registered residential buildings within municipal boundaries; Gurenko et al. (2006) document its construction and its first five reinsurance renewals, and Appendix A.8 reports the rates on line from that programme, the only sustained catastrophe pricing series available for this geography. It held some 11.6 million active policies in mid-2026 and accumulated reserves of around TRY 24.5 billion, and it has sponsored two catastrophe bonds, of USD 400 million in 2013 and USD 100 million in 2015.

The February 2023 Kahramanmaraş earthquakes are the test. Direct physical damage was USD 34.2 billion, which the World Bank expressed as four per cent of 2021 GDP, the pre-event benchmark its damage assessment convention uses; damage and losses together were put at around nine per cent of GDP. Final industry insured loss on a property-lines basis was TRY 116.9 billion, roughly USD 6.2 billion at February 2023 rates. The ratio of the two is of the order of eighteen per cent, though we compute it ourselves from figures struck on different valuation dates in a rapidly depreciating currency, and no source states an insured share for the event. The pool paid TRY 39.1 billion across nearly 528,000 claim files, and this was the only recent event for which its reinsurance programme was called on.

Two ratios have to be kept apart here, and conflating them overstates what the pool itself achieved. The eighteen per cent is all property insurance against direct physical damage, most of it commercial and industrial cover written outside the pool. The pool's own payment of TRY 39.1 billion is about a third of the industry property-lines loss and, converted at February 2023 rates, on the order of six per cent of direct physical damage. Both figures are our own and rest on the valuation-date caveat above. What twenty-three years of compulsory cover, a dedicated state reinsurer, two catastrophe bond issuances and 11.6 million policies produced, then, is a compulsory residential scheme that met a single-digit share of the physical damage from the event it was designed for, within a property insurance market that met under a fifth of it. The limit is partly by construction: the pool writes first-loss residential cover to a capped sum insured and was never intended to indemnify commercial, industrial or infrastructure loss.

The external position around the event cautions against reading too much into any single episode. The central bank lost around USD 30 billion of reserves between January and May 2023, and the lira depreciated by some thirty per cent between the end of May and the end of July, but that adjustment followed the post-election policy turn and not the earthquake, and reserves measured against the adequacy metric moved only from 71.6 per cent in 2022 to 71.8 per cent in 2023. Attributing external movements to a disaster in an economy simultaneously undergoing a monetary regime change is the identification problem Sections 5.4 and 6.1 concede, in its sharpest form.

7. Policy implications

7.1 Participation and protection

The disaster risk financing agenda treats the absence of sovereign catastrophe cover as a gap to be closed. Its diagnostics record which countries lack instruments; its programmes are built to supply them; its results frameworks count them. Morocco's own World Bank operation sets targets in cumulative per-year volume of pre-arranged financing. The Global Shield initiative records requests lodged and packages delivered. The implicit premise throughout is that a sovereign without cover is a sovereign underprotected.

Section 4 gives reason to doubt the premise for part of this region and reason to sharpen it for the rest. If substitution requires a spread of the order Section 4 establishes and an investment-grade sovereign carries a fraction of that, then its absence of cover is not underprotection. But Table 4 also records five economies that cleared the price condition with their standing intact and bought nothing, and for those the agenda's premise is not refuted by the pricing argument at all; it survives, and the question becomes which of the frictions in Section 3.4 is doing the work. A programme counting instruments in place measures something other than protection, and the two diverge most where the condition fails.

There is a further difficulty with the premise and Section 4.4 bears on it, though less than we first thought. The agenda's implicit benchmark is the sovereign that did buy: participation is counted as the behaviour to be reproduced and absence as the deviation from it. Of the seven placements outside the region for which a spread can be sourced, two clear the reference threshold, one is indeterminate and four do not, so participation is at best a noisy indicator of the pricing condition. But three of the four sponsors are floaters for which the comparison is not operative at all, and among the seven the observations separate in the direction the condition implies rather than against it, so this does not support the stronger claim that participation and price are unrelated. What can be said is narrower: a programme counting instruments in place is not measuring protection, and the sovereigns it holds up as having acted are not demonstrably acting on the comparison this paper derives. Whether they are acting on it cannot be settled from six placement dates.

Three cases carry the point without relying on any contested parameter.

Four of the six Gulf states — Saudi Arabia, the United Arab Emirates, Kuwait and Qatar — hold no cover and need none on either the buffer or the price argument, yet they appear in regional protection-gap statistics as uninsured, which they are, and as underprotected, which on the evidence here they are not. Conflating the two overstates the size of the problem the agenda exists to solve, and it does so in precisely the economies whose uninsured losses are largest in absolute terms. The other two Gulf states cut the other way and are worth separating out rather than averaging in: Bahrain holds 2.4 months of import cover and government debt at 134 per cent of GDP, and Oman's official reserves are 68 per cent of the adequacy metric. Neither is an ample-buffer sovereign, and a regional statistic that treats the Gulf as one block gets both halves of this wrong.

Djibouti holds cover and is not protected by it, and the case is the one that most sharply inverts the agenda's premise. Djibouti has no access to international capital markets, so on the corrected argument of Section 3.2 it falls outside the domain of the price comparison rather than deep inside the band: where neither instrument can be acquired at the margin, the comparison is undefined rather than decisively in cover's favour. Table 4 records it as failing condition (i). What it was able to buy was annual limits totalling USD 2.2 million, five hundredths of one per cent of GDP, against perils capable of costing multiples of that, with a donor commitment of USD 2 million across the five-year term meeting the cost. The premium is not disclosed, so the subsidised share is not computable; what is on the record is that the World Bank and the Global Shield Financing Facility paid and the sovereign did not. By the agenda's own metrics Djibouti is a success, a MENAAP sovereign with a parametric policy and regional pool membership; by

any measure of protection it is not.

The agenda's working premise is that the sovereigns which should buy cover are prevented from doing so by a barrier an intervention can lower. Djibouti is the test case for that premise and it fails in the opposite direction from the one usually imagined. The barrier was not information, not capability and not the fixed cost of market access, all of which donors supplied. It was that the economy with the least capacity to fund a premium of any size has a GDP of USD 4.3 billion and no fiscal capacity to pay a catastrophe premium of consequence, so that the protection a donor can afford to subsidise is immaterial at the scale of the sovereign's exposure. That is not a barrier to entry. It is a statement about the size of the balance sheet, and no instrument design addresses it.

The mechanism at work here is not the one usually invoked. Premium subsidy does not relax a constraint that was preventing an otherwise worthwhile transaction. It overrides the price signal that was correctly discouraging one. A subsidy covering most of the premium makes cover cheap to the sovereign whatever the unsubsidised comparison says, so the sovereign buys, the programme records an instrument in place, and the transaction reverts to being uneconomic the moment the subsidy lapses. Subsidy also cannot buy scale, since each additional unit of limit carries its own premium while the fixed cost of access is paid once. The combination produces what Djibouti shows: entry without protection, and entry that is not self-sustaining.

Pakistan is the region's clearest case of a sovereign whose external position a catastrophe overwhelmed, and the timing of the comparison matters. At 0.8 months of import cover and 17.9 per cent of the adequacy metric in FY2023, Pakistan was in no position to fund a premium in foreign currency, whatever its spread; no Pakistani spread after 2014 appears in Table 3, so we make no claim about where it stood relative to the threshold at that date. But cover for the 2022 floods had to be purchased before them, at 2020 or 2021 spreads, and we do not know whether the condition held then. What the case shows is narrower than it first appears: the window in which a sovereign both satisfies the pricing condition and retains the counterparty standing to act on it can close faster than an institutional build-out takes.

Taken together these say the agenda is aimed at the wrong margin in this region. It promotes an instrument that is correctly declined by the sovereigns that can afford it, supplies it at trivial scale to those that cannot, and is unavailable to the one sovereign for which the economics most clearly favoured it.

We are not arguing that sovereign risk transfer is a bad instrument, that the countries which have built it were wrong to, or that the protection gap is imaginary. Morocco's programme delivered: the World Bank records USD 275 million of reinsurance recovery for the earthquake, within USD 300 million met by the solidarity fund against USD 40 million met by the private sector. Section 6.2 notes that the settlement date is not in the record, so the speed of that delivery is not established, only its occurrence and its size. The household layer is a genuine and large gap, and nothing in this paper bears on it. The claim is narrower: the sovereign layer in MENAAP is mostly not where the gap is, the metrics in use cannot tell the difference, and a programme that counted protection instead of participation would allocate differently.

7.2 Two unpublished quantities

Section 4.5 identifies two of the three objects a sovereign would need in order to know whether it is in the band: the annual probability that a catastrophe forces it to the external market, and the price at which it could transfer that event. Neither is published for any country in this region. Of the two placements here, neither discloses its premium, and the Turkish pool states explicitly that its current reinsurance pricing is not disclosed.

A surveillance report that carried both, alongside reserve adequacy, would let a finance ministry answer the question this paper poses for itself in an afternoon. Producing them for a region is a technical assistance task, and it is a cheaper one than arranging the instruments whose value they would establish.

7.3 Surveillance and disclosure

Standard reserve adequacy metrics weight short-term external debt, broad money, exports and other liabilities, and contain no term for catastrophe exposure. Two sovereigns with identical financial vulnerability and radically different hazard exposure are assessed as equally covered.

The recommendation is narrow and implementable: report pre-arranged catastrophe cover as a memorandum item in the reserve adequacy assessment itself, alongside the assessing-reserve-adequacy ratio in the Article IV staff report's external sector annex and in the selected-indicators table, giving notional, trigger, term and counterparty. The addressees are Fund surveillance staff, who already compile that annex, and the Financial Protection Forum and Global Shield secretariats, which hold the contract-level information and do not currently publish it in a form surveillance can use. It requires no change to any metric and costs nothing to adopt. It does not depend on whether any sovereign should buy cover, because it concerns measurement and not purchase; Section 7.4 sets out why Section 4's finding strengthens the case for it rather than weakening it.

Cover should not, however, be credited at par against a net international reserve floor. A floor is a liquidity constraint binding at short horizons; a parametric contract is a contingent claim that may not trigger, carries counterparty and collateral risk, and settles on a timetable this paper has not been able to document for a single contract. Recognition at a haircut reflecting trigger probability, counterparty rating and documented settlement speed, applied against medium-term accumulation paths and not short-term floors, is defensible. On present disclosure almost no contract would qualify, which is an argument for requiring settlement-timing disclosure as a condition of recognition.

Surveillance that reported liquid external assets by holder and by encumbrance would also settle questions the present arrangement hides. Section 5.3 could not resolve Oman's buffer position because official reserves, the investment authority's total assets and the IMF's liquid external buffer measure give three different answers, and the second rests on a press release rather than a published statistical return.

7.4 The composition of concessional support

For the minority of sovereign-years inside the band, the composition of concessional support matters and Proposition 2 governs it. Proposition 2 is asymmetric between the two kinds of support, and Djibouti illustrates the premium clause. USD 2 million of donor premium support induced a purchase the unsubsidised comparison would not have supported, which is what the proposition predicts and what makes premium support the more questionable instrument.

The argument requires the fixed cost to be large relative to the gains, and the one figure on the record, a grant facility of up to USD 5 million for Morocco's modelling work, is small. That figure is a bound rather than a measurement, and the test below is whether the account survives even at it. Regional probabilistic hazard modelling is the component most obviously suited to collective provision, and intermediation by an international financial institution has precedent where a sovereign lacks standing or ratings to access the market directly. That precedent has recently been extended in a way that bears directly on the accessibility condition. The Asian Development Bank's disaster relief bonds of May 2026 place catastrophe risk for the Kyrgyz Republic and Tajikistan, 80 million US dollars each at a 600 basis point risk margin, with the Bank as obligor and the sovereigns as beneficiaries receiving payouts through national social protection systems. No sovereign credit enters the instrument. Neither sovereign has a traded spread against which a margin could be compared, and the Fund's 2025 Article IV for the Kyrgyz Republic records no Eurobond outstanding and uses an estimated yield in its debt sustainability analysis. The structure matters here because it detaches a placement from condition (i) of Section 4.1, which is the condition Djibouti's position turns on and the one no premium subsidy addresses. Whether it also addresses the scale problem of Section 7.1 is a separate question and

this transaction does not answer it: 80 million US dollars is some thirty-six times Djibouti's limit and still small against a sovereign balance sheet. We record the structure as the one development in the period that acts on a condition the pricing argument treats as a domain restriction rather than a price. Regional pooling attacks the fixed cost by spreading modelling and operational costs, though the region's dominant peril, earthquake along the North African and Anatolian systems, is a correlated peak peril of the kind pools diversify least well, and we have no correlation estimates for any pair of regional perils.

None of this applies where the bracket in the entry condition is negative. Section 4.3 establishes that for Saudi Arabia, which has readings in Table 3, and infers it for the United Arab Emirates, Kuwait and Qatar, which appear in neither spread source. There the negative-bracket case of Proposition 2 governs, and the correct recommendation is to accumulate reserves and build fiscal space. Where the bracket was positive, as it was for the five economies in Table 4, the entry condition can be evaluated rather than asserted, and it does not support the fixed-cost account. Two conventions are available for the threshold and we report both, because the answer is not identical under them. Taking the fitted point threshold of 193 basis points from Table 1 Panel A, the bracket runs from 167 basis points at Egypt's 2001 observation to 1,919 at Pakistan's of the fourth quarter of 2008, and the capitalised gain $q[\delta - p(M(pq) - 1)]KA$, at a fifteen-year horizon discounted at five per cent, is between USD 17 and USD 199 million on a USD 100 million programme and between USD 48 and USD 548 million on a programme the size of Morocco's. Taking instead the top of the range that Table 4's own screen uses, 340 basis points, the brackets narrow to 20 to 1,772 basis points and the gains to USD 2 to 184 million at the smaller notional and USD 6 to 506 million at the larger.

Against either, the only fixed cost on the record anywhere in this literature is the grant facility of up to USD 5 million for Morocco's modelling work. On the point threshold the entry condition is satisfied at every one of the eleven observations, by a factor of between three and forty. On the top-of-range threshold it is satisfied at every observation on a Morocco-sized programme, and at nine of eleven on a USD 100 million one, the two exceptions being Egypt's 2001 and 2008 readings, which sit closest to the threshold. Proposition 2 therefore predicts entry at every qualifying observation on the first convention and at all but two on the second.

Entry did not occur. The fixed-cost account therefore fails on the paper's own arithmetic: it would require a fixed cost several times, and at most observations one to two orders of magnitude, above the only documented figure, and we have no evidence of one. The conclusion to draw is not that some other friction in Section 3.4 is established instead. It is that none of the remaining five is, and that the non-purchases are a puzzle this framework poses and cannot close. Unplaceable risk, religious objection to conventional insurance and securitisation, charity hazard, the budgetary salience of a premium paid against a contingent benefit, and weakened incentives for ex ante risk reduction all remain live, and discriminating between them needs the premium and frequency data Section 7.2 says nobody publishes.

That is a weaker policy conclusion than a named friction would be, and it is the one the evidence carries. It also sharpens rather than blunts the surveillance recommendation below: a condition that was satisfied repeatedly and acted on in no case is exactly the situation in which the two unpublished quantities are worth producing, because without them nobody can say whether the sovereigns concerned were leaving money on the table or declining a transaction for reasons the price comparison does not see.

7.5 Limits of the argument

The substitution result is easy to over-read, and four limits bound it.

Catastrophe cover does not substitute for reserves in general. Reserves meet obligations in every state of the world; cover pays on one contingency subject to a trigger. The substitution is specific to the disaster state and partial even there, and Proposition 1 quantifies how partial.

Moral hazard is absent from the framework and should not be. Pre-arranged cover may

weaken incentives for building-code enforcement, land-use restriction and other ex ante risk reduction, and for a sovereign the political economy runs the same way: a funded contingency reduces the salience of prevention. Parametric cover mitigates this at the margin, as noted in Section 3.4, since the payout is independent of realised loss and the sovereign retains the full benefit of reducing damage. Indemnity cover and solidarity funds compensating uninsured households carry the problem in full. We have not modelled it, and a complete treatment of optimal sovereign cover would have to.

Nor does cover substitute for adaptation. These instruments allocate the cost of a loss and do not reduce it. Where the expected loss is rising, the premium rises with it, and at some point the hazard becomes uninsurable at any price a sovereign would pay. Risk transfer buys time to adapt and does not replace adaptation.

Finally, nothing here resolves the protection gap at household level. A sovereign parametric contract compensates the state, not the uninsured homeowner, and the Turkish case in Section 6.3 shows that twenty-three years of compulsory household cover met a single-digit share of the physical damage from the event it was built for. Morocco's dual structure is the only arrangement in the region that addresses the household layer directly, and the distributional question of who ultimately bears the loss is left open by every instrument discussed here.

8. Conclusion

A catastrophe in a sovereign that cannot issue the currency its reconstruction is priced in produces an immediate demand for foreign exchange and a delayed, partial supply of it. Reserves and parametric cover are competing ways of meeting that demand, and no literature has compared them directly. This paper does, and then asks how much of the observed behaviour the comparison explains. The answers to those two questions point in different directions, and saying so is the paper's main contribution.

On the comparison itself the conclusion is firm. Cover is cheaper when the annual probability of a qualifying event, multiplied by the market's loading over expected loss at that frequency, falls below the carry cost of reserves. Neither of the objections usually raised rescues cover: basis risk is not a separate line of cost but it is not free, since a thinner layer is priced at a higher multiple and the cost per unit of expected delivery rises; and the fungibility of reserves nets out only at an interior optimum, so where a buffer is thin the credit is positive and reserves are cheaper still. Both corrections run against cover. Estimating the pricing function on twenty-three sovereign tranches puts the elasticity of the multiple in expected loss at -0.46 pooled and -0.43 once the reinsurance cycle is in the specification and the standard errors are clustered by transaction, and inverting the condition at a one-in-a-hundred-year frequency puts the threshold at 150 to 340 basis points at soft-market prices and 300 to 600 at hard-market prices. The level of that threshold moves 55 per cent with the reinsurance cycle, so it is not a property of a sovereign's balance sheet at all.

On explanatory power the conclusion is the opposite, and we reached it against our own expectations. Applying the condition to every dated observation in the region, Egypt, Pakistan, Jordan, Tunisia and Morocco each satisfied it at one or more observations between 2001 and the end of 2021 while quoted in international markets with their contracting standing intact — eleven observations on the strict screen in Table 4, seventeen on the wider one — and none bought cover at any of them. These are not crisis readings whose spreads record the circumstances foreclosing a transaction, which is how they are usually disposed of; the strict screen exists precisely to remove the readings to which that objection would apply. Morocco is the sharpest case, within the limits of what is observable: it met the condition in 2001 at 518 basis points, placed its contract nineteen years later at a date for which no spread exists in either panel, and renewed in 2023 at 109 basis points against a floor of 300. The price comparison therefore does not explain this region's absence of sovereign catastrophe cover. Where it could be evaluated

and was satisfied, nothing was bought.

Whether the sovereigns that did buy were inside the band when they bought is the first question this invites, and Section 4.4 reports that the available data cannot answer it. Of the seven placements outside the region with a sourceable spread, two clear the reference threshold, one is indeterminate and four do not; but three of the four sponsors are inflation-targeting floaters whose reconstruction is priced in currencies they issue, so Section 3.1's premise and with it the participation condition do not hold for them, and the single sponsor plainly inside the domain has one evaluable placement. The observations also separate almost monotonically in the direction Proposition 1 implies, and every negative verdict reverses at a reserve-drawdown probability we concede we cannot estimate. The regional finding therefore stands as a regional finding, and the sovereign-year panel of Section 4.5 remains the object it requires.

Both findings survive together, and the title claim is the first of them. Sovereign catastrophe cover in this region is not reserve substitution: whatever Morocco purchased in 2020, and the instrument did disburse USD 275 million after the 2023 earthquake, it was not a cheaper way of holding foreign exchange. What does not survive is the deflationary corollary, that the region's near-absence of cover is therefore what the pricing condition predicts. For the four Gulf states with ample buffers it is. For the five economies above it is not. Section 7.4 tests the obvious candidate, the fixed cost of market access, against the entry condition and it fails: at the only documented fixed cost the capitalised gain exceeds it by a factor of three to forty at every qualifying observation, so Proposition 2 predicts entry. The remaining candidates in Section 3.4 are untested rather than excluded.

Two measurement gaps stand behind all of this and they are what a central bank can act on. Nobody in the region publishes the frequency with which an event would force the sovereign to the external market, which is not the frequency of a damaging catastrophe. Nobody publishes the price at which a sovereign could transfer that event: neither of the region's two placements discloses a premium, and the Turkish pool states that its current reinsurance pricing is not disclosed. A finance ministry here cannot evaluate the condition this paper derives, which is a fact about the region's data rather than about its judgement, and the surveillance change in Section 7.3 is cheaper than any of the instrument promotion that currently substitutes for it.

What the regional protection gap is evidence of therefore divides by layer. At the household layer it remains a question of insurance market development, and the Turkish case shows how far even twenty-three years of compulsory cover falls short. At the sovereign layer it is neither the market failure the disaster risk financing literature assumes nor the efficient outcome a pricing argument alone would suggest. It is a gap between a condition that can be written down and a decision nobody in the region is currently equipped to take.

References

Aizenman, J. and Lee, J. (2007) 'International Reserves: Precautionary versus Mercantilist Views, Theory and Evidence', *Open Economies Review*, 18(2), pp. 191–214. doi:10.1007/s11079-007-9030-z

Aslan, I. (2026) 'Monetary Sovereignty and Catastrophe Fiscal Capacity: Turkey, Japan, and the Limits of Disaster Self-Insurance', *Review of Political Economy*. doi:10.1080/09538259.2026.2724127

Aslan, I. (forthcoming) 'Why Risk-Based Premiums Will Not Close the Catastrophe Protection Gap: Conventional Pricing, Reinsurance Capacity, and Public Scheme Design', *Risk Management and Insurance Review*.

Bianchi, J. and Sosa-Padilla, C. (2024) 'Reserve Accumulation, Macroeconomic Stabilization, and Sovereign Risk', *The Review of Economic Studies*, 91(4), pp. 2053–2103. doi:10.1093/restud/rdad075

Borensztein, E., Cavallo, E. and Valenzuela, P. (2009) 'Debt Sustainability Under Catastrophic

Risk: The Case for Government Budget Insurance', *Risk Management and Insurance Review*, 12(2), pp. 273–294. doi:10.1111/j.1540-6296.2009.01168.x

Braun, A. (2016) 'Pricing in the Primary Market for Cat Bonds: New Empirical Evidence', *Journal of Risk and Insurance*, 83(4), pp. 811–847. doi:10.1111/jori.12067

Cavallo, E., Galiani, S., Noy, I. and Pantano, J. (2013) 'Catastrophic Natural Disasters and Economic Growth', *The Review of Economics and Statistics*, 95(5), pp. 1549–1561. doi:10.1162/REST_a_00413

Cebotari, A. and Youssef, K. (2020) *Natural Disaster Insurance for Sovereigns: Issues, Challenges and Optimality*. IMF Working Paper No. 2020/003. Washington, DC: International Monetary Fund. doi:10.5089/9781513525891.001

Clarke, D.J. and Dercon, S. (2016) *Dull Disasters? How planning ahead will make a difference*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780198785576.001.0001

Cummins, J.D. and Mahul, O. (2009) *Catastrophe Risk Financing in Developing Countries: Principles for Public Intervention*. Washington, DC: The World Bank. doi:10.1596/978-0-8213-7736-9

Eichengreen, B., Hausmann, R. and Panizza, U. (2005) 'The Pain of Original Sin', in Eichengreen, B. and Hausmann, R. (eds) *Other People's Money: Debt Denomination and Financial Instability in Emerging Market Economies*. Chicago, IL: University of Chicago Press, pp. 13–47. doi:10.7208/chicago/9780226194578.003.0002

Gurenko, E., Lester, R., Mahul, O. and Gonulal, S.O. (2006) *Earthquake Insurance in Turkey: History of the Turkish Catastrophe Insurance Pool*. Washington, DC: The World Bank.

Hsiang, S.M. and Jina, A.S. (2014) *The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones*. NBER Working Paper No. 20352. Cambridge, MA: National Bureau of Economic Research. doi:10.3386/w20352

Ilzetzki, E., Reinhart, C.M. and Rogoff, K.S. (2019) 'Exchange Arrangements Entering the Twenty-First Century: Which Anchor Will Hold?', *The Quarterly Journal of Economics*, 134(2), pp. 599–646. doi:10.1093/qje/qjy033

Jeanne, O. (2016) 'The Macroprudential Role of International Reserves', *American Economic Review*, 106(5), pp. 570–573. doi:10.1257/aer.p20161013

Jeanne, O. and Rancière, R. (2011) 'The Optimal Level of International Reserves for Emerging Market Countries: A New Formula and Some Applications', *The Economic Journal*, 121(555), pp. 905–930. doi:10.1111/j.1468-0297.2011.02435.x

Lane, M.N. and Mahul, O. (2008) *Catastrophe Risk Pricing: An Empirical Analysis*. Policy Research Working Paper No. 4765. Washington, DC: The World Bank. Available at: http://hdl.handle.net/10986/6900

Nguyen, A.T.-N. and Nguyen, H.M. (2025) 'Currencies in turbulence: exploring the impact of natural disasters on exchange rates', *Journal of International Money and Finance*, 157, 103397. doi:10.1016/j.jimonfin.2025.103397

Noy, I. (2009) 'The macroeconomic consequences of disasters', *Journal of Development Economics*, 88(2), pp. 221–231. doi:10.1016/j.jdeveco.2008.02.005

Obstfeld, M., Shambaugh, J.C. and Taylor, A.M. (2010) 'Financial Stability, the Trilemma, and International Reserves', *American Economic Journal: Macroeconomics*, 2(2), pp. 57–94. doi:10.1257/mac.2.2.57

Ramcharan, R. (2007) 'Does the exchange rate regime matter for real shocks? Evidence from windstorms and earthquakes', *Journal of International Economics*, 73(1), pp. 31–47. doi:10.1016/j.jinteco.2006.12.004

Rodrik, D. (2006) 'The Social Cost of Foreign Exchange Reserves', *International Economic Journal*, 20(3), pp. 253–266. doi:10.1080/10168730600879331

Strebel, B. (2024) Remarks on the MENA protection gap, 34th General Arab Insurance Federation conference, Muscat, reported in *Global Reinsurance*. Available at: https://www.globalreinsurance.com/home/temperatures-are-rising-and-menas-protection-gap-lea

ds-the-world-swiss-re-at-gaif34/1451236.article

von Peter, G., von Dahlen, S. and Saxena, S.C. (2024) 'Unmitigated disasters? Risk sharing and macroeconomic recovery in a large international panel', *Journal of International Economics*, 149, 103920. doi:10.1016/j.jinteco.2024.103920