ADVANCE SOCIAL SCIENCE ARCHIVE JOURNAL
Available Online: https://assajournal.com
Vol. 05 No. 01. Jan-March 2026.Page#.6887-6910
Print ISSN: 3006-2497 Online ISSN: 3006-2500
Platform & Workflow by: Open Journal Systems
https://doi.org/10.5281/zenodo.22642436

Financing Health under Climate Shocks: Risk Pricing, Data Poverty and Financial Protection in Pakistan

Mohammad Asadullah Soomro
Bahria University
asadullahsoomro41@gmail.com
Sabeen Yaqoob
Iqra University
sabeen.yaqoob@iqra.edu.pk
Omama Khan
Mohammad Ali Jinnah University
omamakhani173@gmail.com
Saima Khan
Mohammad Ali Jinnah University, Karachi, Pakistan
saimayounuskhan03@gmail.comm

Abstract
Climate shocks in low-income settings are usually accounted for as damage to physical assets. This paper argues that in health systems financed at the point of use, the larger and more persistent welfare loss falls on household budgets through out-of-pocket medical spending, and that this loss is systematically invisible to the post-disaster accounting on which reconstruction finance is based. Pakistan's 2022 floods are the motivating case: the Post-Disaster Needs Assessment valued damage to health facilities at USD 109 million, roughly seven tenths of one per cent of assessed physical damage, while destroying an eighth of the country's health infrastructure and coinciding with a fivefold rise in malaria incidence in a system where households already financed most care directly. We develop a framework in which two frictions govern whether such losses are pooled. The first is a pricing friction: where the price of climate risk is administratively suppressed or never established, cover is not written. The second is a data friction: where screening is delegated to statistical models, applicants with sparse records are rationed out, a mechanism documented for venture capital by Bonelli (2026) and for clinical risk prediction by Obermeyer et al. (2019). Our central proposition is that in a thin-data health economy these frictions are sequential rather than parallel. A price never charged generates no claims record, and the absent record is precisely what a later algorithm would require, so the same districts are excluded twice. We derive three hypotheses and specify an identification strategy that is executable with existing public microdata: a difference-in-differences design on district flood exposure using the Household Integrated Economic

Survey rounds of 2018–19 and 2024–25, a triple-difference exploiting Punjab's 2025 termination of Sehat Sahulat cover in public hospitals, and a prospective audit of algorithmic underwriting. We set out a four-part design response that couples actuarially priced parametric cover to Pakistan's existing social health insurance and instant-payment infrastructure, and assess its fiscal incidence. The paper is a research design and policy analysis; it reports no estimates of its own, and labels its propositions accordingly.
Keywords: *health financing; catastrophic health expenditure; financial protection; climate shocks; social health insurance; parametric insurance; algorithmic underwriting*
JEL classification: *I13, I18, H51, H84, Q54, O16*

1. Introduction

Post-disaster accounting counts what can be seen. When the Indus basin flooded in the summer of 2022, the joint assessment prepared by the Government of Pakistan with the Asian Development Bank, the European Union, the United Nations Development Programme and the World Bank valued physical damage at USD 14.9 billion, wider economic losses at USD 15.2 billion, and resilient reconstruction needs at USD 16.3 billion (Government of Pakistan et al. 2022). Housing accounted for USD 5.6 billion, agriculture and livestock for USD 3.7 billion, transport and communications for USD 3.3 billion. Damage to health facilities was assessed at USD 109 million, with service losses of USD 34 million and recovery needs of USD 188 million. Together these come to roughly seven tenths of one per cent of assessed physical damage.

That figure is accurate and, as a measure of the health cost of the flood, close to meaningless. It values destroyed buildings and equipment. It does not value the care that households subsequently paid for themselves, nor the care they went without. The distinction matters because Pakistan finances health care overwhelmingly at the point of use: private spending accounted for 64.4 per cent of total health expenditure in 2015–16, of which some 89 per cent was paid directly by households (National Health Accounts, reported in Bashir and Kishwar 2021). In such a system, an event that simultaneously raises the need for care and destroys the supply of it transfers cost onto household budgets almost mechanically. Roughly 13 per cent of the country's health facilities were damaged or destroyed and about 15 per cent of the population were left without access to primary or secondary care; incident malaria rose approximately fivefold, to some 2.6 million cases, and more than 41,000 dengue cases were recorded, concentrated in Sindh and Balochistan (Nasir et al. 2024; World Health Organization 2022). The households facing that disease burden were, for the most part, paying for treatment out of consumption.

This is the loss that the ledger omits, and it is not small. At a 10 per cent threshold of household consumption, the incidence of catastrophic health expenditure in Pakistan was estimated at 4.51 per cent of households in 2015–16 and 13.15 per cent in 2018–19 (Bashir and Kishwar 2021). Those are pre-flood figures. Estimates constructed on the capacity-to-pay definition differ in level but agree on the gradient: incidence in the poorest consumption quintile rose from 8.3 to 13.7 per cent between 2007–08 and 2018–19, against 5.1 to 8.4 per cent in the richest (Ahmed et al. 2024). A climate shock arriving on top of that distribution adds a health burden to the households least able to smooth it, in a system with almost no risk pooling to absorb it.

Why is the risk not pooled? This paper argues that two distinct frictions, each documented rigorously in recent literature but never jointly, govern the answer, and that in a setting like Pakistan's they compound.

The first is a **pricing friction**. Oh, Sen and Tenekedjieva (2026), using an instrumental-variables strategy built from state regulatory rules together with a border-discontinuity design, show that where insurance rate regulation binds, insurers do not adjust rates; they recover the shortfall in less regulated states; rates and risks diverge over the long run; and insurers exit the most tightly regulated markets. The authors trace the pattern to the interaction of rate regulation with insurers' financing constraints. The general lesson is not specific to American homeowners' insurance: where the price of a risk is held below its actuarial cost, or never established at all, the cover is not written.

The second is a **data friction**. Bonelli (2026) shows that venture capitalists adopting machine-learning screening become better at evaluating startups resembling those in historical data, tilt their portfolios towards that pool, and become less likely to finance the innovative ventures that achieve rare major success; plausibly exogenous variation in screening automation supports a causal reading. The health-sector analogue is documented directly: Obermeyer et al. (2019) show that a widely deployed clinical risk-prediction algorithm systematically understated the needs of Black patients because it was trained on health *costs* as a proxy for health *needs*, and less had historically been spent on those patients. In both cases the mechanism is the same. Estimator precision rises with the density of the record; a screening rule applied to a sparse record is conservative; and conservatism is indistinguishable, in its consequences, from exclusion.

The paper's central proposition is that in a thin-data health economy these two frictions are **sequential, not parallel**. A price that is never charged writes no claims record. A district with no claims record presents any subsequent algorithmic underwriter, whether of health microinsurance, digital credit or index-linked cover, with the sparse file its rules are built to decline. The population excluded by the missing price in one period is therefore the population excluded by the missing data in the next. Exclusion compounds. Equivalently, and more usefully for policy, the relation runs in reverse under intervention: cover written today generates the record that widens eligibility tomorrow, so the returns to establishing honest pricing include a data externality that conventional cost-benefit analysis of insurance subsidies omits.

The public-economics conclusion that follows is narrower than "expand insurance." Because the data externality is not internalised by any individual insurer, the private return to writing a first policy in a thin-data district falls short of the social return, which supplies a standard efficiency argument, separate from the usual equity argument, for public subsidy at the extensive margin, and for treating the resulting claims records as a public good rather than as proprietary assets.

Contribution. The paper makes three contributions. First, it reframes climate shocks in low-income settings as *health financing* shocks and shows why standard post-disaster accounting systematically understates them, since that accounting is denominated in destroyed assets and the dominant loss is a transfer onto household budgets. Second, it identifies and formalises the compounding of pricing and data frictions, and derives from it a data externality that changes the welfare arithmetic of insurance subsidy. Third, it specifies an identification strategy that is executable now with existing public microdata, rather than gesturing at data that does not exist: the Household Integrated Economic Survey rounds of 2018–19 and 2024–25 straddle the 2022 flood and both carry health expenditure modules, and Punjab's June 2025 termination of Sehat Sahulat cover in public hospitals provides a second, independent policy discontinuity permitting a triple-difference.

A statement of scope. This paper reports no estimation of its own. It is a conceptual and design contribution: it states propositions, derives testable hypotheses, and specifies the data and identification required to test them. Section 7 is explicit about what this does and does not license the reader to conclude.

Section 2 sets out the Pakistani health financing context and the 2022 shock. Section 3 places the argument in three literatures. Section 4 develops the framework and hypotheses. Section 5 specifies the empirical strategy. Section 6 presents the design response and its fiscal incidence. Section 7 states limitations, and Section 8 concludes.

2. Setting: A Health System Financed at the Point of Use

2.1 Health financing and financial protection

Pakistan's health system is financed principally by households at the moment of illness. Public expenditure on health has historically remained near or below one per cent of GDP, and the residual is met privately, predominantly out of pocket (Bashir and Kishwar 2021). The consequence, familiar from the financial-protection literature since Xu et al. (2003) and Wagstaff and van Doorslaer (2003), is that illness is transmitted directly into consumption: households facing a medical bill reduce food and schooling expenditure, sell productive assets, or borrow on informal terms.

Figure 1 shows the resulting structure. Roughly 57 per cent of total health expenditure is paid by households at the moment of treatment, with no pooling mechanism standing between the illness and the bill.

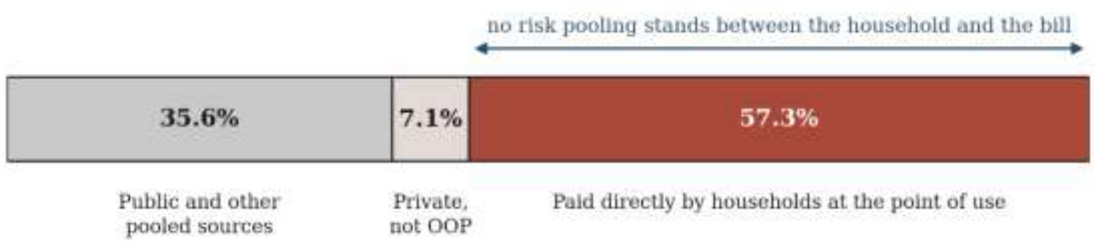

Figure 1. Composition of total health expenditure in Pakistan. The out-of-pocket share is derived from the National Health Accounts finding that private spending was 64.4 per cent of total health expenditure and that 89 per cent of private spending was out of pocket, as reported in Bashir and Kishwar (2021). The derived figure of 57.3 per cent is arithmetic on those two reported shares and is not itself a directly reported statistic.

Two measurement conventions organise the evidence. Under the *budget share* definition, a household incurs catastrophic health expenditure when out-of-pocket medical spending exceeds a fixed share of total consumption, conventionally 10 or 25 per cent. Under the *capacity-to-pay* definition, the denominator is consumption net of a subsistence allowance, with a threshold conventionally set at 40 per cent. The two produce different levels, which is why the Pakistani estimates cited in Section 1 diverge, and it is why any credible empirical work in this setting must report both. Sustainable Development Goal indicator 3.8.2 adopts the budget-share definition at the 10 and 25 per cent thresholds, which makes it the natural primary outcome for internationally comparable work.

Figure 2 juxtaposes the two accountings of the 2022 flood. Panel (a) shows assessed damage by sector, with the health sector's USD 109 million barely visible against housing and agriculture. Panel (b) decomposes that health-sector figure into facility damage, service losses and recovery needs. Neither panel contains a single rupee of household out-of-pocket spending, because the assessment instrument was not designed to capture it. That omission is the paper's point of departure.

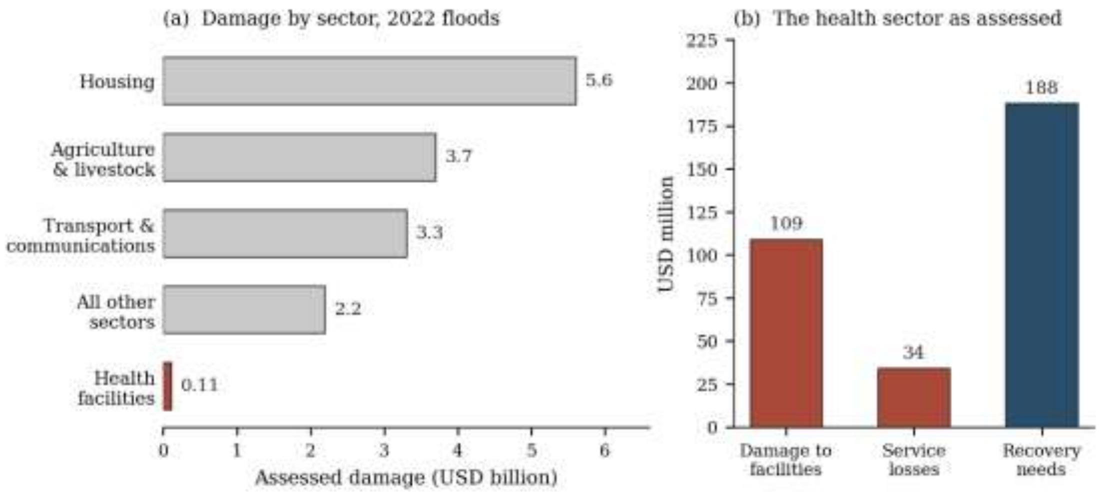

Figure 2. Two accountings of the 2022 floods. Panel (a): assessed physical damage by sector, USD billion, against an assessed total of USD 14.9 billion; the residual category aggregates all remaining sectors. Panel (b): the health sector as assessed, USD million. Neither measure includes household out-of-pocket medical expenditure. Source: Post-Disaster Needs Assessment (Government of Pakistan et al. 2022), health-sector components as reported in World Health Organization (2022).

Figure 3 shows what the pre-flood distribution of financial hardship looked like. Panel (a) reports national incidence of catastrophic health expenditure at the 10 per cent threshold across survey rounds, together with the divergent estimate produced under an alternative definition. Panel (b) shows the quintile gradient: incidence rose in both tails between 2007–08 and 2018–19, but remained roughly 60 per cent higher in the poorest quintile than in the richest throughout.

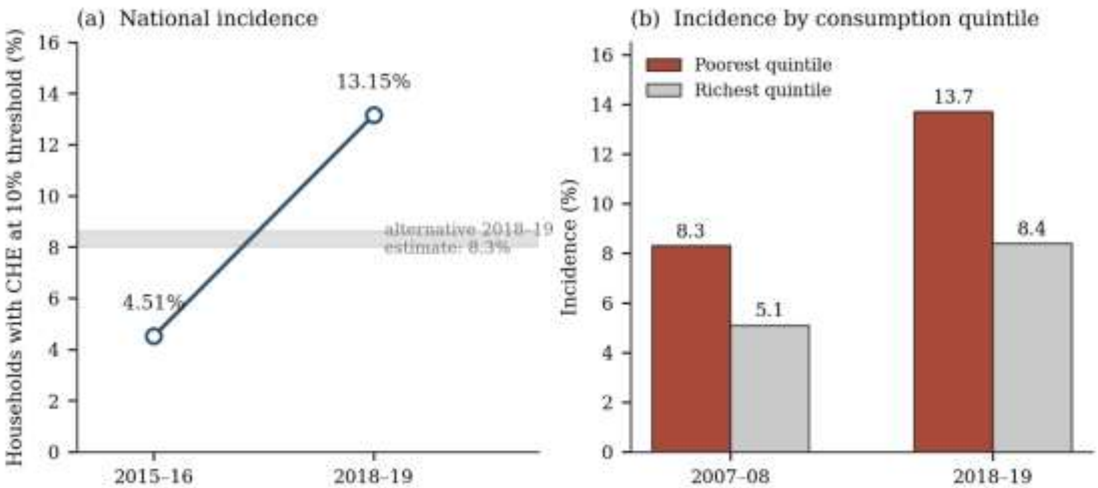

Figure 3. Catastrophic health expenditure in Pakistan before the 2022 floods. Panel (a): national incidence at the 10 per cent threshold of total household consumption. Panel (b): incidence by consumption quintile. Estimates constructed from Household Integrated Economic Survey rounds; levels differ across studies according to whether the budget-share or capacity-to-pay definition is used, which is why both are shown. Sources: Bashir and Kishwar (2021); Ahmed et al. (2024).

2.2 The 2022 floods as a health financing shock

The monsoon of 2022 delivered approximately 190 per cent of usual rainfall nationally in July and August, and several multiples of the norm in parts of Sindh and Balochistan. Roughly a third of the country was inundated; more than 33 million people across 116 districts were affected, over 1,700 died, and close to eight million were displaced (Congressional Research Service 2022; Government of Pakistan et al. 2022). Sindh bore close to 70 per cent of assessed damage, national output fell by an estimated 2.2 per cent of GDP, and an additional nine million people were pushed towards poverty (World Bank 2022).

The health consequences ran on two channels simultaneously, and it is their coincidence that matters analytically.

On the *supply* side, the World Health Organization reported 888 damaged health facilities by late August 2022, of which 180 were completely destroyed; subsequent situation reporting placed total facility damage above 2,000, with the Post-Disaster Needs Assessment finding that over 13 per cent of health facilities were damaged and more than a fifth of those damaged were destroyed outright (World Health Organization 2022). Roughly 15 per cent of the population lost access to primary and secondary care (Nasir et al. 2024). Disease surveillance itself was compromised: the Integrated Disease Surveillance and Reporting System was operational in only 45 of the country's districts at the time and was largely absent from the calamity-hit districts (World Health Organization 2022).

On the *demand* side, need rose sharply. Incident malaria increased approximately fivefold to some 2.6 million cases; more than 41,000 dengue cases were recorded; measles and rubella outbreaks were confirmed across 31 districts; and acute watery diarrhoea, typhoid and skin infections rose steeply in displacement settings (Nasir et al. 2024; World Health Organization 2022). Severe acute malnutrition with complications rose in areas where prevalence already exceeded emergency thresholds.

The pattern has since recurred. Renewed flooding in 2025 had, by 13 September of that year, killed 964 people, injured 1,181 and displaced more than 4.3 million, with the collapse of sanitation infrastructure again driving increases in diarrhoeal illness, malaria, acute fever and skin infections across Sindh, southern Punjab and Khyber Pakhtunkhwa (Fahim, Qureshi and Sethi 2025; Jan, Jan and Jan 2025). The policy question is therefore not how to finance recovery from a singular event but how to finance a recurring one, which is the argument for pre-arrangement developed in Section 6.
Figure 4 places the two channels side by side.

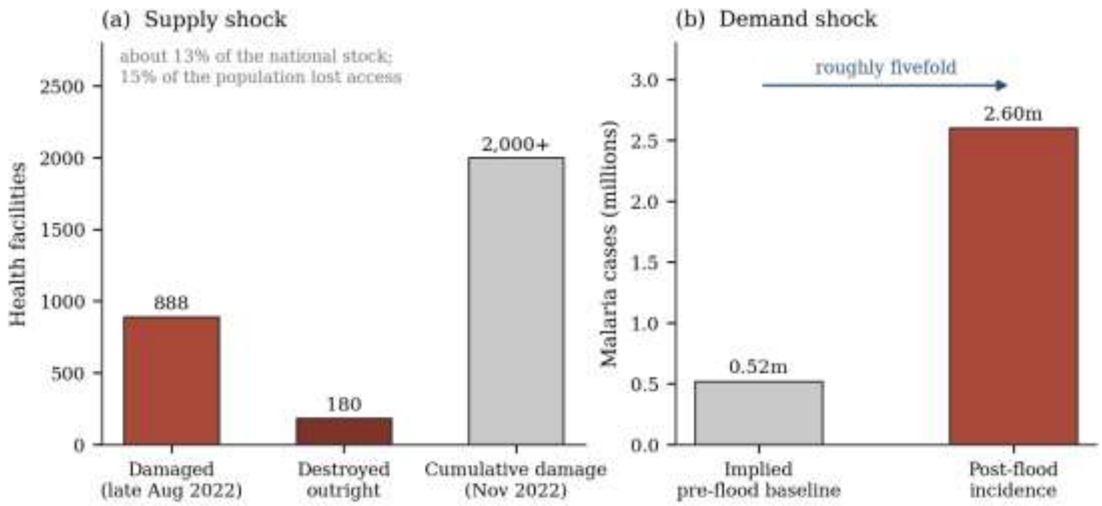

Figure 4. The 2022 floods as a two-sided shock. Panel (a): health facilities damaged and destroyed, with the November 2022 cumulative figure drawn from later situation reporting. Panel (b): malaria incidence. The post-flood figure of 2.6 million is reported; the pre-flood baseline is implied by the reported fivefold increase rather than separately observed, and is shown on that basis. Not shown: more than 41,000 dengue cases, and measles and rubella outbreaks confirmed across 31 districts. Sources: World Health Organization (2022); Nasir et al. (2024).

A shock that raises need while destroying supply, in a system financed at the point of use, produces a predictable set of household responses: forgone care, treatment sought from more expensive or lower-quality private providers as public facilities close, distress financing, and asset depletion. Each is a welfare loss. None appears in a damage assessment denominated in destroyed buildings. This is the measurement gap the paper’s empirical strategy is designed to close.

2.3 Existing institutions: what a design response can build on

Three existing institutions determine what is feasible, and a design that ignores them will not be implemented.

Social health insurance. The Sehat Sahulat Programme, launched in Khyber Pakhtunkhwa in 2015 and expanded nationally from 2019, provides cash-free inpatient cover, typically around PKR 460,000 per family per year, administered through third-party administrators, principally the State Life Insurance Corporation, across empanelled public and private hospitals (Commonwealth Fund 2026; Khan et al. 2022). The evidence base is mixed but broadly positive on financial protection, with benefits concentrated among poorer households (Ahmad et al. 2025).

Three features constrain its usefulness against climate shocks. First, it is predominantly an inpatient benefit, whereas the flood-associated disease burden was overwhelmingly outpatient. Second, it is fiscally fragile: by March 2023 the State Life Insurance Corporation had warned the Punjab government of possible closure over unpaid premium invoices reported at approximately PKR 83.5 billion, against a provincial allocation of PKR 125 billion for FY2022–23 of which only about PKR 31 billion had been released (Dawn 2023).

Third, and consequently, coverage is not stable. The precise nature of the 2025 change in Punjab matters for what follows and is frequently misdescribed, so it is worth stating exactly. A notification issued by the Punjab Health Initiative Management Committee directed all government hospitals to end Universal Health Insurance services under the programme by 30 June 2025, after which the State Life Insurance Corporation ceased processing reimbursements for public-sector claims; the stated rationale was to avoid duplication with the Chief Minister's selective initiatives, such as the children's heart surgery and dialysis and transplant programmes. Cover was **retained at empanelled private hospitals** within the province. Separately, from 1 July 2025, Punjab cardholders lost portability: empanelled hospitals in Islamabad, Azad Jammu and Kashmir, Sindh, Khyber Pakhtunkhwa, Balochistan and Gilgit-Baltistan were instructed to stop treating them, with only patients already under treatment permitted to continue. Provincial officials publicly denied that the programme had been discontinued, maintaining that free treatment continued (Dawn 2025; ARY News 2025).

This is therefore a *partial and channel-specific* withdrawal, not a termination. That distinction is what makes it useful as an identifying event, and it is also what makes its interpretation delicate, since the replacement initiatives substitute for cover in some conditions and not others.

Targeting infrastructure. The Benazir Income Support Programme maintains a national socioeconomic registry with established proxy-means targeting and payment mechanisms. Any premium subsidy or contingent transfer can be routed through this architecture rather than through a new one, which materially reduces both administrative cost and the political difficulty of establishing eligibility.

Payment rails. Raast, the State Bank's instant-payment system, had processed 1.9 billion transactions worth PKR 44.3 trillion by June 2025, with more than 2.6 million merchants accepting instant payment; digital channels carried 92 per cent of retail payment volumes by the third quarter of FY26, and registered mobile-banking and wallet users rose from 96 million to 132 million in the year to March 2026 (State Bank of Pakistan 2025, 2026). The first digital retail bank was licensed in January 2025.

The rails permit a payout to reach a household within days rather than months. They also determine what runs on them. Digital lenders and insurtechs will screen by model, trained on whatever records the formal economy has generated. That returns the argument to the data friction, and to the reason it cannot be treated separately from the pricing question.

3. Related Literature

3.1 Financial protection, health shocks and consumption smoothing

The proposition that health risk is insurable and that private markets will under-provide the insurance is the founding result of health economics (Arrow 1963). The empirical literature on financial protection descends from Wagstaff and van Doorslaer (2003) and Xu et al. (2003), who established the measurement of catastrophic and impoverishing health payments, and from O'Donnell et al. (2008), who standardised its implementation in household survey data. Wagstaff et al. (2018) document that progress in reducing catastrophic spending has been slow and uneven across 133 countries.

On the consequences of unpooled health risk, Gertler and Gruber (2002) show that Indonesian households smooth consumption well against minor illness and poorly against major illness, precisely the pattern that makes catastrophic events the relevant object of insurance. Mohanan (2013) uses bus-accident injuries in India as a quasi-experiment and finds that households protect consumption largely by taking on debt rather than by reducing expenditure. The result matters here because debt-financed recovery is a deferred welfare loss, not an avoided one.

On whether insurance helps, the evidence is more qualified than advocacy usually allows. Finkelstein et al. (2012), exploiting the Oregon lottery, find substantial improvements in financial protection and self-reported health from Medicaid coverage. Banerjee et al. (2021), in a large-scale randomised evaluation of Indonesia's national scheme, find that enrolment in voluntary contributory health insurance is low without substantial intervention and that assistance with enrolment and subsidy both raise it, with implications for the design of any contributory scheme in a comparable setting. Sood et al. (2014) find reductions in mortality and out-of-pocket spending from India's RSBY-type coverage. The general lesson is that coverage on paper and financial protection in practice diverge, and that the divergence is a design problem.

3.2 Catastrophe risk, pricing and the protection gap

A parallel literature explains why markets for catastrophe risk remain thin. Froot (2001) documents that catastrophe reinsurance is supplied sparingly and priced well above actuarial value, and attributes this to capital market imperfections. Koijen and Yogo (2015) show that financing frictions internal to insurers bend pricing away from fundamentals. On the demand side, Cole et al. (2013) find that trust, liquidity and salience, rather than price alone, govern Indian household take-up of rainfall insurance, while Mobarak and Rosenzweig (2013) show that index products interact with informal risk-sharing in ways that may either crowd it out or complete it. Barnett, Barrett and Skees (2008) set out the case for index-based transfer against poverty traps, and Jensen and Barrett (2017) review the disappointing take-up record and the basis-risk problem that explains much of it. Clarke and Dercon (2016) draw the policy conclusion that pre-arranged, rule-based finance dominates post-disaster improvisation. Hallegatte et al. (2017) quantify how conventional loss accounting understates welfare losses by weighting the assets of the poor at market value; the argument of Section 2.1 is a health-financing instance of that general critique.

To this, Oh, Sen and Tenekedjieva (2026) add the supply-side channel with unusually clean identification: rate suppression produces cross-subsidy, divergence of price from peril, and eventual exit.

3.3 Algorithmic screening and the allocation of care and capital

A third literature asks what changes when screening is delegated to a model. Goldstein, Jiang and Karolyi (2019) frame the research agenda. Berg et al. (2020) show that digital footprint data predict default about as well as bureau scores, a genuine promise for thin-file borrowers. Fuster et al. (2022) show that the gains from machine learning in credit are distributed unequally across borrower groups. Bonelli (2026) supplies the allocation result for venture capital described above.

Within health specifically, Obermeyer et al. (2019) provide the canonical demonstration that a model trained on an imperfect proxy for need reproduces and amplifies the historical pattern of under-provision, and that the defect is invisible without an audit designed to detect it. Mullainathan and Obermeyer (2022) show that the same machinery, differently deployed, can identify systematic error in human clinical judgement. The pair of results is important for the policy design in Section 6: the appropriate response to algorithmic exclusion is neither prohibition nor unaudited adoption, but mandatory auditing against an outcome that is not itself contaminated by historical spending.

3.4 The gap

Financial-protection scholarship measures the consequences of unpooled health risk but is largely silent on the supply-side reasons that pooling does not emerge. Catastrophe-insurance scholarship analyses pricing and regulation, almost entirely in high-income property markets, and rarely treats health as the insured loss. Algorithmic-finance scholarship analyses selection and fairness but treats the training data as given rather than as itself the product of prior pricing decisions. To our knowledge no study has joined the suppression of risk prices to the data poverty that subsequently biases algorithmic underwriting, or derived the joint design implications for health financing in a climate-exposed low-income economy. That junction is the gap this paper addresses.

Table 1 distils the two anchor studies and states the transposition each requires before it can bear weight in this setting.

Table 1. The anchor results and their transposition to health financing in Pakistan.

Dimension	Oh, Sen and Tenekedjieva (2026)	Bonelli (2026)	Obermeyer et al. (2019)
Setting	US homeowners' insurance across state regulatory regimes.	US venture capital firms adopting data technologies for screening.	A commercial risk-prediction algorithm applied to a large US patient population.
Method	Instrumental variables from state regulatory rules; border-discontinuity design.	Panel comparison around adoption with plausibly exogenous variation in screening automation.	Comparison of algorithmic risk scores against realised health need at equal predicted risk.
Core finding	Rates are not adjusted where regulation binds and are recovered elsewhere; price and risk diverge; insurers exit tightly regulated markets.	Adopters improve within the pool resembling historical data, tilt towards it, and finance fewer rare major successes.	Training on health spending as a proxy for health need reproduces historical under-provision as predicted low risk.
Mechanism invoked here	Binding price constraint means cover is not written (F1).	Estimator precision rises with data density, so thresholds favour the well recorded (F2).	The training label itself encodes prior unequal access (F2, health-specific).
Transposition required	Federal rather than state regulation; the binding case is usually the absence of a price rather than a ceiling on one.	Underwriter objective differs from a VC's; the relevant outlier is an underserved household, not a high-growth firm.	Closest in domain; the constraint is that Pakistan's spending data are sparser still, so the bias may be larger and harder to detect.

4. Framework: Two Frictions and Their Compounding

4.1 The pricing friction

Consider the supply of cover against a health-relevant climate risk in district *d*. An insurer faces expected loss $E[L_d]$, a cost of risk capital κ and an administrative loading c. Cover is written only if the permitted premium covers the load:

$$P_d \geq E[L_d]\,(1+\kappa) + c$$

Equation (1)

Where a ceiling $\bar{P}_d < E\underline{L}_d + c$ binds, the insurer has two options and no third: cross-subsidise the shortfall from unregulated segments, or withdraw. This is the **pricing friction (F1)**. The subsidy implicit in a binding ceiling does not disappear; it relocates. It moves first into the premiums of other policyholders, then into insurer balance sheets, and finally into the withdrawal of supply that Oh, Sen and Tenekedjieva (2026) document.

Two features distinguish the Pakistani case from the American one and must be carried into any transplant. First, regulation is federal, through the Securities and Exchange Commission of Pakistan, so there is no analogue to cross-state variation in rate filings; the relevant variation lies in provincial scheme rules, premium subsidies and benefit ceilings. Second, and more fundamentally, the binding constraint in most districts is not a ceiling on an existing price but the *absence of any established price at all*. Formally, this is the corner case $\bar{P}_d = 0$: no product exists. The economics is the same, since cover goes unwritten either way, but the policy implication differs, since the remedy is establishing a price rather than relaxing one.

4.2 The data friction

Consider an underwriter estimating the quality or risk ϑ_i of applicant *i* from n_i historical observations, funding or covering only applicants whose lower confidence bound clears a threshold ϑ^*:

$$\mathrm{Var}(\hat{\theta}_i) = \sigma^2/n_i, \qquad \text{accept } i \iff \hat{\theta}_i - z\sigma/\sqrt{n_i} \geq \theta^*$$

Equation (2)

Because estimator precision rises with the density of the record, two applicants of identical true quality are not treated identically: the one with the longer record clears the threshold more easily. Automating the rule therefore tilts allocation towards the well-documented. This is the **familiarity or data friction (F2)**, the mechanism Bonelli (2026) identifies in venture capital, Fuster et al. (2022) in consumer credit, and Obermeyer et al. (2019) in clinical risk prediction.

The health-sector version has an additional and more troubling feature. In credit, the training label (default) is at least a direct measure of the outcome of interest. In health, as Obermeyer et al. (2019) demonstrate, the available label is frequently *spending*, which is a measure of historical access rather than of need. Where access has been unequal, a model trained on spending will encode that inequality as though it were epidemiology. In Pakistan, where the flood-affected districts had the weakest surveillance coverage before the flood and lost what they had during it, this is not a hypothetical concern.

4.3 Why the two frictions compound

The paper's central claim is that F1 and F2 are sequential: the first manufactures the conditions for the second.

Prices suppressed or never established mean policies not written. Policies not written mean claims histories never generated. Histories never generated mean that n_i in equation (2) is smallest exactly where expected loss $E[L_d]$ in equation (1) is largest. The algorithms deployed several years later are then trained, necessarily, on the urban, salaried, already-documented segment of the economy, and they learn, correctly given their evidence, to decline everyone else. A flood-exposed district is excluded twice: once by the price that was never set, and once by the record that was, in consequence, never written. Figure 5 sets out the schematic.

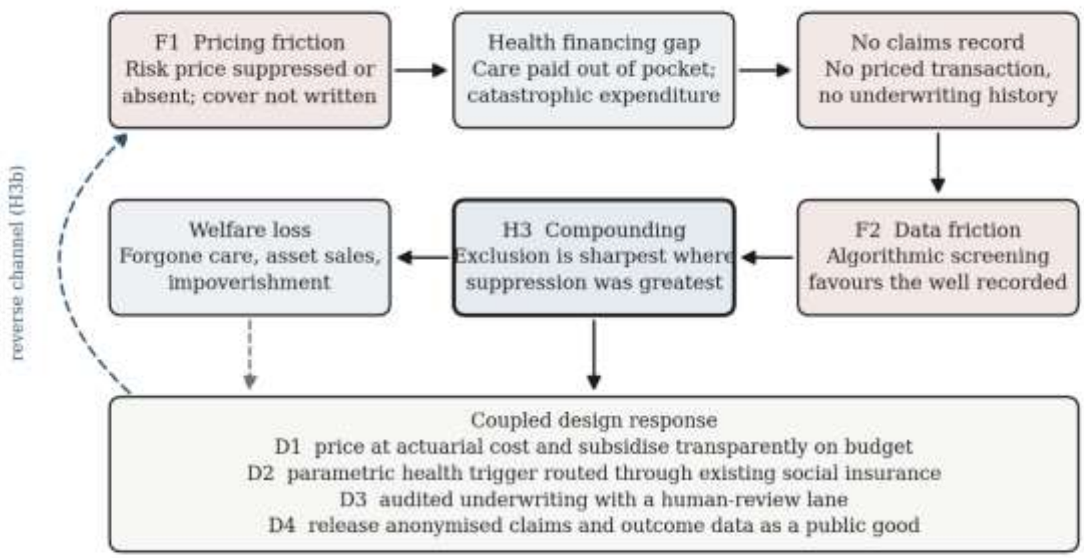

Figure 5. Framework schematic. The pricing friction (F1) and the data friction (F2) compound through the health financing gap (H3), and are addressed by a coupled design response (D1–D4). The dashed reverse channel (H3b) indicates that intervention on F1 relaxes F2 with a lag. Conceptual diagram: the boxes assert logic, not estimates.

This has a corollary that is both testable and optimistic. If suppression starves a district of records, then cover generates them. A district provided with honestly priced, well-recorded parametric protection today should, several years hence, present underwriters with a denser file and a narrower gap between model-based and human decisions. The compounding runs in reverse under intervention, and the lagged relaxation is an estimable quantity.

It also has a welfare implication that does not depend on distributional preferences. The claims record generated by writing a first policy in a thin-data district is an input to every subsequent underwriter's estimation problem, not only to the writing insurer's. Its social value therefore exceeds its private value, and the wedge is a positive externality of exactly the kind that justifies public intervention on efficiency grounds. Standard appraisals of premium subsidy weigh the transfer against the deadweight cost of the revenue raised; they omit this term. Where it is large, which is where n_i is near zero and the marginal informational return to an additional record is correspondingly high, the efficient subsidy at the extensive margin is larger than conventional analysis implies, and the case for public ownership of the resulting data is correspondingly stronger.

4.4 Hypotheses

H1 (pricing friction). Districts and product lines in which premiums are most suppressed relative to modelled risk, or in which no priced product exists, exhibit thinner cover, larger implicit cross-subsidy and faster insurer withdrawal following a major flood, holding physical exposure constant.

H2 (data friction). As underwriting and credit screening automate, the approved pool shifts towards applicants with long digital or claims histories and away from first-time, rural, female and thin-file applicants; and, conditional on approval, the incidence of good outcomes among marginal applicants is not lower for the excluded group. Exclusion following that pattern reflects estimator conservatism rather than genuine risk.

H3 (compounding). The exclusion measured under H2 is sharpest in precisely the districts identified under H1, because price suppression starved those districts of the records an algorithmic screen rewards.

H3b (reverse channel). Districts receiving honestly priced, well-recorded cover exhibit, with a lag, denser applicant files and a narrower model-versus-officer approval gap.

H4 (health financing incidence). Flood exposure raises the incidence of catastrophic and impoverishing health expenditure, and the increase is attenuated where households held social health insurance coverage at the time of the shock.

H4 is the hypothesis that makes the framework a health-economics contribution rather than an insurance-market one, and it is the hypothesis for which the data required are already in the public domain.

5. Empirical Strategy and Data

This section specifies what can be estimated. The designs are ordered by data availability: Design I is executable immediately with released public microdata; Design II becomes executable as the 2026–27 district round is released; Design III requires regulatory cooperation and ethical approval. Figure 6 places the designs against the survey and policy timeline, and Table 3 maps each design to its identification strategy and sources.

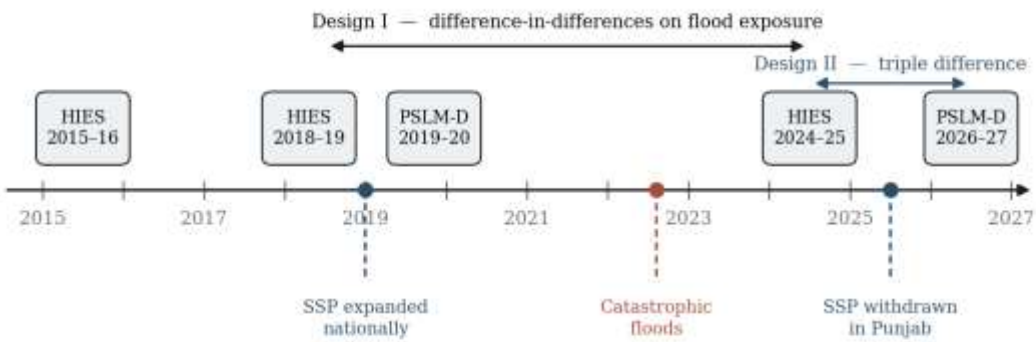

Figure 6. Survey rounds, policy events and identification windows. HIES denotes the provincially representative Household Integrated Economic Survey; PSLM-D the district-representative Pakistan Social and Living Standards Measurement survey. Design I compares flood-exposed and unexposed districts across the 2018–19 and 2024–25 rounds; Design II adds variation from Punjab's June 2025 termination of Sehat Sahulat cover in public hospitals.

5.1 Data

Household microdata. The Household Integrated Economic Survey is the principal source. The 2018–19 round is the last full pre-flood round; the 2024–25 round, the first fully digital HIES, was completed and released in January 2026 (Pakistan Bureau of Statistics 2026). Both contain the consumption and health expenditure modules required to construct catastrophic and impoverishing payment measures on both conventions. The district-representative PSLM rounds of 2019–20 and 2026–27 provide the finer geographic resolution needed for exposure matching, the latter currently in field.

Exposure. District flood exposure is constructed from satellite-derived inundation extent, using Sentinel-1 synthetic aperture radar composites and UNOSAT flood analyses for the July to October 2022 window, aggregated to district boundaries as the share of populated area inundated, and cross-checked against National Disaster Management Authority calamity declarations. Continuous exposure is preferred to a binary treatment indicator because the calamity declaration is itself an administrative decision plausibly correlated with unobservables.

Health system and coverage. Facility damage and functionality from the health-sector components of the Post-Disaster Needs Assessment and WHO situation reporting; Sehat Sahulat enrolment and utilisation from programme administrative data and the third-party administrator; disease incidence from IDSR and provincial surveillance, with the caveat about differential surveillance coverage noted in Section 5.5.

5.2 Outcome measurement

The primary outcomes are the standard financial-protection measures, reported on both conventions so that the results are not an artefact of definition:

- **Catastrophic health expenditure**, budget-share definition, at 10 and 25 per cent of total household consumption (SDG indicator 3.8.2).

- **Catastrophic health expenditure**, capacity-to-pay definition, at 40 per cent of consumption net of a subsistence allowance defined on the food share, following O'Donnell et al. (2008).
- **Impoverishing health expenditure**: the share of households pushed below the national poverty line by out-of-pocket payments, and the associated poverty gap.
- **Out-of-pocket share** of household consumption, as a continuous outcome robust to threshold choice.
- **Forgone care**: reported illness episodes without treatment, which is essential because unmet need suppresses measured expenditure. A household too poor to seek care records *low* out-of-pocket spending, so catastrophic expenditure measures understate hardship exactly where hardship is greatest. Any estimate that omits this margin will understate the effect of the shock.

Table 2. Outcome definitions, thresholds and rationale.

Outcome	Construction	Threshold	Why included
CHE, budget share	Out-of-pocket health payments for all household members, over total household consumption expenditure	10% and 25%	SDG indicator 3.8.2. Primary outcome, chosen for international comparability
CHE, capacity to pay	Same numerator, over consumption less a subsistence allowance defined on the food share	40%	Follows O'Donnell et al. (2008). Reported alongside the budget-share measure because published Pakistani estimates diverge in level according to which is used
Impoverishing health expenditure	Households falling below the national poverty line once out-of-pocket payments are netted out of consumption	National poverty line, adjusted to survey-round prices	Captures the margin at which a health payment becomes a poverty event. Report headcount and poverty gap
Out-of-pocket share	Out-of-pocket payments over total consumption, left continuous	none	Threshold-free. Guards against a result that turns on where the threshold was placed
Forgone care	Reported illness episodes for which no treatment was sought, over all reported episodes	none	Unmet need suppresses measured spending, so omitting this margin biases every measure above towards zero

5.3 Design I: flood exposure and financial protection

The first design tests H4. Let y_{hdt} denote an outcome for household *h* in district *d* at survey round $t \in$ {2018–19, 2024–25}, and Flood_d the continuous exposure measure:

$$y_{hdt} = \alpha_d + \lambda_{pt} + \beta\,(\text{Flood}_d \times \text{Post}_t) + \mathbf{X}'_{hdt}\gamma + \varepsilon_{hdt}$$

Equation (3)

with district fixed effects α_d, province-by-round effects λ_{pt}, and household covariates **X**. Under H4, $\beta > 0$ for catastrophic and impoverishing expenditure. Standard errors are clustered at the district level, with a wild cluster bootstrap given the modest number of clusters.

The coefficient of interest for the paper's central argument is the interaction with pre-shock coverage:

$$y_{hdt} = \alpha_d + \lambda_{pt} + \beta_1(\text{Flood}_d \times \text{Post}_t) + \beta_2(\text{Flood}_d \times \text{Post}_t \times \text{Covered}_{hd}) + \mathbf{X}'_{hdt}\gamma + \varepsilon_{hdt}$$

Equation (4)

where $\beta_2 < 0$ would indicate that social health insurance attenuated the financial impact of the shock.

That is the quantity most directly relevant to the design response in Section 6.

5.4 Design II: the Sehat Sahulat withdrawal as a policy discontinuity

Punjab's June 2025 termination of Sehat Sahulat cover in public hospitals removed a major financing channel from a large, previously covered population at a date unrelated to flood exposure. This permits a triple-difference in which the third dimension is exposure to that termination:

$$y_{hdt} = \alpha_d + \lambda_{pt} + \delta\,(\mathrm{Flood}_d \times \mathrm{Post}_t \times \mathrm{Withdrawn}_p) + \text{lower order interactions} + \mathbf{X}'_{hdt}\gamma + \varepsilon_{hdt}$$

Equation (5)

Three features of the policy change shape the design, and each cuts both ways.

The treatment is channel-specific, which is an advantage. Because cover was terminated in public hospitals but retained at empanelled private hospitals, the shock does not fall uniformly on all cardholders. Households whose care was concentrated in public facilities, disproportionately poorer and rural ones, lost more than households already using the private network. Pre-shock reliance on public facilities therefore provides a continuous measure of exposure intensity within Punjab, which is a considerably stronger design than a binary provincial indicator and permits within-province comparison that absorbs provincial policy shocks common to both groups.

The loss of portability is a second, cleaner margin. From 1 July 2025 Punjab cardholders could no longer be treated in other provinces. This affects a narrower population, namely cross-provincial care-seekers concentrated in tertiary referrals, and is unlikely to be confounded by the substituting provincial initiatives, which operate only within Punjab. It therefore serves as a useful robustness margin.

The replacement programmes confound the sign. The Chief Minister's initiatives substitute for the terminated cover in specific high-cost conditions (paediatric cardiac surgery, dialysis, transplant) and not in others. The net effect on financial protection is therefore ambiguous *a priori* and heterogeneous by condition. Analysis must be stratified by whether the household's care falls inside or outside the substituted conditions; pooling them would average a substitution effect against a withdrawal effect and could yield a null that means nothing.

The identifying assumption is that, absent the termination, financial-protection outcomes in Punjab would have evolved in parallel with those in provinces retaining public-hospital cover, conditional on flood exposure. This is contestable. Two things discipline it. First, the change appears to have been driven by provincial fiscal arrears and programmatic substitution rather than by flood recovery, which supports exogeneity with respect to flood exposure while raising a separate concern that provincial fiscal stress is itself correlated with capacity to fund recovery. Second, the pre-period permits a placebo: exposure to the termination should not predict differential trends in financial protection before June 2025. If it does, the design fails and should be reported as failing.

5.5 Design III: auditing algorithmic underwriting

The third design tests H2 and H3 directly, and requires regulatory cooperation. Within the State Bank's regulatory sandbox, matched thin-file applications for health-linked microinsurance or credit are routed in parallel to an algorithmic screen and to experienced human officers. Both decisions are recorded; only one is implemented; and outcomes are tracked. The estimand is the model-versus-officer approval gap and, critically, its geography: H3 predicts that the gap is largest in the districts identified under H1.

The audit must be constructed to avoid the failure mode Obermeyer et al. (2019) identify. Evaluating a model against historical approvals or historical spending tests only whether it reproduces past decisions. The audit must therefore track a *realised outcome*, whether repayment or health service utilisation and health status, among applicants the model would have declined but the officer approved. Without that arm, the exercise cannot distinguish estimator conservatism from correct risk assessment, which is precisely the distinction H2 asserts.

5.6 Threats to identification

Four threats deserve statement in advance rather than in response to referees.

Differential measurement. Surveillance coverage was weakest in the districts most affected, so administrative disease incidence is measured with error correlated with treatment. This is why household survey outcomes, not administrative health data, are the primary outcomes; the administrative series is used descriptively only.

Selective migration. Close to eight million people were displaced. If displacement is selective on the outcome, so that the households worst affected left the sampling frame, estimates from a repeated cross-section are biased, and the sign of the bias is towards zero. Bounding exercises using district population change between the 2017 and 2023 censuses, and comparison of household composition across rounds, are necessary rather than optional.

Compositional change in the survey. The 2024–25 round is the first fully digital HIES, conducted after the 2023 census. Instrument and frame changes can generate apparent effects. Comparison of stable, non-health outcomes across rounds provides a check.

Anticipation and the parallel-trends assumption. With two survey rounds, pre-trends cannot be inspected in the estimating sample. The 2015–16 and 2018–19 rounds permit a pre-period placebo on the same specification, and this should be reported whether or not it is favourable.

Table 3. Empirical agenda: designs, identification and data.

Design	Hypothesis and question	Identification	Key data sources	Status
D-I. Household panel of repeated cross-sections	H4: does flood exposure raise catastrophic and impoverishing health expenditure, and does prior coverage attenuate it?	Two-way fixed effects with province-by-round shocks; continuous satellite exposure; placebo on the 2015–16 to 2018–19 pre-period; migration bounding.	HIES 2018–19 and 2024–25; PSLM-D 2019–20; Sentinel-1 and UNOSAT inundation; NDMA records.	Executable now
D-II. Triple difference on coverage termination	H1, H4: does removal of social health insurance amplify the financial impact of climate shocks?	Flood exposure interacted with Punjab's June 2025 termination of public-hospital cover; exposure intensity measured by pre-shock reliance on public facilities; portability loss as a second margin; stratified by substituted conditions; pre-June-2025 placebo.	HIES 2024–25; PSLM-D 2026–27; SSP administrative enrolment and utilisation; PHIMC notifications.	On release of 2026–27 round

D-III. Prospective underwriting audit	H2, H3: how large is the model-versus-officer gap on matched thin-file applications, and is it largest in H1 districts?	Randomised parallel routing within the SBP regulatory sandbox; realised-outcome tracking on the officer-approved, model-declined arm.	Participating digital banks and insurtechs; BISP registry covariates; Raast metadata under consent.	Requires approval

6. Design Response and Fiscal Incidence

The framework yields a package of four coupled components. They are coupled by construction: each addresses one friction and generates the input the next requires.

D1. Price at actuarial cost; subsidise transparently and on budget. The remedy for F1 is not a lower price but a visible subsidy. Premiums should be set to satisfy equation (1), with affordability addressed through explicit premium support paid on budget and targeted through the existing BISP registry. This converts an implicit, regressive and undocumented cross-subsidy into an explicit, targetable and auditable transfer. The political economy is harder than the arithmetic: a subsidy that can be seen acquires opponents who can see it, whereas a cross-subsidy buried in premiums has none. That difficulty is a reason to design for it, not a reason to prefer opacity.

D2. Parametric cover for health-relevant flood risk, routed through existing social insurance. Triggers defined on satellite rainfall and river-gauge thresholds permit payment without loss adjustment, which is what makes disbursement in days rather than months feasible; payment travels over Raast. Two design choices distinguish this from a generic parametric proposal. First, the benefit should be structured to cover the *outpatient* and transport costs that dominated the 2022 disease burden, since the existing Sehat Sahulat benefit is principally inpatient and would not have covered the malaria and diarrhoeal episodes that constituted most of the shock. Second, the product should be administered as a rider on the existing scheme rather than as a parallel programme, which economises on enrolment infrastructure and writes the claims record into an existing dataset rather than a new silo.

Basis risk is the central technical objection to any parametric design and should not be minimised (Jensen and Barrett 2017): a household may suffer loss without the trigger firing. Two mitigations are available. The trigger can be calibrated against historical district-level health utilisation rather than rainfall alone, and a residual discretionary layer can be retained for verified non-triggered losses. Neither eliminates basis risk; both bound it.

Sovereign and regional parametric facilities, among them the African Risk Capacity, the Caribbean Catastrophe Risk Insurance Facility and the Pacific catastrophe risk pool, have operated against rainfall, drought and cyclone triggers for over a decade, so the design question is adaptation rather than invention. Notably, nothing in a parametric design requires a claims history, which is why it can begin in districts where F1 has left no record: it relaxes the pricing friction and simultaneously begins generating the data that will later relax the data friction.

D3. Algorithmic accountability as a licence condition. Digital banks and insurtechs would publish inclusion audits alongside default and loss ratios, reporting approval rates by first-time status, district, gender and file thickness, and would maintain a human-review lane for applicants the model has not seen. That requirement follows directly from equation (2). Following Obermeyer et al. (2019), audits must be conducted against realised outcomes rather than against historical approvals or spending, since a model trained on spending in an unequally served population will encode that inequality as need.

D4. Claims and outcome data as a public good. Anonymised aggregates of flood outcomes, health utilisation, and payment records should be released for model training under a defined consent and disclosure-control architecture. This is the component that internalises the externality identified in Section 4.3, and it is the one most likely to be resisted, since the data are commercially valuable to whoever holds them. That is precisely the argument for treating them as public infrastructure.

6.1 Fiscal incidence

A design response that does not state who pays is not a policy proposal. Three observations about incidence follow from the framework, though quantification requires the estimates Section 5 specifies.

First, the counterfactual is not zero. Pakistan already acts as insurer of last resort, financing post-disaster response through borrowing, budget reallocation and appeals. The relevant comparison is not "subsidy versus no subsidy" but "pre-arranged and priced versus improvised and unpriced," and the pre-arranged option is cheaper for a given level of protection because it is financed before the event, when the sovereign is not a distressed borrower (Clarke and Dercon 2016).

Second, the incidence of the current arrangement is regressive and undocumented. Uninsured recovery is financed by household asset sales and informal borrowing, which fall on the households least able to bear them, and by contingent public liabilities that are not scored until they crystallise. Making the transfer explicit does not create a cost; it reveals and reallocates one.

Third, the marginal fiscal cost of D2 is bounded by the parametric structure. A trigger-based payout has a known maximum exposure per event, which is what makes it reinsurable through regional pools and the Loss and Damage Fund. That complement is essential, since a single correlated catastrophe would otherwise exhaust the domestic market the scheme exists to develop.

Table 4. Mapping the design to the Sustainable Development Goals.

Goal	Channel in the framework	Illustrative indicator
SDG 3. Good health and well-being	Financial protection against health shocks is the core object of the design; parametric cover addresses the outpatient burden that dominates post-flood morbidity.	SDG 3.8.2: population with household expenditure on health exceeding 10% and 25% of total consumption.
SDG 1. No poverty	Payouts within days interrupt the asset-sale and distress-borrowing channels of uninsured recovery; the nine million pushed towards poverty in 2022 are the counterfactual.	Impoverishing health expenditure headcount; share of affected households receiving payout within 14 days of trigger.
SDG 5. Gender equality	Payouts routed to accounts held by women; inclusion audits report the female share of approvals; maternal and child care dominate the outpatient burden.	Female share of policyholders and of algorithm-approved first-time applicants.
SDG 8. Decent work and growth	Protection against health-driven asset depletion preserves the productive capacity of smallholders and informal enterprises through shocks.	Thin-file share of new lending; enterprise survival, insured versus uninsured, after flood exposure.
SDG 10. Reduced inequalities	The data-externality argument implies that unaudited automation widens existing gaps; D3 and D4 are the corrective.	Approval-rate gap between thin-file and thick-file applicants of equal realised risk.
SDG 13. Climate action	Pre-arranged health-linked adaptation finance rather than post-hoc appeal.	Share of modelled flood-related health loss covered by pre-arranged instruments.

SDG 17. Partnerships	Regional reinsurance pools and the Loss and Damage Fund carry the correlated tail no domestic market can hold.	Reinsured share of national catastrophe exposure; disbursements to pre-arranged schemes.

7. Discussion and Limitations

Six limitations deserve statement, and the first is the most consequential.

The paper reports no estimates. Its propositions are disciplined by rigorous external evidence, but for Pakistan they remain hypotheses. Nothing here establishes that flood exposure raised catastrophic health expenditure, nor that algorithmic underwriting excludes flood-exposed households; those are the claims Section 5 is designed to test. A reader should treat the framework as a structure for interpretation and the empirical section as a protocol, not as findings.

External validity of the anchor results. Oh, Sen and Tenekedjieva (2026) identify their effects in the American homeowners' insurance market, which has dense actuarial data, mature regulation and near-universal product availability. Pakistan has none of these. Section 4.1 argues that the corner case $\bar{P}_{\mathrm{d}} = 0$ preserves the mechanism, but that argument is a conjecture about generalisability, not an extension of their identification. Similarly, Bonelli (2026) identifies effects among venture capitalists, whose objective function differs materially from a health insurer's. The health-sector transposition rests more securely on Obermeyer et al. (2019), which is closer in domain but narrower in scope.

The compounding claim is the hardest to identify. H3 asserts a sequential relationship over a horizon of years. Establishing it requires either long panels that do not yet exist, or the sandbox audit of Design III, whose external validity is limited by the sample of participating institutions. It is entirely possible that F1 and F2 are correlated through a common cause, district poverty, rather than sequentially linked. Distinguishing the two requires variation in price suppression that is independent of poverty, and the paper does not claim to have identified a fully convincing source of it.

Measurement of forgone care. Section 5.2 notes that unmet need suppresses measured expenditure. Household surveys measure reported illness and reported treatment imperfectly, and reporting itself varies with education and access, in the same direction as the treatment. This biases estimates of the effect of the shock towards zero and may do so substantially.

Political feasibility. D1 requires converting hidden cross-subsidies into visible budget lines, and D4 requires firms to release commercially valuable data. Punjab's 2025 termination of public-hospital cover under an established programme, preceded by two years of unpaid premium invoices to the scheme's administrator, is evidence that health-financing commitments in this setting are durable neither fiscally nor across political cycles. A design whose benefits accrue over years and whose costs are visible in the first budget faces an adverse political-economy gradient, and the paper does not have a solution to that problem.

Data access and ethics. Designs II and III require data held by regulators and firms under legitimate duties of confidence. The consent architecture, anonymisation and disclosure control described in Section 5 are binding design constraints, not procedural additions. Design III would require institutional review, participant consent and regulator approval before execution, and its randomisation of decision routing raises questions about the treatment of applicants that must be resolved before, not during, implementation.

8. Conclusion

The 2022 floods were recorded as a loss of USD 14.9 billion in physical assets, of which health facilities accounted for roughly seven tenths of one per cent. That accounting is not wrong; it answers a different question from the one that matters for household welfare. In a health system financed at the point of use, a shock that destroys an eighth of the country's health infrastructure while raising malaria incidence fivefold transfers its cost onto household budgets, and it does so most heavily on households that were already, before the flood, among the most likely to face catastrophic medical spending.

Whether that cost is pooled depends on two frictions that this paper argues are one. Where the price of risk is suppressed or never established, cover is not written; where cover is not written, no record accumulates; and where no record accumulates, the algorithmic underwriters now arriving on Pakistan's digital payment infrastructure will decline the same households a second time. The claim that the frictions are sequential rather than parallel is the paper's central proposition, and it carries a practical corollary: because cover generates the data that widens future eligibility, the social return to establishing honest pricing in a thin-data district exceeds the private return, and the gap is an efficiency argument for subsidy that conventional appraisal omits.

The corollary also runs the other way, which is the more useful direction for policy. Cover written today is a record written today. A district given honestly priced, well-recorded parametric protection should, within a few years, present a denser file and a narrower gap between what a model and a human officer will approve. That is a testable prediction, and Section 5 specifies how to test it with data that already exist.

Pakistan is unusually well placed to attempt this. It holds a functioning instant-payment system, a national targeting registry, an operating social health insurance scheme, and a household survey series that brackets the largest climate shock in its recent history. What it lacks is a priced product and the data that only a priced product generates. The design proposed here is intended to supply both at once, and the empirical agenda is intended to establish whether it works.

Declarations

Funding. [No external funding was received / grant details.]

Conflicts of interest. The author declares none.

Data availability. No new data were generated. All statistics are drawn from the cited public sources. Figures 1 to 4 are constructed from values reported in the cited assessments and published studies; construction code is available on request.

Ethics. The proposed underwriting audit (Design III) would require institutional review, participant informed consent and regulator approval prior to execution.

Author contributions. [Single author.]

References

Ahmad, S., et al. (2025). The law of diminishing marginal returns in health insurance: evidence on child health outcomes from Pakistan's Sehat Sahulat program. *Frontiers in Public Health*, 13. doi 10.3389/fpubh.2025.1729440.

Ahmed, S., et al. (2024). Socioeconomic inequalities in out-of-pocket and catastrophic health expenditures in Pakistan. *International Journal of Public Health*, 69. doi 10.3389/ijph.2024.1607313.

ARY News (2025). Punjab suspends Sehat Card treatment outside province. 1 July. Lahore.

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5), 941–973.

Banerjee, A., Finkelstein, A., Hanna, R., Olken, B. A., Ornaghi, A. and Sumarto, S. (2021). The challenges of universal health insurance in developing countries: experimental evidence from Indonesia's national health insurance. *American Economic Review*, 111(9), 3035–3063. doi 10.1257/aer.20200523.

Barnett, B. J., Barrett, C. B. and Skees, J. R. (2008). Poverty traps and index-based risk transfer products. *World Development*, 36(10), 1766–1785.

Bashir, S. and Kishwar, S. (2021). Incidence and determinants of catastrophic health expenditures and impoverishment in Pakistan. *Public Health*, 197, 42–47.

Berg, T., Burg, V., Gombović, A. and Puri, M. (2020). On the rise of FinTechs: credit scoring using digital footprints. *Review of Financial Studies*, 33(7), 2845–2897.

Bonelli, M. (2026). Data-driven investors. *Review of Financial Studies*, 39(7), 1909–1969. doi 10.1093/rfs/hhaf078.

Clarke, D. J. and Dercon, S. (2016). *Dull Disasters? How Planning Ahead Will Make a Difference*. Oxford: Oxford University Press.

Cole, S., Giné, X., Tobacman, J., Topalova, P., Townsend, R. and Vickery, J. (2013). Barriers to household risk management: evidence from India. *American Economic Journal: Applied Economics*, 5(1), 104–135.

Commonwealth Fund (2026). International health care system profiles: Pakistan. New York: Commonwealth Fund.

Congressional Research Service (2022). *Pakistan's 2022 Floods and Implications for U.S. Interests*. Report IF12211. Washington, DC.

Dawn (2023). Sehat Sahulat Programme faces closure in Punjab over unpaid dues. Report on State Life Insurance Corporation correspondence with the Government of Punjab, March. Karachi.

Dawn (2025). Universal Health Insurance discontinued in Punjab's public hospitals from 30 June. Report on the Punjab Health Initiative Management Committee notification, June. Karachi.

Fahim, A., Qureshi, H. A. and Sethi, A. (2025). Public health collapse during floods in Pakistan: a call for national and global response. *Pakistan Journal of Medical Sciences*, 41(10). doi 10.12669/pjms.41.10.13355.

Finkelstein, A., Taubman, S., Wright, B., Bernstein, M., Gruber, J., Newhouse, J. P., Allen, H. and Baicker, K. (2012). The Oregon health insurance experiment: evidence from the first year. *Quarterly Journal of Economics*, 127(3), 1057–1106.

Froot, K. A. (2001). The market for catastrophe risk: a clinical examination. *Journal of Financial Economics*, 60(2–3), 529–571.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T. and Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *Journal of Finance*, 77(1), 5–47.

Gertler, P. and Gruber, J. (2002). Insuring consumption against illness. *American Economic Review*, 92(1), 51–70.

Goldstein, I., Jiang, W. and Karolyi, G. A. (2019). To FinTech and beyond. *Review of Financial Studies*, 32(5), 1647–1661.
Government of Pakistan, Asian Development Bank, European Union, United Nations Development Programme and World Bank (2022). *Pakistan Floods 2022: Post-Disaster Needs Assessment*. Islamabad.
Hallegatte, S., Vogt-Schilb, A., Bangalore, M. and Rozenberg, J. (2017). *Unbreakable: Building the Resilience of the Poor in the Face of Natural Disasters*. Washington, DC: World Bank.
Jan, R., Jan, M. and Jan, N. (2025). From 2022 to 2025: preventing the escalating health crisis of Pakistan's floods. *Annals of Global Health*. doi 10.5334/aogh.4955.
Jensen, N. and Barrett, C. (2017). Agricultural index insurance for development. *Applied Economic Perspectives and Policy*, 39(2), 199–219.
Khan, S. A., et al. (2022). Contextualising Sehat Sahulat Programme in the drive towards universal health coverage in Khyber Pakhtunkhwa, Pakistan. *Khyber Medical University Journal*, 14(1), 63–70.
Koijen, R. S. J. and Yogo, M. (2015). The cost of financial frictions for life insurers. *American Economic Review*, 105(1), 445–475.
Mobarak, A. M. and Rosenzweig, M. R. (2013). Informal risk sharing, index insurance, and risk taking in developing countries. *American Economic Review*, 103(3), 375–380.
Mohanan, M. (2013). Causal effects of health shocks on consumption and debt: quasi-experimental evidence from bus accident injuries. *Review of Economics and Statistics*, 95(2), 673–681. doi 10.1162/REST_a_00262.
Mullainathan, S. and Obermeyer, Z. (2022). Diagnosing physician error: a machine learning approach to low-value health care. *Quarterly Journal of Economics*, 137(2), 679–727. doi 10.1093/qje/qjab046.
Nasir, K., et al. (2024). The urgency of climate-resilient health systems in Pakistan: lessons from the 2022 floods. *International Journal of Public Health*, 69. doi 10.3389/ijph.2024.1607981.
Obermeyer, Z., Powers, B., Vogeli, C. and Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. doi 10.1126/science.aax2342.
O'Donnell, O., van Doorslaer, E., Wagstaff, A. and Lindelow, M. (2008). *Analyzing Health Equity Using Household Survey Data*. Washington, DC: World Bank.
Oh, S. S., Sen, I. and Tenekedjieva, A.-M. (2026). Pricing of climate risk insurance: regulation and cross-subsidies. *Journal of Finance*, 81(3), 1161–1215. doi 10.1111/jofi.70029.
Pakistan Bureau of Statistics (2026). *Household Integrated Economic Survey 2024–25: Final Report*. Islamabad.
Sood, N., Bendavid, E., Mukherji, A., Wagner, Z., Nagpal, S. and Mullen, P. (2014). Government health insurance for people below poverty line in India: quasi-experimental evaluation of insurance and health outcomes. *BMJ*, 349, g5114.
State Bank of Pakistan (2025). *Annual Payment Systems Review FY2024–25*. Karachi: Digital Financial Services Group.
State Bank of Pakistan (2026). *Quarterly Payment Systems Review, Q3 FY2025–26*. Karachi.
Wagstaff, A. and van Doorslaer, E. (2003). Catastrophe and impoverishment in paying for health care: with applications to Vietnam 1993–1998. *Health Economics*, 12(11), 921–933. doi 10.1002/hec.776.
Wagstaff, A., Flores, G., Hsu, J., Smitz, M.-F., Chepynoga, K., Buisman, L. R., van Wilgenburg, K. and

Eozenou, P. (2018). Progress on catastrophic health spending in 133 countries: a retrospective observational study. *Lancet Global Health*, 6(2), e169–e179.
World Bank (2022). Pakistan: flood damages and economic losses over USD 30 billion and reconstruction needs over USD 16 billion. Press release, 28 October. Islamabad.
World Health Organization (2022). *Pakistan Floods Situation Reports* and *Health Cluster Appeal*. Cairo and Islamabad: WHO Regional Office for the Eastern Mediterranean.
Xu, K., Evans, D. B., Kawabata, K., Zeramdini, R., Klavus, J. and Murray, C. J. L. (2003). Household catastrophic health expenditure: a multicountry analysis. *The Lancet*, 362(9378), 111–117. doi 10.1016/S0140-6736(03)13861-5.