SAFWAN ALMODHAYAN



OPINION AND ANALYSIS

Editorial date: November 2, 2026

The Insurer That Paid Before Seeing the Damage

One Storm, Three Answers

On July 3, 2024, Hurricane Beryl moved along Jamaica's southern coast, its center remaining offshore. Two days later, the finance minister reported that the center had passed roughly forty-five miles south of Kingston, that parts of the island had suffered substantial damage, and that fruit remained on trees in other areas. In the days that followed, insurance documents revealed something requiring more attention than the wind reports: nothing was due under the sovereign catastrophe bond worth $150 million, yet about $16.3 million was due under the Caribbean Catastrophe Risk Insurance Facility's tropical-cyclone policy and a further $10.3 million under its excess-rainfall policy. There were not two storms, and the insurers did not disagree about when the disaster occurred. There was one storm and three different ways of asking it a question.[1][2][3]

In a public-finance office, this looks like a file in disaster-risk management. In a home that has lost its roof or road access, it can look like a cruel paradox: how can a financial contract know that a hurricane happened yet see no reason to pay? The exact answer is not that the computer failed to notice the destruction. It was never assigned to observe it. Its task was to perform a calculation under terms agreed before hurricane season. That differs from evading an established indemnity, but it reveals a foundational choice in insurance: the interval between event and money can be shortened if the payment is determined by a measure of the event rather than a measurement of every loss it caused.

When an Index Replaces Loss Inspection

This is called parametric insurance. In traditional indemnity insurance, the most sensitive work begins after the incident: establishing damage, determining its causal relationship to the insured peril, valuing what the policy accepts, and settling the claim. A typical parametric contract states in advance a payment or a formula if an observable event crosses a defined threshold. The trigger might be earthquake intensity, a hurricane's central pressure and track, rainfall measured across a particular time and place, or a modeled loss derived from those observations. This does not mean all parametric contracts are alike or verification disappears. Verification moves from inspecting each home to validating event data, calculating the index, and confirming the contract's eligibility.

To understand the epistemic price of that shortcut, three things casually grouped under “disaster” must be separated. First is the natural phenomenon itself: air pressure, wind speed, millimeters of rain. Second is the exposure of things to it: where is the house, at what elevation, and do hills or trees stand before it? Third is their vulnerability: roofing material, drainage condition, insulation quality, household reserves, and evacuation speed. The hurricane can be one event while the losses are thousands of different events. Wind speed is not physically equal to the value of damage, just as the value of damage is not economically equal to the money a municipality needs in the first seven days.

Without pretending to solve it with a simple equation, the idea can be written this way: the index is generated by hazard data and a calculation method; loss is generated by hazard, exposure, vulnerability, and response; payment is generated by the index and policy terms. The relationship among the three chains is statistical and must be built and tested. It is not identity. Insurance specialists call the gap between the index result and the actual need for protection **basis risk**. A crop may burn while the index threshold remains unmet; or the trigger may produce more money than the insured's direct loss, or a payment without an equivalent direct loss. The first case is especially dangerous for someone who bought coverage to avoid being left alone on the worst day. The second is not necessarily a contractual error: the payment may be intended for emergency expenses that need not match a repair invoice.

Jamaica and Layers of Protection

An absence of payment is not, by itself, evidence of unacceptable basis risk. Someone buying protection designed for a rare, devastating hurricane cannot infer, merely because a less destructive hurricane occurred, that the instrument promised coverage. Jamaica's Ministry of Finance said on July 5 that the catastrophe bond formed the highest layer of a system including reserves, emergency financing, and tropical-cyclone and excess-rainfall policies, and that Beryl's track and intensity did not meet that layer's conditions. When the World Bank

issued the $150 million bond in April 2024, it explained that it covered named tropical cyclones under specified location and intensity rules over four hurricane seasons. Excluding certain events can be a disclosed design, not a discovery made after disaster.14

The fairness of a design is not tested by its terms alone. Was there a gap between layers that nothing filled, and did the buyer know where it was? The numbers that emerged from Beryl become more important here than the number that did not. CCRIF's annual report shows that the Jamaican government became entitled to about $26.59 million under the tropical-cyclone and rainfall policies. The catastrophe bond did not pay it; two mechanisms lower in the protection architecture did. The exact question was whether the system as a whole performed its function at the realized scale of damage, not whether every instrument in it paid.3

A later event changes the meaning of that first silence. Hurricane Melissa struck Jamaica in October 2025. On November 7, the World Bank announced that an independent calculation agent, AIR Worldwide, had used track and central-pressure data from the US National Hurricane Center and found that the conditions had been met for a **full $150 million payment** under the 2024 bond. The contract that did not pay for Beryl did not abandon its promise under Melissa; it responded to different inputs. In May 2026, the World Bank announced the pricing of a new $200 million bond to replace the exhausted coverage. This sequence does not prove the former bond was ideal, but it prevents the hasty conclusion that its first nonpayment showed fraud or incapacity.56

Malawi: The Crop Changed, the Model Did Not

That story alone could inspire too much confidence in indices. Thousands of kilometers away, another experience supplies a harder test. During Malawi's 2015/2016 season, farmers suffered a severe drought. The government had bought a policy from African Risk Capacity, or ARC, based on its Africa RiskView model. In the first calculation, the model said the number affected fell below the level required for payment. Conditions in the field indicated a broad food crisis, while the contract did not release its benefit. This was not simply a difference between a higher and a lower coverage layer. It was a contradiction between the peril as people lived it and the peril as the model described it.7

What did the computer get wrong? It did not emerge that a satellite had misrecorded rainfall. According to the insurer, a joint investigation found that the model assumed farmers grew maize requiring 120 to 140 days to mature, while use of a shorter-cycle variety, nearer ninety days, had spread. Different plant-growth schedules made the same distribution of rainfall affect the crop very differently. Accurate meteorological observation could not rescue a model carrying an outdated description of human activity. After field studies and revisions to assumptions and the related policy, ARC announced in November 2016 that about $8.1 million

was due, subject to approval of the implementation plan through which the money would be used. The advantage parametric insurance was designed to deliver—early financing—was delayed.78

This episode is a rare kind of knowledge in applied economics because it exposes error where it normally hides. “Maize” was not a minor agricultural detail entered incorrectly. It was a variable translating a chain of rainfall observations into an estimate of human need. A review confined to precipitation measurements might have praised the model. A review confined to household suffering might have condemned it without finding the cause. Joining the two levels is the real scientific work. Correct data are not enough when the thing one believes is being measured has changed since the model was last updated.

Nor does the story end with a corrected table and a celebration of an algorithm that finally worked. In May 2017, ActionAid published a sharply critical report. It said millions had faced food insecurity, the payment arrived late and was small beside total need, and questioned both the premiums' value and what adaptation spending might have achieved instead. ARC responded that the amount had been adjusted through a contractual technical mechanism rather than a negotiated grant, and said it had strengthened its review of inputs with participating countries. The disagreement should not be erased to manufacture a reassuring ending: the final settlement may be correct under the corrected contract, while the question remains of who bore the consequences during the months the correction required.89

This makes “speed” more complex than the count of days between disaster and transfer. The time to complete a payment includes observing the event, auditing its data, calculating the index, confirming policy eligibility, moving funds, and then delivering the benefit to the affected person. An insurer may pay a government within two weeks while the government needs several more to reach an isolated household. A contract promising speed but requiring its assumptions to be reopened after a dispute may lose much of its financial value. In a disaster, delay costs more than interest on an account. A farmer may sell a productive cow to buy food; a child may leave school; a village may borrow to finance water. The first loss becomes the cause of the next.

Which Part of the Risk Remained Unprotected?

Northern Kenya illuminates another side of the problem. Researchers from the International Livestock Research Institute and their colleagues examined livestock insurance linked, in one structure, to indices derived from satellite-observed vegetation. The scientific intuition is clear: when drought weakens pasture, livestock loss becomes more likely, and a vast region can be measured from space more cheaply than every animal can be visited. Yet a cow can fall

ill or be stolen; one grazing area can weaken while another does not; and a herder may protect a herd through a decision no sensor observes. These cases are not noise to be removed from a herder's life so it fits the model. They are part of what is supposed to be insured.

In a 2016 study, Nathaniel Jensen, Christopher Barrett, and Andrew Mude analyzed household data over eight half-year seasons. They found that the product reduced households' exposure to large covariate shocks by an average of **63%**, yet on average left insured households with about **69% of their original risk associated with high-loss events** once other sources of variation were included. The figures are not mathematically contradictory: the first concerns a particular component of shared risk, while the second concerns total residual risk from multiple sources. This is more honest than calling the product simply successful or unsuccessful. It protected against something important, and it did not protect against other important things.10

The question sold to insured people by some commercial presentations should therefore be reordered. It is not “Did the index release money?” but “Which part of the distribution of losses did that money reduce, for whom, and when?” A product can achieve a high statistical correlation between a rainfall index and aggregate losses across ten provinces while remaining weak for a household at the edge of a valley struck by a passing flood. Tables can be accurate on average and practically unfair where the person affected has no average waiting to compensate them. If the purpose is to protect a government budget rather than a home's wall, the broad index may be more suitable. The beneficiary and purpose determine how much error is acceptable, not the elegance of the statistical curve alone.

“Accuracy” itself must be unpacked. A rain gauge can measure excellently at its own location, while a summer storm falls on one district and misses another nearby. A satellite map can cover a whole country, while the area summarized by one cell hides variation in elevation, soil, and development. In a January 2024 technical feasibility study of smallholder insurance in El Salvador, Guatemala, and Honduras, the World Bank stressed proximity between the observing station and the insured. It cited an indicative range of three to five kilometers to reduce basis risk when station data are used, together with the need for a reliable long time series, often twenty to twenty-five years or more. This is not an engineering rule valid for every city and rainfall regime. It is a reminder that the quality of a number is not the same as the suitability of its spatial domain.11

Then comes the side rarely drawn on a promotional poster: the right time scale for measuring harm. Crops respond not only to total rain but to its timing relative to germination, flowering, and harvest. Four consecutive days of rain may be more damaging than a wet month whose rainfall is evenly distributed. Urban floods interact with how full drainage networks were before the storm, the position of inlets, and water levels in valleys. A contract based only on

monthly rainfall totals has taken an implicit position on the physical processes producing loss. That position can be tested, but it does not become correct when “artificial intelligence” is added to the dashboard.

Designing Layers That Reduce Basis Risk

Seen this way, CCRIF's development of its contracts is an experiment worth studying. In 2017 it introduced Aggregate Deductible Cover, or ADC, for tropical-cyclone and earthquake policies. The feature permits a limited payment when modeled loss falls below the principal policy's attachment point, subject to further conditions: for example, when the modeled loss is between 50% and 99% of the attachment point, or between 10% and 49% alongside a disaster alert from a specified source. CCRIF explicitly says the aim is to reduce the prospect of no payment despite damage on the ground. This is not comprehensive protection for every missed disaster; the maximum amount is governed by limits set out in the program.12

CCRIF has also developed mechanisms for extreme rainfall that address circumstances a single signal does not capture well, including a trigger associated with a wet rainfall season and one for severe localized events. The experimental insight is that a single index is not immutable destiny. If errors cluster near a threshold, across small spatial areas, or after the soil becomes saturated, a contract can be layered and draw on multiple signals. Every added signal brings observation, verification, and pricing costs and may create new errors. The objective is not to multiply indices for decoration, but to show that an addition lowers the probability of leaving the person in need without the intended amount.12

For a case that resists reduction to theory, consider Grenada's water utility. When Beryl passed on July 1, 2024, water and sewerage networks on the island and its dependencies were damaged. On July 23, CCRIF announced that the National Water and Sewerage Authority, NAWASA, had received **US$2,201,833** under a new parametric cover designed for water utilities, to help restore service. The insured entity did not have to finish a final inventory of every pipe before coverage was triggered. Yet one cannot leap from the announcement that a utility received money to the claim that water returned to every home that day. Civil engineering, labor, materials, access routes, and service priorities stand between payment and restarted pumps; an insurance policy cannot compress them into itself.13

Why might parametric insurance suit a water utility better than a precise promise to repair each home? Its early need can be defined differently: initial liquidity to rent equipment, bring in repair teams, purchase fuel, and secure temporary alternatives while the full loss picture emerges. If the payment slightly exceeds the cost of that initial response, the contract may still be performing its function. If the declared purpose is to compensate the exact value of every

broken window, atmospheric pressure alone will not settle the valuation. Contract design begins by naming **the function the money performs** and then choosing an appropriate measurement—not by admiring a new sensor and looking for something to sell around it.

The Edge Test and Protection When Needed

The fairness question appears at a small point on a curve, not at its center. Imagine cumulative-rainfall readings of 99 and 101 millimeters and a 100-millimeter payment threshold. If the contract pays a large lump sum once the threshold is crossed, two close readings can separate a full amount from zero. That does not make the two-millimeter difference trivial in every context; some systems have sharp physical limits. But the designer should show that the financial jump reflects a real jump in damage or need rather than a pricing table's preference for round numbers. Graduated payments, several thresholds, or a layer supplementing near-threshold payments can soften the boundary if price and structure permit.

Call this **the Edge Test**. Move the simulated hurricane track a few kilometers, vary the rain reading within the instrument's error margin, alter the germination period by two days, and ask: how much do payments change, and how much do actual losses change with them? If money jumps while need does not, or need jumps while money stays still, a point requiring redesign has been found. The Edge Test is not presented here as an approved regulatory standard. It is an implementable analytical proposal requiring independent loss data and actual contracts before any success rate can be attributed to it.

A second test, closer to the lives of insured people, can be called the **Protection-When-Needed Test**. Define in advance the loss or financing-need threshold the policy is meant to mitigate, then review historical events and whatever future simulations are feasible. Calculate the rate of nonpayment when that threshold is exceeded, the share of need funded by payouts, the distribution of time until benefit arrives, and the variation among districts or farms—not merely the national average payment. Separate cases in which the disaster lay above the layer the policy ever covered from those in which it fell within scope but the index represented it poorly. Combining the two gives the metric false severity or grants the design an acquittal it does not deserve.

The Price of Coverage and the Limits of Speed

A professional underwriter may answer that every modification increasing expected payouts also increases the premium. The basic trade-off is real. Widening coverage or allowing payment for more frequent events transfers additional risk to the funder. Jamaica's 2024 bond was not a refuge from every loss; under its terms, investors received a specified annual risk

margin for bearing the risk of major events. Money received by a stricken country is not financially created by an index. It comes from a contract moving priced risk onto other balance sheets. Reducing basis risk therefore does not mean making coverage broader without cost. It means making explicit the trade among premium price, the loss of missed protection, and the cost of delay.

The opposite error is equally dangerous: demanding that every parametric contract possess a precision even indemnity insurance lacks. Traditional insurance has exclusions, valuation disputes, and administrative delays, and it can remain inaccessible to whole groups of farmers or to public utilities with limited capacity to document every asset. Parametric coverage does not enter a perfect world and spoil it; it addresses real shortcomings. It has an important advantage: when event conditions are objective and independently verifiable, there is less room to dispute individual damage values before payment. That advantage creates a new responsibility. Who reviews the index's integrity, ensures the model uses current data, and explains in advance which events the policy will not pay for?

That responsibility grows as climate and human life change together. Rainfall during the next contract need not be distributed as it was over the preceding thirty years, while settlement and infrastructure can move faster than a loss model. A new drainage channel can reduce damage in a district even if rainfall stays constant; paving land can accelerate runoff without an obvious change in the volume of rain. A model learning from old claims reports also learns the history of who could document a loss and object, and who could not reach a settlement office. It may improve predictive accuracy on historical data while reproducing the same historical blindness toward those absent from the record.

It is not useful to place all of this under “algorithmic bias” as though that were a self-contained cause. There are contract design, instrument accuracy, changing agriculture, population movement, infrastructure maintenance, financing goals, and response timing. Good work identifies the layer in which disparity arose. Malawi was not repaired by adding more satellite imagery; it required asking farmers which variety they were actually planting. Jamaica was not repaired by correcting Beryl data; other protection layers performed their function while the catastrophe bond remained reserved for a more severe event. “Basis risk” names several different errors, and what looks like basis risk from a distance may instead be an intentional boundary around the insured peril.

From a Global Index to an Arab City

How might this knowledge move into Arab cities? Imagine a pilot—not an approved product here—to provide emergency funding after urban floods. It might combine radar and ground-station rainfall, elevation maps, drainage capacity, and the location of critical roads to estimate the early need for service restoration. The designer must decide whether the beneficiary is a

municipality paying to reopen roads, a factory stopped by water, or a household that lost its contents. Each has a different loss function. A measure based on citywide average rainfall may mistreat a low-lying district beside a flood channel. The policy does not become fairer because the index receives an Arabic name or an elegant interface.

In such a hypothetical case, three domains can be tested before sale is considered. First, do the data actually warn of damage in the intended locations, or merely describe weather at a broader scale? Second, does the contract cover the layer of loss the client needs, or does the client imagine compensation for every damage after buying only early financing? Third, is there a published way to review the model and data sources and challenge measurement errors without turning every claim into an individual dispute that restores traditional insurance's delay? The premium must then be compared with a cash reserve, emergency credit, or existing drainage improvement. A utility should not be insured against water that could be diverted more durably at lower cost.

A good contract, then, does not promise to know damage from the sky. It discloses the limits of its knowledge and buys precious time within them. A government may reasonably pay a premium for dependable financing during the first days of a major disaster. A poor household may reasonably decline a policy tied to a station too distant for local rainfall patterns, however attractive the promise of a fast transfer. The better path may be a protection network combining an initial reserve, broad-risk insurance, and a field plan for reaching people—not one index promising to replace entire institutions with numbers.

Jamaica's experiences with Beryl and Melissa teach that zero can be the correct execution of a contract and $150 million can be the correct execution of the same contract. Malawi teaches that a disciplined calculation can describe a crop people no longer grow. Northern Kenyan livestock teaches that protecting one part of risk does not make the rest disappear. Grenada's water utility teaches that rapid payment can be valuable even when nobody claims it equals the final damage bill. These are not scattered proofs of speed's virtue or indices' vice. They are different limits around one task: converting an event we cannot stop into a defined capacity to act when it occurs.

Ultimately, a contract does not deserve to be called smart because it dispenses with a loss adjuster walking among homes with a notebook. It deserves the name when it knows, with modest and testable precision, what its instrument can see, what it cannot see, and what it actually promised to protect its holders against. On the morning of a disaster, people do not merely need a number that wins a data test. They need money arriving while it can still make a difference, and an honest explanation when it does not. The heaviest loss an insurer leaves behind is not always an unpaid amount. It may be trust whose premiums were paid for years, only for its owners to discover that the contract's definition of rain was narrower than the sky above them.

References

1. Jamaica Ministry of Finance, Update on the Disaster Risk Financing Instruments, July 5, 2024.
2. CCRIF, Jamaica Excess Rainfall Event Briefing: Hurricane Beryl, July 12, 2024.
3. CCRIF, Annual Report 2024–2025, November 2025.
4. World Bank, World Bank Returns to the Cat Bond Market, Providing Financial Protection to Jamaica, April 25, 2024.
5. World Bank, Hurricane Melissa triggers 100% payout of $150 million catastrophe bond for Jamaica, November 7, 2025.
6. World Bank, World Bank Prices Catastrophe Bond for Disaster Protection for Jamaica, May 18, 2026.
7. African Risk Capacity, Malawi case statement, November 14, 2016.
8. African Risk Capacity, Response on the Malawi experience, July 10, 2017.
9. ActionAid, The Wrong Model for Resilience, May 24, 2017.
10. Nathaniel Jensen, Christopher Barrett, and Andrew Mude, Index insurance quality and basis risk: evidence from northern Kenya, October 2016.
11. World Bank, Feasibility study for smallholder insurance in Central America, January 2024.
12. CCRIF, Payout record and explanation of ADC and rainfall triggers, accessed November 2, 2026.
13. CCRIF, Payout to Grenada's National Water and Sewerage Authority, July 23, 2024.