The Economics of Sovereign Parametric Insurance in Low- and Middle-Income Countries

Michael R. Carter and Marcos M. Sugastti*

February 1, 2026

Abstract

The increased frequency and severity of natural disasters has spawned the creation of multilateral risk management facilities that offer sovereign parametric insurance contracts that provide governments with budgetary support for infrastructure replacement and excess social protection payments that accumulate in the wake of hurricanes and droughts. While the argument for prearranged financing is compelling, there is a paucity of economic analysis concerning how much coverage a government should optimally purchase, especially taking into account the fact that parametric insurance will sometimes over- or under-predict true losses. We answer this question using a formal model of the decision to purchase parametric insurance to cover stochastic social protection payments. Assuming that the government has a fixed social protection budget for shock responses and that its goal is to maximize the expected well-being of the drought-affected population, we show that optimal insurance coverage is highly sensitive to both the predictive accuracy of the parametric disaster index and the pricing of the insurance relative to its actuarially fair level. Using realistic parameters drawn from a scheme designed to insure Kenya's social protection program for its drought-prone regions, we further show that while the optimal amount of insurance is positive, it is lower than typically imagined. Our analysis underscores the need to improve the predictive accuracy of parametric insurance and to stress-test it with rigorous public finance models.

Keywords: Sovereign Parametric Insurance, Adaptive Social Protection ***JEL* Codes:** O23, Q54

*Carter: Department of Agricultural and Resource Economics, University of California Davis, Davis, CA, 95616 USA, NBER and University of Cape Town; Sugastti: Division of Global Agriculture and Food Systems, University of Edinburgh.

1 Introduction

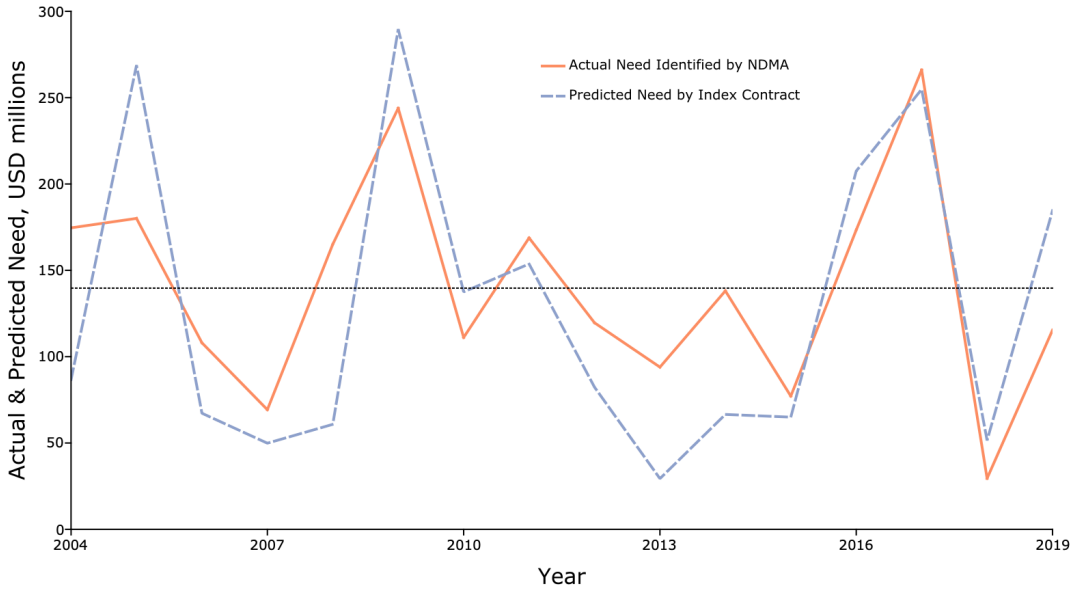
Natural disasters strain the budgets of governments confronted by broken infrastructure and an increased number of citizens in need of income support. In response to these stochastic budgetary needs, a new class of sovereign parametric insurance contracts has emerged. These contracts deliver governments budgetary support when a parametric index correlated with damage (e.g., wind-speeds for hurricanes and rainfall deficits for droughts) crosses a pre-determined threshold.[1] Clarke and Dercon (2016) make the case that such insurance can make disasters politically "dull" and enable faster disaster response. In this paper, we complement their work by analyzing when it makes public finance sense to purchase sovereign index insurance coverage and how much coverage should be optimally purchased. Motivated by an actual sovereign index contract designed to help Kenya manage its adaptive social protection program, we show that optimal insurance coverage is highly sensitive to both the predictive accuracy of the disaster index and the pricing of the insurance relative to its actuarially fair level. When predictive accuracy is high and insurance is priced close to its actuarially fair value, a government with a fixed social protection budget for shock responses will optimally choose a high coverage level if its goal is to maximize the economic well-being of its drought-affected population. However, as the accuracy of the index declines or insurance markups increase, our analysis shows that optimal insurance coverage falls quickly.

Similar to other countries, Kenya devised an adaptive social protection strategy that expands cash transfers when natural disasters occur.[2] To help fix the ideas that underlie

[1] Several specialized funds, owned by member states and further supported by international development agencies, now operate in different regions, including Africa Risk Capacity (ARC), the Pacific Catastrophe Risk Insurance Company (PCRIC), and the Caribbean Catastrophe Risk Insurance Facility (CCRIF).

[2] The Hunger Safety Net Programme (HSNP) is Kenya's social safety net initiative targeted specifically at the country's arid and semi-arid regions, which are home to a large share of its poorest population. Initiated in

Figure 1: Actual Versus Predicted Social Protection Need Under Sovereign Parametric Insurance



this paper, the solid orange line in Figure 1 graphs the social protection budget, measured in millions of USD, that Kenya's National Drought Management Agency (NDMA) would have needed to finance shock-responsive cash transfer scale-ups following droughts in Kenya's drylands over the 2004–2019 time period. The horizontal dotted line represents average need over this period of $140 million annually, while the dashed blue line graphs the annual payouts that would have occurred under a proposed sovereign parametric insurance contract.[3] The parametric index for this contract was based on satellite-based measures of drought.[4]

2007 as a fairly conventional unconditional cash transfer program, HSNP later came under the administration of Kenya's National Drought Management Authority (NDMA). Under NDMA's management, the program scales up payments to pre-qualified households when drought or flood indices maintained by the agency signal increased need. Details on the program are provided in National Drought Management Authority (2024), which documents the scaling mechanism and explicitly acknowledges the financing challenge when it notes that payments are put on hold when the available budget is insufficient to meet the indicated need.

[3]The contract was designed by ARC (see Note 1). Details on how the contract would have performed historically were provided by the International Livestock Research Institute.

[4]The broad approach of using satellite data to predict drought and social protection need in Kenya's arid regions builds on a micro-level index insurance product called Index-based Livestock Insurance (IBLI). Chantarat et al. (2013) discusses the initial IBLI logic and design, while Jensen et al. (2024) summarize subsequent refinements and use of this approach for micro level index insurance contracts.

Impressionistically, the need predicted by the parametric insurance index tracks the actual need declared by NDMA tolerably well. At the same time, the index under-predicts need in multiple years. The most extreme example is 2008, when predicted need stood at only $60 million, while actual need topped $160 million. In the index insurance literature, this kind of under-prediction is known as downside basis risk or false negative events. Also visible in the figure are years in which the index over-predicts need. While this kind of upside basis risk event (or false positive events) may seem less concerning than under-predictions, it is also economically damaging. Payouts resulting from need over-prediction increase the overall insurance premium, and when insurance is marked up over its actuarially fair price, markups on these payouts represent a net leakage of funds from the social protection system that will never reach the target population. In summary, these two forms of misalignment have asymmetric welfare consequences: over-payments raise insurance costs through higher premiums, while under-payments fail to fully compensate households for drought-induced losses.[5]

Against this backdrop, this paper develops the economic foundations for answering when and how much insurance a sovereign entity should buy as a function of the markup cost and the accuracy of the underlying parametric insurance contract. Following a literature review (Section 2), Section 3 presents the optimization problem for a government that chooses a level of insurance coverage to maximize the expected utility of its drought-affected population, subject to a fixed social protection budget for shock responses that must cover any direct payouts plus insurance premium. Section 4 then approaches this problem numerically, under the assumption that the government has a fixed budget equal to the expected amount of funding required to offset drought-induced losses, evaluated relative to a reference

[5]This framework highlights a tension between two functions of social protection. One is income stabilization, which seeks to smooth resources across states of the world. The other is poverty alleviation, which prioritizes closing poverty gaps in periods of high need. Parametric insurance can support the former by shifting resources into bad states, but may undermine the latter when basis risk is substantial and transfers fail to reach households when need is greatest.

consumption level equivalent to the poverty line of $1.90 per person per day. We use this reference level solely as a normalization for welfare comparisons.[6] Note that this assumption implies that the government could fully insure its shock-response obligation if insurance were priced at its actuarially fair level.

Under a base case scenario calibrated to the risk and parametric index accuracy levels in Figure 1, we then use simulation analysis to solve for the optimal level of insurance coverage. Assuming no insurance markup, optimal government policy under this scenario is to insure only against the most extreme need events that occur 29% of the time. This amount falls well below full insurance, even though full coverage is affordable given the assumed budget. Under this optimal contract, the certainty equivalent utility of affected people rises to $1.82 per person per day, a small improvement over what the government could attain by purchasing no insurance and directly transferring money to the drought-affected population instead. Importantly, if the government were to buy more insurance than is optimal—a lower-deductible contract covering the most extreme 78% of events—it would reduce the certainty equivalent well-being of the affected to $1.81 per person per day, leading to no discernible gains over what the government could achieve without any insurance.

Section 5 then expands the analysis to examine how insurance markups and variation in predictive accuracy of the contract impact optimal insurance demand. When the markup rises to 25%, the layer of risk covered by the optimal contract falls to the most extreme 10% of events. Even if the insurance markup remains zero, optimal insurance coverage falls to a similar level when basis risk worsens. The combination of higher basis risk and a 50% markup pushes optimal insurance to essentially zero.

Despite the conceptual appeal of sovereign index insurance as a budgetary instrument for

[6]This normalization should not be interpreted as a claim about the adequacy of baseline social protection transfers in the absence of a shock. In practice, baseline social protection budgets are insufficient to lift all households to this reference level. Our analysis therefore interprets the poverty line as a welfare benchmark, and the resulting welfare metrics should be understood as relative comparisons across financing arrangements rather than absolute measures of achieved well-being.

low- and middle-income countries, our analysis shows that the optimal amount of insurance is highly sensitive to both the pricing of the insurance and the reliability of the underlying parametric insurance index. Section 6 concludes the paper with reflections on technological options for improving insurance index accuracy, and on institutional procedures that governments might follow to ensure that scarce social protection resources are used to maximize the well-being of vulnerable populations.

2 Disaster Risk Management and Insurance Quality

Given the increased frequency and intensity of climate-related shocks that disproportionately affect vulnerable and poor individuals, governments may consider engaging in anticipatory risk financing as a social protection priority. A strand of the literature provides a theoretical discussion on possible approaches to do so.

2.1 Public Finance and Disaster Risk Management

Economists have long argued that governments are uniquely positioned to pool and absorb large covariate risks that would otherwise overwhelm private agents. For example, Arrow and Lind (1978) postulate that because governments can diversify risk across the entire population, they can behave as if risk-neutral. That is, they may effectively ignore uncertainty when making public investment decisions, thereby eliminating the need for *ex ante* risk financing.

In contrast, more recent work (e.g., Ghesquiere, 2007) argues that the assumption of a risk-neutral government does not hold in many real-world settings. Liquidity constraints, institutional rigidities, and political frictions often prevent governments from delivering timely and adequate responses following large shocks. These limitations create a rationale for *ex*

ante financial preparedness.

Natural disasters are precisely the kind of covariate shock that resist diversification and are macroeconomically disruptive. These events increase public expenditure needs for emergency response, infrastructure reconstruction, and expanded social protection. In low- and middle-income countries, where fiscal space is limited and sovereign borrowing can be costly or slow, natural disasters can lead to severe welfare losses unless governments proactively manage their financial risks. This threat has become more urgent as climate-related shocks increase in frequency and intensity, disproportionately harming vulnerable and poor people.

Governments seeking to manage disaster risk before shocks materialize have a suite of financial instruments at their disposal.[7] These include budgetary tools, contingent credit, and market-based risk transfer solutions. Contingency or reserve funds are budget lines earmarked for emergency response. Contingent credit refers to *ex ante* loan agreements that provide countries with access to liquidity following a shock and is typically offered by multilateral development banks and international financial institutions. Market-based risk transfer instruments, such as sovereign parametric insurance, allow governments to shift a portion of disaster risk to private sector actors (Beazley et al., 2019). In practice, these tools differ in their cost, speed, scalability, and reliability. A central challenge for policymakers is to determine how to layer and balance these different instruments to support timely and effective disaster response.

Clarke et al. (2017) model how governments should allocate resources across various disaster risk financing instruments. These authors formalize a model in which governments choose among budget allocations and reallocations, contingent credit, sovereign borrowing, and insurance to minimize the cost of disaster response. They ultimately argue in favor of

[7]The World Bank has invested heavily in a set of disaster risk management tools designed to help governments financially prepare for climate shocks and other disasters, including thruough emergency borrowing and sovereign parametric insurance instruments. For more details, see Bjerde and Mountfield (2024): https://blogs.worldbank.org/en/voices/unlocking-new-crisis-response-tools-build-more-resilient-future.

a tiered approach in which each instrument is used sequentially, with insurance reserved for when cheaper instruments are exhausted. A key limitation of their framework is the assumption that insurance payouts perfectly match the government's stochastic expenditure needs. In their application exercise, they recommend using insurance to cover only high-severity events corresponding to the most extreme 7% of shocks, based on a numerical assumption of a 50% markup rate. The paper does not consider how this finding might change if index accuracy were to vary. Moreover, even under the assumption that payouts align with actual need, the authors conclude that insurance is best suited for infrequent but costly disasters.

While the Clarke et al. (2017) model offers valuable insights, it focuses on aspects such as cost and timing while abstracting away from the reliability of the instruments themselves. In particular, it largely assumes away the insurance-specific problem of basis risk—that is, a potential mismatch between the index or model used and the actual situation on the ground, leading to payouts that do not align with public expenditure needs. This assumption overlooks a defining feature of parametric insurance: the index that triggers payouts may only be imperfectly correlated with actual damages or social protection need, and may fail to fully capture the economic or social impact of a disaster. This basis risk introduces a fundamental wedge between financial preparedness and realized protection. Section 2.2 presents real-world examples of how this wedge has undermined past efforts to promote sovereign parametric insurance.

This reliability gap is not just a technical detail, but it is central to determining whether insurance delivers real public value and to shaping effective policy design. Questions such as how much insurance to buy, which layers of risk to cover, or how to value an imperfect insurance contract relative to other social protection regimes remain largely unaddressed. As we show in this manuscript, incorporating basis risk into a social welfare optimization framework substantially alters optimal insurance choice and reveals important links between coverage, cost, and predictability.

In the next section, we review how basis risk has been studied in the microinsurance literature, where concerns about index reliability, contract design, and behavioral response have received more sustained attention. We then build on these insights to develop a model of sovereign index insurance contract choice that incorporates both financial constraints and the potential for mismatches between payouts and need.

2.2 Index Insurance Quality

Ironically, the core strength and weakness of index insurance stem from the same source, namely, that damages and financial needs are not directly measured but are instead estimated using near real-time remote sensing data. This feature is a strength because it allows objective, verifiable criteria to trigger payments, allowing indemnity payments to be issued quickly. However, it is also a weakness, as the estimates may be inaccurate. As illustrated in the introduction, the underlying index may either under-predict or over-predict actual need. Commonly referred to as "basis risk," this feature of index insurance has largely been ignored in sovereign insurance literature—explicitly so by Clarke et al. (2017).

In contrast, the literature on micro index insurance—that is, index insurance sold to cover the losses of individual farms or other businesses—has given substantial attention to the basis risk problem. Key discussions of basis risk in micro insurance appear in Clarke et al. (2012) and Clarke (2016), the latter of which argues that basis risk can, in theory, eliminate demand for index insurance altogether. Several empirical studies support this claim, showing that basis risk substantially depresses index insurance demand (see, for example, Jensen et al. (2016), and the summaries in Miranda and Farrin (2012), Jensen and Barrett (2017), and Carter et al. (2017)).

The recognized importance of basis risk in microinsurance has led to various efforts to

improve the predictive accuracy of insurance indices.[8] Some approaches aim to create insurance zones that more accurately capture the losses experienced by individuals within them (Benami et al., 2025; Estefania-Salazar et al., 2025). More directly related to this paper, a complementary literature has emerged that applies standard expected utility theory to evaluate the quality of index insurance. This framework allows researchers to assess whether the purchase of a given index insurance contract improves the expected economic well-being of the insured party, and to determine which among competing contracts offers the greatest expected welfare gain.

Building on ideas developed in Carter and Chiu (2018), Kenduiywo et al. (2021) define the Relative Insurance Benefit (RIB) measure of index insurance quality as:

$$RIB = 100 \times \frac{C^I - C^N}{C^P - C^N}, \tag{1}$$

where C^N is the certainty equivalent of the expected utility an individual would experience without insurance, C^I is the certainty equivalent under an index insurance contract, and C^P is the certainty equivalent under a hypothetical perfect (i.e., zero basis risk) insurance contract.[9] This approach is easily adapted to our public finance problem, where we assume the government's goal is to maximize the expected utility of its citizens who are vulnerable to droughts (see Section 3.2).

The RIB measure defined in Equation 1 has a convenient interpretation as it captures the percentage of the potential welfare gain from insurance that is actually achieved by a parametric insurance contract. As discussed by Kenduiywo et al. (2021) and illustrated using microdata from Kenya, both false negative and false positive events reduce the RIB.[10]

[8] Benami et al. (2021) summarize recent progress in this area.

[9] Given a time series of household income data in the absence of insurance, along with the ability to calculate the corresponding insurance index over the same period, each of these terms can be easily estimated.

[10] When $RIB < 0$, the index insurance contract is not worth buying as it reduces the expected utility of the insured household or business.

Kenduiywo et al. (2021) show how the RIB measure can be used to compare alternative remote sensing technologies, while more recent studies by Benami et al. (2025) and Estefania-Salazar et al. (2025) apply the RIB framework to identify optimal index insurance zone designs.

While it is clearly more difficult to accurately predict losses at the micro scale than at a national scale, sovereign index insurance programs have seen high-profile "basis risk events" (Johnson, 2021). As discussed by Chamberlain and Bernards (2024), the ARC Sovereign Risk Fund (see Note 1 above) faced financial strain after consecutive instances of need under-prediction prompted many participating countries to either withdraw from the pool or reduce their coverage levels. Given this experience, it is simplistic to dismiss basis risk as irrelevant in the sovereign insurance context. The goal of this paper is to offer an economic framework—akin to the RIB measure—that supports informed public decision-making about how much a nation should optimally rely on sovereign index insurance to manage climate and disaster risk.

3 Optimal Financing of Social Protection Under Budget Constraints

This section models the decision problem faced by a government that seeks to finance shock-responsive transfers for a population of citizens *ex ante* vulnerable to droughts. The govenment has a fixed and limited social protection budget allocated for this purpose. The challenge is compounded by the fact that the number of people requiring assistance in any given period is random. To give structure to the problem, we assume that the government's annual budgetary allotment equals the expected amount of funding required to offset drought-induced losses for the vulnerable population. This assumption, of course, implies that the annual

budget will not suffice to fully compensate affected households in years when drought shocks generate scale-up needs that exceed their average level.

To manage such budgetary shortfalls, the government has access to a sovereign parametric insurance policy, sold at some (non-negative) markup over its actuarially fair price. The government's public finance problem is to decide how much of its budget to allocate toward the purchase of insurance. Any portion of the budget not spent on insurance is available to directly meet social protection need, effectively constituting a reserve fund. We analyze this public finance decision under the assumption that the goal of the government is to deploy its fixed social protection budget to maximize the expected economic well-being of its drought-exposed population.

In practice, when shocks occur governments scale up existing social protection schemes and offer additional cash transfers to households affected by shocks. These scale-ups may be vertical, by increasing transfer amounts to individuals already receiving pre-shock support, or horizontal, by extending coverage to people not previously covered by baseline transfers. Typically, regular, non-shock-contingent transfers are budgeted *ex ante*, while droughts trigger additional, time-limited financing needs associated with compensating households for shock-related income losses. Here, we focus on the additional social protection obligation generated by random shocks.

In our analysis, we also ignore targeting issues—such as how the government identifies and transfers resources to those in need—by implicitly assuming that these challenges are the same regardless of the source of finance. We further assume that all households identified as requiring shock-response assistance in a given period face the same drought-induced loss and receive transfers according to a uniform allocation rule.[11] In this way, minimizing aggregate

[11]These assumptions deliberately characterize a best-case scenario for sovereign parametric insurance: we ignore heterogeneity in losses, delivery frictions, and targeting errors that would otherwise increase the welfare losses associated with under-payments. In practice, Kenya's NDMA does not follow a uniform allocation rule with the HSNP, its social protection program that provides regular cash transfers to chronically poor

drought-induced losses and maximizing the expected well-being of drought-vulnerable individuals become equivalent objectives. Appendix A analyzes a dual version of this problem in which the government must meet the drought-induced need of all vulnerable citizens and allocates its budget to minimize the expected economic cost of fulfilling that fixed social protection obligation.

3.1 The Shock-Responsive Social Protection Finance Problem

For the analysis, we denote by $p_t \in [0, N]$ the total number of people requiring drought-related assistance (individuals eligible for drought-response cash transfer scale-ups) in period t, where N is total number of people who are vulnerable to drought.[12] The probability density function $\phi_p(p)$ generates the number of people requiring assistance and $\overline{p}$ denotes the expected number of people affected by a drought. To keep matters simple, we assume that all people hit by a drought experience the same drought-induced loss and, after accounting for baseline social protection transfers, have income $y_0 < \underline{y}$, where $\underline{y}$ is the money-metric reference income level (equal to the \$1.90 poverty line) used solely for welfare normalization purposes. Let $G = (\underline{y} - y_0)$ denote the associated need gap for all individuals. If the government were to fully offset drought-induced losses for all affected people, then in period

households in arid and semi-arid regions and implements temporary cash transfer scale-ups during droughts. Their decisions regarding when and where scale-ups occur, as well as how limited resources are allocated across counties and households, follow pre-defined allocation quotas and prioritization rules based on poverty indicators, population shares, and administrative criteria. When fiscal resources are insufficient to cover all eligible households, these rules determine which households or locations receive assistance first, resulting in unequal transfers across affected households (National Drought Management Authority, 2024). Our abstracting from these prioritization mechanisms and assuming uniform per-capita transfers therefore yields a conservative benchmark for the value added by insurance, as it minimizes dispersion in post-transfer outcomes and attenuates the welfare costs of funding shortfalls. If insurance struggles to add value in this setting, it will do so even more once these real-world frictions are introduced.

[12]The variable p_t should be interpreted as the realized number of people requiring shock-response assistance in period t, drawn from a fixed population of individuals *ex ante* vulnerable to drought of size N. This formulation is consistent with ARC practices, where drought exposure is assessed at the level of predefined vulnerability polygons (such as livelihood zones or administrative units) and translated into estimates of affected populations using static vulnerability profiles previously calibrated from household survey data. Once these profiles are fixed, shock realizations are mapped into aggregate affected population counts, abstracting from household-specific losses at the time of payout (African Risk Capacity, 2016).

t, it would require funds in the amount $T_t = p_t G$. Note that the maximum possible social protection expenditure is therefore $T^{max} = NG$.

The annual budget allotment for social protection scale-ups due to droughts is $\overline{T} = \overline{p}G$, meaning that while on average the government has adequate funding to meet its shock-responsive social protection obligation, it will have insufficient current funds to do so whenever $p_t > \bar{p}$. By allowing the government to move resources from good states of the world (small p_t) to bad states (large p_t), sovereign parametric insurance has the potential to help the government better fulfill its social protection obligation by smoothing expenditures over time. How well insurance meets this potential depends on its cost and its reliability (that is, the basis risk problem mentioned above). If the government purchases insurance with deductible level δ,[13] it will have an available budget, B_t, to meet social protection needs of:

$$B_t = I(\delta, \tilde{p}_t) + \left(\overline{T} - \pi(\delta, m)\right),$$

where $I(\delta, \tilde{p}_t)$ are the insurance indemnity payments, which depend on the coverage or deductible level δ and the parametric insurance index $\tilde{p}_t$ that predicts the number of drought-affected people and triggers payments. The insurance premium $\pi(\delta, m)$ depends on both the coverage or deductible chosen and the markup rate m. Section 3.2 below discusses the nature of the insurance contract in more detail.

Our setup aligns with the role of a government agency tasked with administering a fixed, earmarked social protection budget for shock-responsive cash transfer scale-ups. This approach differs from that of Clarke et al. (2017), who model the government's disaster response as part of a broader public finance problem, allowing resources to move flexibly between

[13]The deductible level is the fraction of aggregate shock-response need that the government chooses not to insure. For example, if $\delta = 50\%$, the insurance only begins making payments when more than 50% of the vulnerable population is estimated to be hit by a drought. If the estimated affected share is 75%, the insurance company would pay $G \times 0.25 \times N$, corresponding to the portion of aggregate need above the deductible threshold.

uses—including social protection, infrastructure, and other public spending—and evaluating the opportunity cost of disaster response for social protection versus allocating funds elsewhere. The question we address is not how to minimize social protection costs in the event of a disaster, but how to deploy a committed level of funding to maximize the welfare of those affected by a shock under stochastic need. In our framework, the government's choice of insurance contract directly determines the size of the reserve fund: the premium is subtracted from the fixed budget, and the remaining funds are used to respond to realized need directly. This tighter social protection mandate reframes insurance as a tool for stabilizing the well-being of those vulnerable to shocks across states of the world, rather than simply smoothing liquidity. As a result, in our analysis, the value and risk of insurance depend not only on its cost, but also on how well it aligns with actual need.[14]

To study our problem, we assume that the government chooses whether and how much insurance to purchase in to maximize a Benthamite social welfare function focused only on the drought-affected population, defined as:

$$\begin{aligned} W_t\left(\delta, \tilde{p}_t | \bar{T}\right) &= \sum_{i \in P_t} u(y_{it}) \\ &= p_t \times u\left(\min\left[\underline{y}, y_0 + \frac{B_t}{p_t}\right]\right), \end{aligned}$$

where P_t is the set of drought-affected people in period t, $u(\cdot)$ is a conventional individual utility function.[15] The second line of the expression follows from our assumptions that all people requiring shock-response assistance experience the same drought-induced loss and are treated uniformly under the scale-up allocation rule. If the budget is insufficient to meet

[14] While our model focuses exclusively on sovereign insurance as the primary alternative to holding a reserve fund, a contingent credit line could be incorporated into the same framework as a complementary instrument in future extensions.

[15] This social welfare function exhibits what might be termed imperfect empathy for drought-affected people, in that it cares about their welfare only when their incomes fall below the reference income line, but not once they exceed it.

aggregate shock-response needs (i.e., $B_t < p_t G$), the available funds are split equally among affected individuals, and each receives an income of $y_0 + \frac{B_t}{p_t}$. If the budget is sufficient (i.e., $B_t \geq p_t G$), then each affected household is fully compensated for its drought-induced loss and reaches the reference income level, $\underline{y}$. Importantly, we assume that any insurance payouts the government receives in excess of realized shock-response need in a given period do not reach the target population but are instead reallocated for other purposes. In particular, these funds are not carried forward or saved for future periods. In this sense, from the perspective of the shock-response budget, such excess payouts therefore constitute a net leakage.

Given this notation, we assume that the government chooses how much insurance to purchase in order to maximize expected economic well-being of a drought-affected citizen:

$$\max_{\delta} E_{\tilde{p}} \left[W_t \left(\delta, \tilde{p}_t | \bar{T} \right) \right]$$

$$st:$$

$$B_t = I(\delta, \tilde{p}_t) + \left(\overline{T} - \pi(\delta, m) \right) \tag{2}$$

$$\pi(\delta, m) \leq \bar{T}$$

where the expectation is taken over the distribution of the random parametric insurance index, $\tilde{p}$. The second constraint simply requires that the contract must be affordable to the government given its annual social protection budget of $\bar{T}$. Before turning to the analysis of this optimization problem in Section 4, we first provide additional detail on the structure of the parametric sovereign insurance contract.

3.2 Basis Risk and the Parametric Insurance Contract

Parametric insurance issues payouts based on a verifiable index, $\tilde{p}_t$, that is intended to closely reflect actual damages without the need for time-consuming and often contentious processes

of loss verification. Because the index is an estimate (based on remote sensing satellite data in most cases), there will be some imprecision associated with it. For simplicity, we assume that the index estimate of the number of drought-affected people is equal to the true value plus random noise:

$$\tilde{p}_t = p_t + \omega_t, \tag{3}$$

where we assume that $\mathrm{E}[\omega_t] = 0$ and that ω_t is generated by a random process governed by the probability distribution function ϕ_ω.[16] Denote the joint distribution of p_t and ω_t as $\phi_{p,\omega}(p_t, \omega_t)$.

Given the index, the sovereign parametric contract offers indemnity payments in period t according to the following function:

$$I_t(\delta) = \max[0, \tilde{T}_t - \delta T^{max}],$$

where $\tilde{T}_t = G \times \tilde{p}_t$ is the estimated amount of shock-response need, and $T^{max} = G \times N$ is the total amount of money required to fully fund the drought-related needs of the entire vulnerable population, N. Note that the contractual deductible, $\delta \in [0, 1]$, determines the shock reduction burden that the government shoulders. If $\delta = 0$, then the insurance would reimburse the government for the drought-related needs of every person estimated to be affected. At the other extreme, a 100% deductible ($\delta = 100\%$) would imply no insurance coverage. A deductible of $\delta = \frac{\bar{p}}{N}$ would imply that the insurance would only cover shock-response need in excess of the average amount.

We assume that the market offers the government a parametric insurance contract priced

[16] These assumptions generously imply that the index is unbiased and correct on average. Importantly, this formulation captures uncertainty in aggregate scale-up need rather than heterogeneity in individual-level losses.

at its actuarially fair level, marked up at rate m:

$$\pi(\delta, M) = (1 + m) \times E\left(I_t(\delta)\right)$$

where expectations are take over the joint distribution of p_t and ω_t . The markup is intended to cover administrative and other costs needed to offer the insurance, and it is important for the government because every dollar spent on it reduces the government's expected shock-responsive cash transfer scale-up budget B_t by a dollar.

In summary, we can now rewrite the maximand for Problem 2 as:

$$W^* \equiv \max_{\delta} E \int \int \left[W_t\left(\delta, \tilde{p}_t | \bar{T}\right)\right] \phi_{p,\omega}(p_t, \omega_t)\, d\omega\, dp.$$

Because optimized social welfare is expressed in terms of the expected utility of a drought-affected person, it is straightforward to create an analogue measure to the index insurance quality measure developed for individual microinsurance, as defined by expression 1. To make the sovereign insurance quality measure more interpretable, W^* can be straightforwardly transformed into its corresponding certainty equivalent. Denoting this certainty equivalent as C^{δ^*}, the relative benefit of sovereign index insurance (RBS) can be written as:

$$RBS = 100 \times \frac{C^{\delta^*} - C^1}{C^\zeta - C^1}, \tag{4}$$

where C^1 indicates the certainty equivalent income of drought-affected people when no insurance is purchased (i.e., when the deductible is $\delta = 100\%$) and C^ζ is the certainty equivalent income of drought-affected people that could be obtained under a perfect, zero basis risk contract. Note that this ideal contract would arise if the prediction error ω_t were a degenerate random variable equal to zero in every period. The RBS measure thus quantifies how much of the potential benefit of sovereign index insurance can be achieved under any given

implementable insurance contract. Akin to the microinsurance RIB measure, the RBS measure allows for the comparison of competing contract designs—for example, those based on alternative satellite sensors or processing models.

4 Numerical Analysis of Optimal Insurance Coverage

The ability to find analytical solutions to the optimal contract choice Problem 2 is limited by the distributional assumptions we make on the stochastic processes driving the prediction index, $\tilde{p}$. While normality assumptions can potentially be analyzed, they would introduce a number of undesirable properties. Instead, we choose more flexible distributional forms and pursue our analysis using numerical simulations consistent with those distributions, as discussed in detail in Appendix B. A key advantage of this approach is that it allows us to calibrate our simulations to the empirical reality of the sovereign index insurance contract offered to Kenya, as illustrated in Figure 1.

After first laying out our numerical assumptions in Section 4.1, the remainder of this section uses simulation methods to derive the social welfare-maximizing insurance purchase decision under a base case scenario of modest basis risk and insurance priced at its actuarially fair level (zero markup), as presented in Section 4.2. Section 5 then explores the sensitivity of optimal insurance coverage to these assumptions.

4.1 Numerical Assumptions and Calibration

This section lays out the numerical assumptions that we will use to analyze our model of optimal government purchase of a sovereign parametric insurance contract designed to underwrite a budget-constrained government's ability to meet the needs of its citizens. All assumptions are at least roughly based on the empirical reality of the arid and semi-arid

Table 1: Numerical Assumptions

Individual Parameters		**Aggregate Parameters**		**Basis Risk Distributional Parameters**	
Reference Income Line, $\underline{y}$	USD 1.9	*Vulnerable Population,* N	5 m	*PDF for* ω	$\omega \sim N(0, \sigma^2)$
Income when Affected, y_0	USD 1.4	*Average Number of Affected,* $\bar{p}$	1.55 m	*Base Case*	$\sigma = 800{,}000$
Need Gap, G	USD 0.5	*State Budget Allocation for Social Protection,* $\bar{T}$	USD 776k per day	*Low Basis Risk*	$\sigma = 600{,}000$
Utility Function, $u(\cdot)$	CRRA	*PDF for* $\phi_p(p_t) = \{\dots\}$	$k = 2.4727$, $\lambda = 0.3497$	*High Basis Risk*	$\sigma = 1{,}000{,}000$

areas of Northern Kenya.

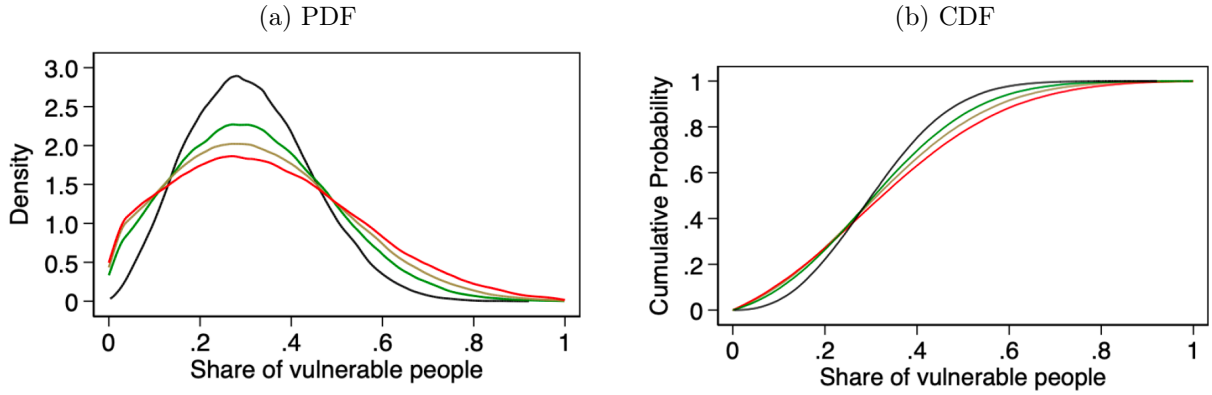
Table 1 details the parameters that we use for the simulation. The first column presents the individual-level assumptions, including a need gap of USD 0.5 for people affected by droughts, defined relative to a reference income line of USD 1.9 per day. The middle column describes our population assumption that 5 million individuals are vulnerable to being hit by a drought. Using 2004/05–2019/20 data from NDMA,[17] we find that the mean number of people in need for each period is $1,551,353$, with a standard deviation of $1,109,570$. Based on this data, we parameterize a Weibull distribution—as shown in the table—so that its mean and standard deviation match these numbers. Given these assumptions, we endow the government with a daily budget allocation of USD 775,677, an amount just sufficient to close the USD 0.5 daily need gap for the expected number of drought-affected individuals.

To capture imperfections in the estimation of need by a parametric insurance contract, and consistent with Equation 3, we add an additive error term, ω_t, to the Weibull draws of true need.[18] For our standard, medium basis risk base case scenario, we consider the

[17]Kenya's National Drought Management Agency (NDMA) is an agency of the Government of Kenya mandated to establish mechanisms which ensure that drought does not result in emergencies. NDMA chairs, with World Food Program (WFP), the Kenya Food Security Steering Group (KFSSG), a multiagency body comprised of government departments, UN agencies and NGOs concerned with food and nutrition security, which produce bi-annual assessments of the number of people in arid and semi-arid regions who require emergency relief during the Long Rain/Long Dry and Short Rains/Short Dry seasons. In this manuscript, we assume that these figures represent the "true" social protection needs. While this assumption can be debated, it provides a useful benchmark against which to measure the costs and benefits of the different insurance and other public finance approaches that we analyze.

[18]Truncation was performed by iteratively redrawing predicted need values that fell outside the 0 to 5,000,000 bounds until all estimates were within range.

Figure 2: True Versus Estimated Need Under Different Levels of Basis Risk

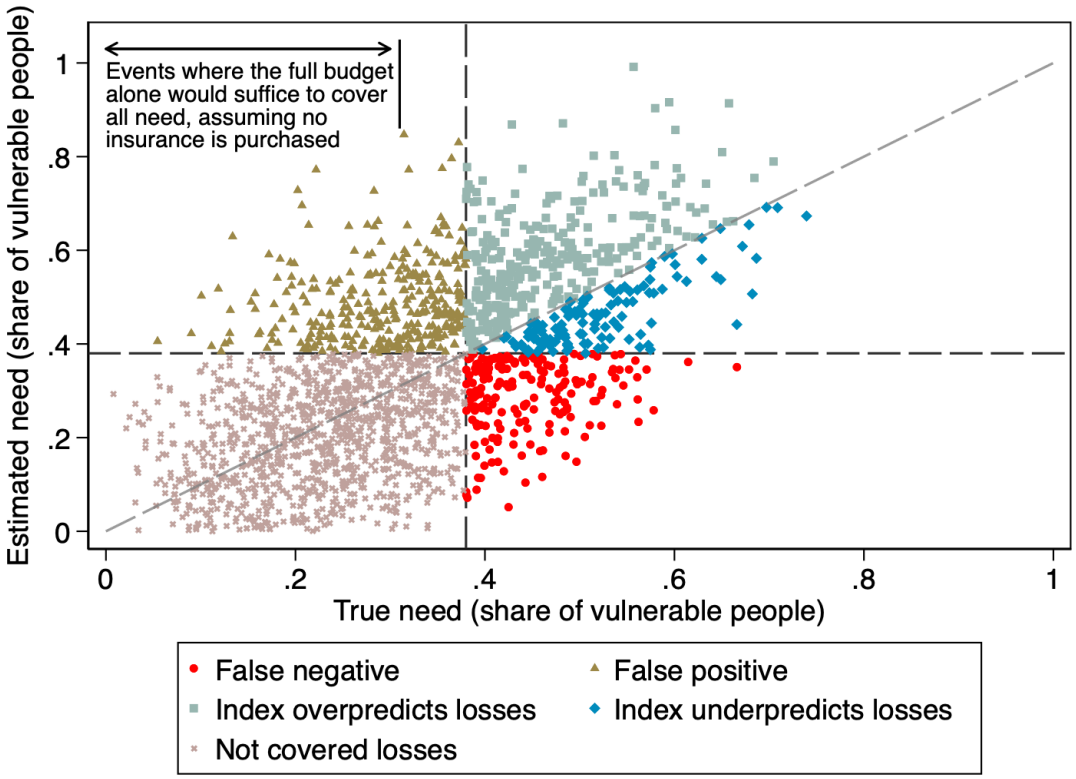


imprecision term $\omega_{t,M}$ to be i.i.d. with $\mu_{2,M} = 0$ and $\sigma_{2,M} = 800,000$.[19] The two panels in Figure 2 illustrate the implications of these assumptions. Focusing for now only on the base case scenario, we see that the probability density function (Figure 2a) for the estimated number of people (the brown curve) has thicker left and right tails than the true distribution of need (the solid black curve). The cumulative distribution functions in Figure 2b more fully draw out the implications of this additive prediction error specification. As shown, the index is about 10 percentage points more likely to overpredict the frequency of severe events in which 60% or more of the vulnerable population is hit by a drought. Similarly, it overestimates the frequency of small shocks—those affecting fewer than 20% of the vulnerable population—predicting a 25% likelihood compared to the true 20%. Both types of prediction error reduce the value of insurance. While overprediction leads to oversupply of funds, which are costly to the government, underprediction leads to undersupply, which prevents affected people from covering their drought-induced needs.

Figure 3 provides a clearer depiction of how basis risk impacts the reliability of sovereign parametric insurance. Drawing on the graphical device used by Jensen et al. (2016), the figure plots actual versus predicted need using the medium basis risk parameters from the base case

[19]This assumption is empirically grounded in observed discrepancies between predicted and actual social protection need, based on the hypothetical ARC contract and NDMA data.

Figure 3: Prediction Errors in True Versus Estimated Need



scenario. The average share of the population that is hit by a drought in any given period is 31%. As indicated in the upper left quadrant of the figure, if the government purchased no insurance, it would be able to fully meet shock-response needs only in periods when fewer than 31% of the population is affected. As the figure shows, many realizations fall to the left of this vertical threshold, meaning that in a large number of states of the world, setting aside the entire budget as a reserve fund would be sufficient to cover shock-response need. Under our distributional assumptions and simulation procedure, the share of realizations in which need exceeds the reserve fund's capacity defines a 48% uncovered risk layer of the most extreme events. These more severe outcomes highlight the critical role of sovereign insurance in stabilizing well-being by providing budgetary support precisely when shock-response need is greatest.

For illustrative purposes, we assume a sovereign contract with a 38% deductible ($\delta = 0.38$), meaning that insurance only issues payments when it is estimated that at least 38% of the vulnerable population has been hit by a drought. The horizontal dashed line marks this threshold for the insurance index, while the vertical dashed line divides the space into regions where actual needs exceed this level. The 45-degree line represents points where actual and predicted need coincide. Points below the 45-degree line indicate need under-prediction, while points above reflect need over-prediction.

Using the base case data generation process, we created a random sample of actual and predicted needs, which are plotted in Figure 3. The brown triangles in the northwest quadrant represent false positive events in which the contract would issue some payments despite true need being below the 38% deductible threshold. The red dots in the southeast quadrant mark false negative events, meaning no insurance payment would be made despite true need reaching as high as 50–60% of the population. The square markers in the northeast quadrant indicate cases where both estimated and true need exceed the 38% threshold, but the amounts differ, resulting in either under-compensation (blue squares) or over-compensation (teal squares).

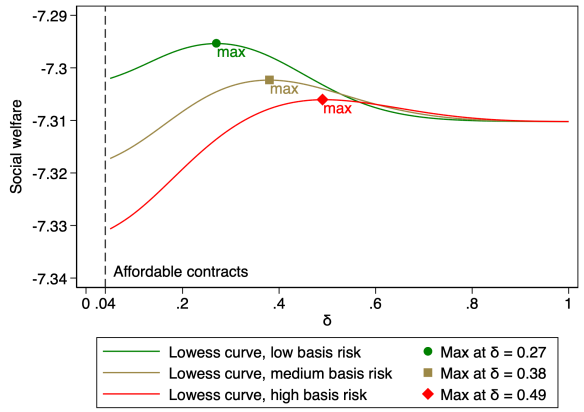
While Figure 3 gives a visual sense of the importance of basis risk, it does not, on its own, allow us to infer what level of insurance coverage should be optimally purchased by a government seeking to maximize the well-being of drought-affected households. The next section addresses this question directly.

4.2 Optimal Base Case Sovereign Insurance Under an Actuarially Fair Premium

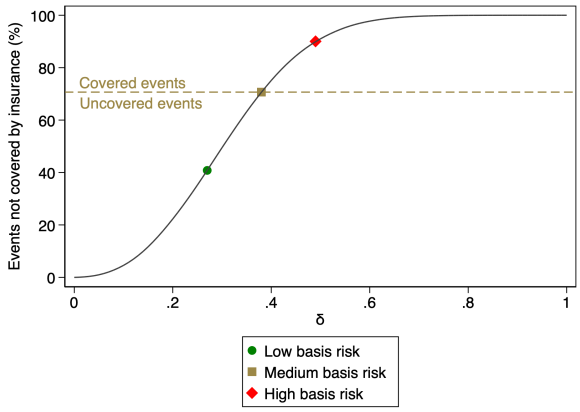
We begin our analysis by assuming that governments can purchase sovereign index insurance at its actuarially fair price. While this scenario is not realistic, it allows us to focus first on

Figure 4: Optimal Deductible and Coverage Levels Under No Markup

(a) Social Welfare Maximizing Deductible



(b) Optimal Coverage Level



the impact of basis risk on the optimal demand for insurance. For this analysis, we use the base case data generation process and simulate 500 histories, each spanning 200 years. As further detailed in Appendix B, we optimize Problem 2 by performing a grid search over a sequence of values for δ ranging from zero to one.

Figure 4a displays the result of this exercise. The brown curve shows how the value of the objective function changes as the deductible increases. Note that, given the government's budget constraint, the lowest deductible contract that it can afford corresponds to $\delta = 4\%$.[20] As shown in the figure, the social-welfare maximizing level of insurance under these parameters is achieved when the contract features a 38% deductible. Again, this deductible reflects the percentage of the largest possible loss—namely, all 5 million vulnerable households being hit by a drought.

Figure 4b provides an alternative way of visualizing the optimal level of insurance coverage in the base case scenario. An optimal deductible of 38% implies that the government relies on its reserve fund to cover about 71% of events—those less severe—and purchases insurance

[20]This affordability bound, as depicted in the figure, is actually specific to the high-basis-risk case. The affordability threshold varies with basis risk, since higher basis risk leads to greater variability in payouts and therefore higher premiums for the same deductible value.

only for the remaining 29% risk layer associated with the most extreme outcomes. This result can be benchmarked against the application in Clarke et al. (2017), where insurance is used to cover only the most extreme 7% of events, based on a 50% markup and the assumption of complete alignment between payouts and need. In contrast, our base case assumes no markup, making insurance more affordable. Crucially, our framework also incorporates basis risk, which affects not only the cost of insurance but also its effectiveness in matching true need. As a result, the optimal risk layer in our framework is endogenously determined by both price and predictive accuracy. In our setting, greater affordability and moderate basis risk lead to broader insurance coverage than that recommended by Clarke et al. (2017). This difference underscores the importance of evaluating sovereign parametric insurance not only by its cost but also by how well it aligns with true social protection objectives. In Section 5, we return to this comparison under the assumption of positive markups and examine how insurance coverage recommendations shift when variations in both premiums and index reliability are jointly considered.

To gauge the impact of sovereign index insurance on the well-being of drought-affected households, we calculate the certainty equivalent (CE) of income of an affected person across the full set of 500 simulated histories.[21] This metric translates the objective function into an interpretable economic welfare metric. Table 2 reports the CE income of affected people under various policy regimes. By construction, when the government provides no social protection, the CE income of the affected is \$1.40. If the government pursues a conventional social protection policy where it uses its budget $\bar{T}$ to close the shock-response need gap as best it can without using insurance, the CE rises to \$1.81. The optimal parametric insurance contract modestly improves welfare further, increasing the CE income of the affected to \$1.82. Finally, the table shows that when the government deviates from the optimal contract—

[21] Using the utility function in Table 1, the CE is calculated as the inverse of expected utility, that is, $CE = u^{-1}(\frac{\sum_{h=1}^{500}\sum_{t=1}^{200} p_{th} \times u(y_{th})}{\sum_{h=1}^{500}\sum_{t=1}^{200} p_{th}})$, where p_{th} represents the number of affected people in year t of history h and y_{th} is the income level they receive after any social protection payments.

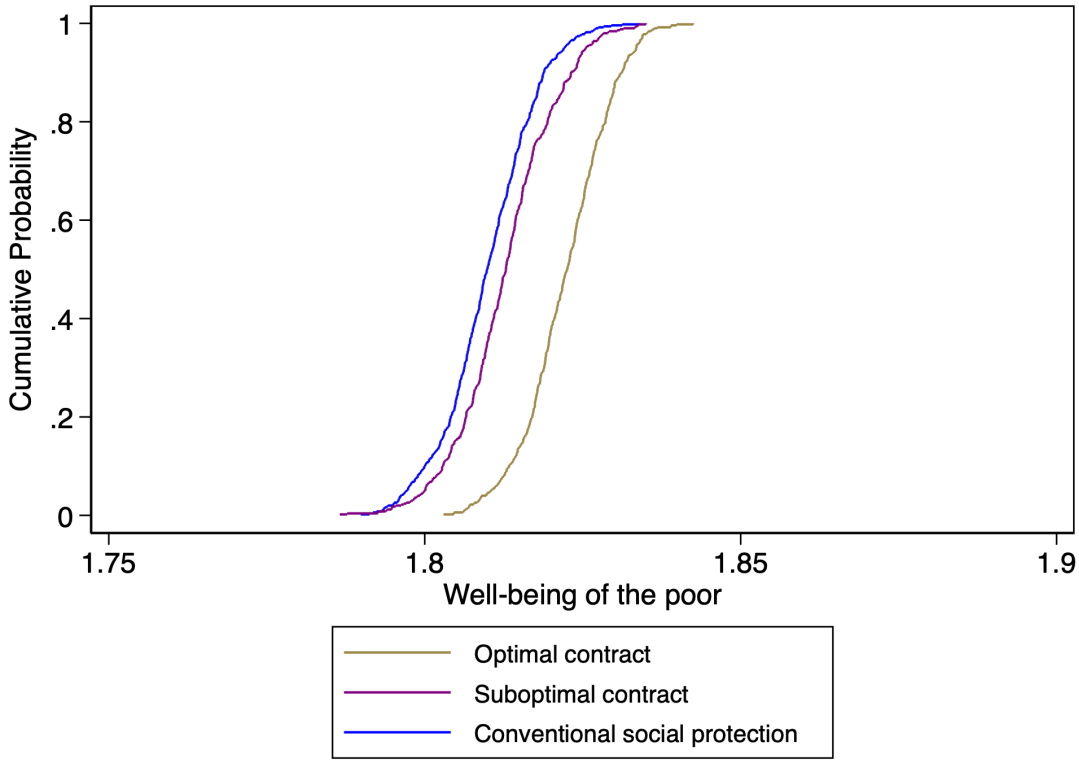
Table 2: Insurance Performance Under the Base Case (No Markup)

	δ (%)	**Risk Layer Covered** (%)	**Certainty Equivalent Income of Affected (USD)**	**Relative Benefit of Sovereign Index Insurance (RBS)** (%)
No Social Protection	–	–	1.40	–
Conventional Social Protection	100	0	1.81	–
Optimal Parametric Insurance	38	27	1.82	14
Under-insurance	60	2	1.81	6
Over-insurance	20	78	1.81	3
Perfect Insurance	0	100	1.90	100

either by under-insuring (e.g., using a 60% deductible) or over-insuring (e.g., choosing a 20% deductible)—the CE drops back to \$1.81. This result inducates that both forms of suboptimal insurance yield no discernable gains over the conventional social protection regime.

Figure 5 presents the cumulative distribution functions of CE income across all 500 simulated histories. As shown, the optimal contract (brown line) consistently delivers the highest CE income, with outcomes never falling below \$1.80. In contrast, the suboptimal contract (purple line), which corresponds to the over-insurance scenario in Table 2, and the conventional social protection regime (blue line) both lead to more frequent low-welfare outcomes. In these cases, CE falls below \$1.80 in approximately 5–10% of histories. Notably, in the over-insurance scenario, the worst-case outcomes are slightly worse than under the no-insurance regime. This finding suggests that purchasing excessive insurance in the presence of basis risk can result in an inefficient use of limited funds. As we discuss in Section 5, these risks become even more pronounced when contracts are priced with a markup or when basis risk increases.

Figure 5: Empirical CDFs of Well-Being of the Affected by Droughts Across Histories, No Markup and Medium Basis Risk



Note: The optimal contract corresponds to a 38% deductible under the base case. The suboptimal contract featured here reflects an over-insurance scenario in which a 20% deductible is chosen.

5 Stress-Testing Optimal Sovereign Parametric Insurance with Variable Basis Risk and Markups

In this section, we consider alternative scenarios to the base case by varying the level of basis risk and the insurance markup rate. As shown in Table 1, we specify the imprecision term for the low-basis-risk case ($\omega_{t,L}$) and for the high-basis risk case ($\omega_{t,H}$) to be i.i.d. with $\mu_{2,L} = 0, \sigma_{2,L} = 600,000$ and $\mu_{2,H} = 0, \sigma_{2,H} = 1,000,000$, respectively. Both follow the same additive functional form used in the base case. For markups, we evaluate results under 25% ($m = 0.25$) and 50% ($m = 0.5$) rates. We analyze how these changes affect need prediction,

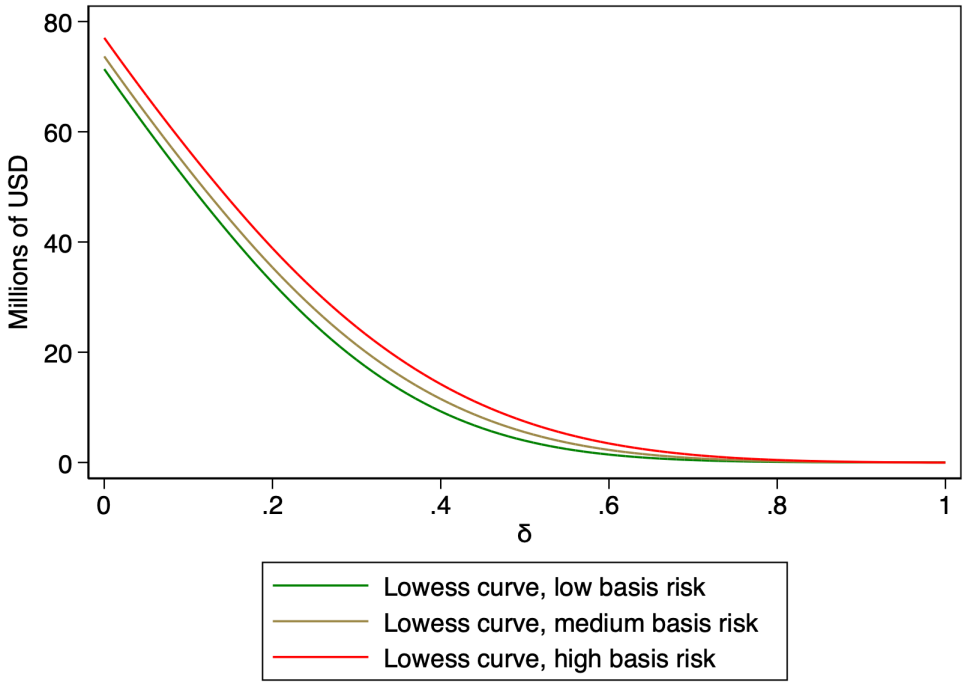
contract premiums, optimal contract choice, and coverage levels under the optimal contract relative to the base case.

Figure 2 above illustrates how estimated need under different levels of basis risk deviates from the true social protection need. Panel (a) shows the probability density functions, while Panel (b) displays the corresponding cumulative distribution functions. In the low-basis-risk case (green lines), the estimated need closely follows the true need (black lines), with a narrow distribution centered around the actual values. As basis risk increases, moving from medium (brown lines) to high levels (red lines), the estimated need distribution becomes progressively more dispersed. This greater dispersion implies that the index misclassifies events more frequently, leading to both underestimation and overestimation of need. In panel (b), we see that under high basis risk, the cumulative distribution rises more gradually than the true distribution, indicating an inflated frequency of small-shock predictions and an underrepresentation of moderate to severe events. In summary, as previously discussed, higher basis risk reduces the reliability of index insurance by increasing both over-predictions and under-predictions, ultimately weakening the ability of insurance to provide support when it is most needed.

For each possible value of the deductible δ, we calculate the insurance contract premium by setting it equal to the expected indemnity payments, evaluated separately for the three different basis risk levels. Figure 6 plots these premiums under the assumption of no markup ($m = 0$). The x-axis represents the deductible, which determines the threshold above which the contract begins to pay out. At low values of δ, the contract covers a broader set of events, resulting in higher expected indemnity payments and, consequently, higher premiums. As δ increases, payouts occur less frequently (only when estimated need exceeds a higher threshold) so expected indemnity payments fall, and premiums decline accordingly.

The level of basis risk also affects premiums. In the low-basis-risk case, where estimated need closely tracks true need, payouts are relatively predictable and aligned with

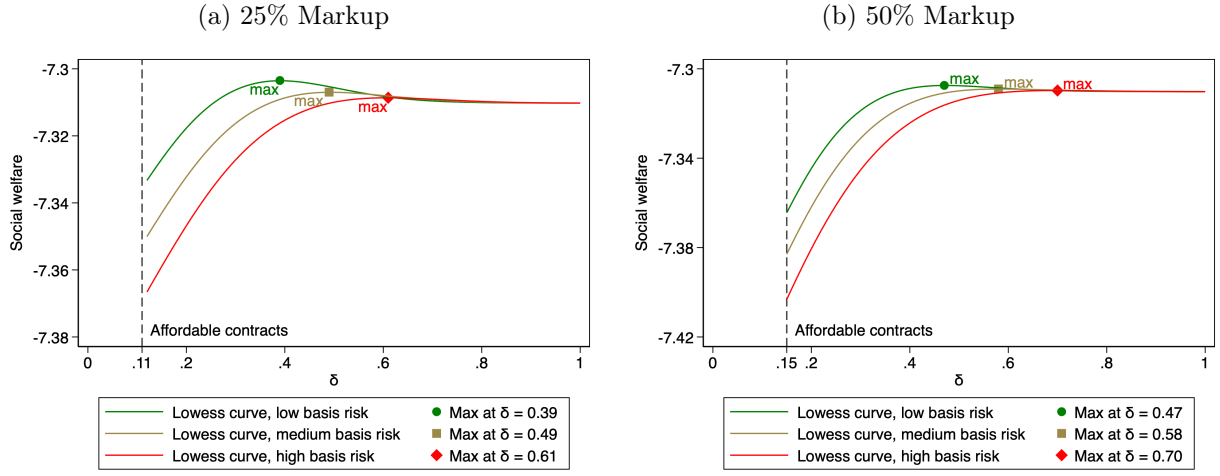
Figure 6: Insurance Premiums Under No Markup



realized events, resulting in lower premium costs. As basis risk increases, prediction errors become larger and more frequent, raising the likelihood of both under-compensation and over-compensation. In our framework, indemnity payments are a truncated, convex function of the index, so that increases in the variance of index error raise expected payouts even when prediction errors are mean zero. As a result, even without a markup, this greater uncertainty raises the expected indemnity payments and leads to higher actuarily fair premiums, as seen in the upward shift of the curves from green to brown and red.

For every combination of basis risk level and markup rate, we calculate the expected social welfare across all deductible values and identify the optimal contract as the one that maximizes this objective function. Figures 7a and 7b complement Figure 4a above plot social welfare as a function of δ under higher markup rates. Figure 4a displays results for a 0% markup, 25% and 50% markup cases are presented in Figures 7a and 7b, respectively. In each

Figure 7: Social Welfare Maximization and Optimal Deductible Choices Under Alternative Markups



case, the deductible that maximizes social welfare is highlighted for each basis risk level.[22]

Table 3 summarizes the results for each scenario, including the optimal deductible δ, the associated risk layer covered, the CE income of the affected, and the RBS measure we introduced in Section 3.2. In each case, we also report outcomes under a perfect insurance contract (one that operates without basis risk) as a benchmark for evaluating the value and limitations of more realistic index insurance contract options.[23]

Perfect insurance, by construction, eliminates basis risk and guarantees full coverage of shock-response need whenever it arises, specifically under no markup. As shown in the table, in this case, this benchmark contract consistently produces a CE income of \$1.90. However, under a positive markup rate, full coverage becomes unaffordable, and the chosen deductible under perfect insurance rises up to 33%. Notably, even with perfect insurance, the welfare

[22]Note that not all contracts are feasible given the government's fixed budget. Deductibles below a certain threshold fall outside the affordability range, and this threshold increases as the markup rises. This budget constraint therefore limits the government's ability to insure against more frequent, moderate shocks, reinforcing the tendency to choose high-deductible, low-coverage contracts.

[23]The perfect insurance benchmark should be interpreted as an upper bound on attainable welfare gains rather than a feasible policy option. It provides a useful reference point for assessing how rapidly the value of insurance erodes as prediction errors and pricing wedges are introduced.

Table 3: Optimal Deductibles and Welfare Outcomes for Sovereign Index Insurance Across Basis Risk and Markup Scenarios

	Low Basis Risk	**Medium Basis Risk**	**High Basis Risk**
No Markup			
Perfect Insurance			
Deductible (%)	0	0	0
Risk Layer Covered (%)	100	100	100
CE Income of Affected (USD)	1.90	1.90	1.90
Gain over Conventional Social Protection (USD)	0.09	0.09	0.09
Optimal Insurance			
Deductible (%)	27	38	49
Risk Layer Covered (%)	59	29	10
CE Income of Affected (USD)	1.83	1.82	1.82
Relative Benefit of Sovereign Insurance (RBS) (%)	26	14	7
25% Markup			
Perfect Insurance			
Deductible (%)	24	24	24
Risk Layer Covered (%)	68	68	68
CE Income of Affected (USD)	1.84	1.84	1.84
Gain over Conventional Social Protection (USD)	0.03	0.03	0.03
Optimal Insurance			
Deductible (%)	39	49	61
Risk Layer Covered (%)	27	10	2
CE Income of Affected (USD)	1.82	1.81	1.81
Relative Benefit of Sovereign Insurance (RBS) (%)	32	15	8
50% Markup			
Perfect Insurance			
Deductible (%)	33	33	33
Risk Layer Covered (%)	42	42	42
CE Income of Affected (USD)	1.82	1.82	1.82
Gain over Conventional Social Protection (USD)	0.01	0.01	0.01
Optimal Insurance			
Deductible (%)	47	58	70
Risk Layer Covered (%)	12	3	0
CE Income of Affected (USD)	1.82	1.81	1.81
Relative Benefit of Sovereign Insurance (RBS) (%)	87	41	19

Note: Conventional social protection yields a CE income of the affected of $1.81.

gain over conventional social protection is modest once a markup is imposed, underscoring how cost alone can substantially limit the value of insurance, regardless of the reliability of the index. These benchmark results provide important context for evaluating the performance of more realistic, index-based contracts under less-than-ideal conditions.

As for the optimal contracts in the presence of basis risk, when insurance is priced at its actuarily fair value (i.e., no markup), the optimal δ ranges from 0.27 in the low-basis-risk case to 0.49 in the high-basis-risk case. This pattern reflects the fact that when basis risk is

high, the government is better off choosing contracts that cover only the most severe events, as these trigger payouts less frequently and reduce the likelihood of costly misclassifications. In such cases, retaining a larger share of the social protection budget as a reserve fund is preferable to committing resources to a low-deductible, high-premium contract that is poorly aligned with true need. At this pricing level, the optimal contract modestly increases CE income over conventional social protection, raising it from \$1.81 to \$1.82–\$1.83 depending on basis risk. Although this welfare gain may seem small in absolute terms, it reflects a nontrivial improvement under tight budget constraints and imperfect prediction.

Similarly, for a given basis risk level, an increase in the markup rate also leads to a higher optimal deductible. For instance, under medium basis risk, the optimal δ rises from 0.38 with no markup to 0.49 with a 25% markup and to 0.58 with a 50% markup. As higher markups make insurance more expensive, the government shifts toward higher-deductible contracts to reduce its spending on premiums, even though this choice results in fewer payouts.

Comparing social welfare levels across markup scenarios, we find that under no markup, social welfare is maximized at moderate deductible levels. When the markup rises to 25%, peak social welfare occurs at a higher δ across all basis risk levels, and social welfare declines overall relative to the no-markup case. This pattern intensifies under a 50% markup, as optimal deductibles move further to the right and social welfare drops more steeply, reflecting how higher premiums reduce the efficiency and value of the insurance contract.

Now we turn to the levels of coverage associated with the optimal contract deductibles. Note that the optimal deductible δ^* corresponds to a threshold number of people $p^*(\delta^*) = \delta^* N$, such that payouts happen only when the number of people affected by a drought exceeds this threshold. Recall that the number of people in need follows a probability density function $\varphi_p(p)$. Given its associated cumulative distribution function $\Phi_p(p)$, the term $\Phi_p\left(p^*(\delta^*)\right)$ represents the risk layer of events not covered by insurance—that is, the share of least severe events for which insurance does not yet kick in under the optimal contract. Conversely, the

expression $1 - \Phi_1\left(p^*(\delta^*)\right)$ describes the risk layer that is covered by the optimal contract. This metric captures how restrictive or generous the contract is in triggering payouts. As shown in Section 4.2, this covered risk layer is approximately 29% in the base case.

As shown in the table, the risk layer covered declines steeply as either the markup or basis risk increases. Under no markup, the optimal contract covers 59% of events in the low-basis-risk case, 29% in the medium-basis-risk case, and just 10% in the high-basis-risk case. Similarly, for a given level of basis risk, higher markup rates lead to reduced coverage. For example, under medium basis risk, the risk layer covered falls from 29% of events with no markup to 10% with a 25% markup and only 3% with a 50% markup. In the extreme case (50% markup and high basis risk), the optimal contract effectively covers 0% of events, implying that insurance becomes nearly irrelevant in that scenario. This result stands in contrast to Clarke et al. (2017), who recommend insuring the most extreme 7% of events under a 50% markup, regardless of index reliability. In our framework, by contrast, optimal insurance choice under the same markup rate varies with basis risk. A key result is that under high basis risk, we find that purchasing insurance delivers virtually no value, and the optimal policy is to allocate the entire social protection budget to a reserve fund. Overall, these results illustrate how sovereign parametric insurance can quickly become hollowed out when it becomes too inaccurate or too expensive.

Because of the skewed distribution of drought events, even a small increase in the deductible—reflecting the government's optimal adaptation to either increased basis risk or a higher markup—can lead to a large drop in the share of events covered. In particular, when basis risk is high, the optimal contract covers only rare, extreme events, even when there is no markup. When the markup rate is high, the government also optimally chooses to insure only infrequent events, regardless of basis risk. These findings highlight how higher basis risk and markup rates can substantially reduce the value that index insurance offers to governments. The implication is that, depending on these factors, sovereign parametric insurance

may be best suited as a targeted tool for managing tail risks, rather than as a broad-based social protection instrument.

Finally, the RBS offers a normalized measure of the welfare gain achieved by the optimal contract relative to perfect insurance.[24] While one might expect RBS to fall as insurance becomes more costly or inaccurate, we observe that RBS actually increases with markup. Under no markup, RBS ranges from 26% (low basis risk) to just 7% (high basis risk). However, at a 50% markup, RBS rises to 87%, 41%, and 19% across low, medium, and high basis risk, respectively. This pattern emerges not because the optimal contracts perform substantially better, but because the benchmark deteriorates sharply as markup increases. As perfect insurance becomes more expensive and requires higher deductibles to remain affordable, its superiority over realistic contracts diminishes. In effect, the rising RBS values under high markup rates reflect a shrinking gap between the best-case and feasible scenarios.

6 Conclusion

There is a compelling case for prearranging financing to address the excess social protection needs that emerge in the wake of droughts, hurricanes, and other natural disasters. In response, a number of multilateral facilities have emerged offering parametric insurance contracts that provide insured governments with budgetary support when an underlying index of wind speeds, drought, or similar indicators, signals that increased social protection expenditures are required. These contracts are built on principles similar to those used to offer parametric or index insurance to individual farmers, businesses, and financial institutions.

Like their micro-level counterparts, sovereign parametric insurance contracts face the

[24] Although we apply the RBS measure to a single contract design, it is intended as a general tool for comparing alternative sovereign insurance contracts, including differences in the underlying index, deductible or coverage structure, and premium pricing. The framework can be readily extended to assess the value of improved index accuracy or modifications to contract design.

problem of so-called "basis risk events," in which losses occur but are inaccurately predicted by the underlying index. While the contract inaccuracy problem has received significant attention in the microinsurance space, this study is the first we know to analyze the issue for sovereign parametric insurance. Our optimization model demonstrates that the optimal amount of parametric insurance a government should buy is highly sensitive to the accuracy of the index—that is, to the degree of basis risk. Using realistic assumptions about basis risk derived from a sovereign parametric contract developed for arid regions of Kenya, we find that the optimal amount of insurance, while nonzero, is lower than is typically presumed.

To illustrate, under medium basis risk and zero markup, we find that the government should insure only the most extreme 29% of events. This result contrasts with Clarke et al. (2017), who recommend covering only the most extreme 7% of events under a 50% markup, assuming no basis risk. In our model, however, applying that same markup level in the presence of high basis risk leads to an optimal insurance coverage near zero. This comparison highlights the importance of jointly considering both contract cost and predictive reliability when evaluating the usefulness of sovereign parametric insurance. Taken together, our results suggest that sovereign insurance is best understood as a financing instrument whose value depends critically on index accuracy and pricing, rather than as a substitute for targeting or program design. By abstracting from heterogeneity and delivery frictions, we provide a conservative benchmark for insurance performance.

This finding suggests that the recent advances in remote sensing (Benami et al., 2021) and machine learning tools (Estefania-Salazar et al., 2025), which have been applied to micro-level index insurance, should also be used to improve the predictive accuracy of sovereign parametric insurance indexes. Our results further underscore the need for governments to adopt a rigorous economic approach when evaluating how to integrate parametric insurance products into their adaptive social protection schemes.

References

African Risk Capacity, "Africa RiskView Methodology Technical Note: Drought," Technical Report, African Risk Capacity 2016.

Arrow, Kenneth J and Robert C Lind, "Uncertainty and the evaluation of public investment decisions," in "Uncertainty in economics," Elsevier, 1978, pp. 403–421.

Beazley, Rodolfo, Ana Solórzano, and Valentina Barca, "Study on Shock Responsive Social Protection in Latin America and the Caribbean: Summary of Key Findings and Policy Recommendations.," 2019.

Benami, Elinor, Ella Kirchner, Andrew Hobbs, Michael R. Carter, and Zhenong Jin, "Get in the Zone: The Risk-Adjusted Welfare Effects of Data-Driven vs. Administrative Borders for Index Insurance Zones," 2025.

_ , Zhenong Jin, Michael R. Carter, Aniruddha Ghosh, Robert J. Hijmans, Andrew Hobbs, Benson Kenduiywo, and David B. Lobell, "Uniting remote sensing, crop modelling and economics for agricultural risk management," *Nature Reviews Earth & Environment*, February 2021, *2* (2), 140–159.

Bjerde, Anna and Ed Mountfield, "Unlocking new crisis response tools to build a more resilient future," 2024.

Carter, Michael, Alain de Janvry, Elisabeth Sadoulet, and Alexandros Sarris, "Index Insurance for Developing Country Agriculture: A Reassessment," *Annual Review of Resource Economics*, 2017, *9* (1), 421–438.

Carter, M.R. and T Chiu, "Quality standards for agricultural index insurance: An agenda for action," *State of Microinsurance*, 2018, (4).

Chamberlain, James and Nick Bernards, "Insurance and the contradictions of the climate-development-finance nexus: The case of the African Risk Capacity," *Competition & Change*, 2024, *0* (0), 10245294241226985.

Chantarat, Sommarat, Andrew G Mude, Christopher B Barrett, and Michael R Carter, "Designing index-based livestock insurance for managing asset risk in northern Kenya," *Journal of Risk and Insurance*, 2013, *80* (1), 205–237.

Clarke, Daniel J., "A Theory of Rational Demand for Index Insurance," *American Economic Journal: Microeconomics*, February 2016, *8* (1), 283–306.

Clarke, Daniel J and Stefan Dercon, *Dull Disasters? How planning ahead will make a difference*, Oxford University Press, 2016.

_ , Olivier Mahul, Richard Poulter, and Tse-Ling Teh, "Evaluating sovereign disaster risk finance strategies: a framework," *The Geneva Papers on Risk and Insurance-Issues and Practice*, 2017, *42*, 565–584.

Clarke, Daniel, Olivier Mahul, Kolli Rao, and Niraj Verma, "Weather Based Crop Insurance in India," Policy Research Working Paper 5985, World Bank, Washington D.C. 2012.

Estefania-Salazar, Enrique, Michael R. Carter, Eva Iglesias, and Alvaro Escribano, "The Economic Gains from Breaking Boundaries: Reformulating Index Insurance Zones to Maximize Lower Tail Dependence," *NBER Working Paper*, 2025, (w32618).

Ghesquiere, Francis, *Sovereign natural disaster insurance for developing countries: A paradigm shift in catastrophe risk financing*, Vol. 4345, World Bank Publications, 2007.

Jensen, Nathaniel and Christopher Barrett, "Agricultural Index Insurance for Development," *Applied Economic Perspectives and Policy*, 2017, *39* (2), 199–219.

Jensen, Nathaniel D., Christopher B. Barrett, and Andrew G. Mude, "Index Insurance Quality and Basis Risk: Evidence from Northern Kenya," *American Journal of Agricultural Economics*, 2016, *98* (5), 1450–1469.

Jensen, Nathaniel D, Francesco P Fava, Andrew G Mude, Christopher B Barrett, Brenda Wandera-Gache, Anton Vrieling, Masresha Taye, Kazushi Takahashi, Felix Lung, Munenobu Ikegami et al., *Escaping poverty traps and unlocking prosperity in the face of climate risk: Lessons from index-based livestock insurance*, Cambridge University Press, 2024.

Johnson, Leigh, "Rescaling index insurance for climate and development in Africa," *Economy and Society*, 2021, *50* (2), 248–274.

Kenduiywo, Benson K., Michael R. Carter, Aniruddha Ghosh, and Robert J. Hijmans, "Evaluating the quality of remote sensing products for agricultural index insurance," *PLOS ONE*, October 2021, *16* (10), e0258215.

Miranda, Mario J. and Katie Farrin, "Index Insurance for Developing Countries," *Applied Economic Perspectives and Policy*, 2012, *34* (3), 391–427.

Museru, Malimu, Francois Toerien, and Sean Gossel, "The Impact of Aid and Public Investment Volatility on Economic Growth in Sub-Saharan Africa," *World Development*, 2014, *57*, 138–147.

National Drought Management Authority, "The Hunger Safety Net Programme Consolidated Operations Manual," Technical Report, National Drought Management Authority 2024.

Appendices For Online Publication

Appendix A Minimizing the Cost of Meeting a Fixed Social Protection Obligation

This appendix presents a dual formulation of the problem analyzed in the main text. Unlike the fixed-budget setup, here the government commits to covering the entirety of shock-response social protection need in each period. This formulation is more aligned in spirit with Clarke et al. (2017), as it captures the fiscal trade-off between meeting social protection needs and preserving resources for public investment. The government relies on a three-layer financing structure: *ex ante* budget allocation (reserve fund), sovereign parametric insurance, and emergency *ex post* budget reallocation. The reserve fund is fed by a fixed endowment each year, equal to the expected level of social protection need. From this endowment, the government purchases an insurance contract, subject to a deductible δ, as defined in the main text. The leftover funds after paying the premium constitute the reserve fund. In periods when social protection need exceeds the combined resources from the reserve fund and indemnity payments, the government covers the shortfall by diverting resources away from its public investment budget through emergency *ex post* reallocation. Conversely, when realized social protection need is smaller than the available resources, the excess is funneled back into public investment.

This framework introduces a trade-off already highlighted by Clarke et al. (2017): reallocating funds from public investment toward social protection imposes macroeconomic costs. Specifically, variation in public investment caused by emergency budget reallocations affects GDP growth through two channels. First, there is a level effect, since reduced investment reduces output. Second, there is a volatility effect, as unpredictable public investment flows

undermine long-term growth. We quantify both effects using the framework and estimates by Museru et al. (2014), who study the impact of public investment volatility on growth using data from 26 Sub-Saharan African countries. Importantly, we assume that funds drawn from the pre-arranged social protection budget do not trigger these GDP penalties, since they are planned in advance and introduce no uncertainty.

Given this three-layer financing structure, the government's optimization problem becomes one of choosing the insurance contract—that is, the deductible δ—that minimizes expected GDP loss. This formulation shifts the focus from stabilizing household well-being, as in the main text, to minimizing the macroeconomic cost of fully meeting social protection obligations, thereby making explicit the public finance externalities of disaster response.

In practice, this exercise requires simulating the effects of insurance and reallocations on public investment. To calculate the objective function of GDP losses due to budget diversions toward social protection, we first estimate a predicted public investment curve to smooth variation in the real data. Disaster risk financing leads to new values of public investment, where in period t,

$$\text{Public investment after reallocation}_t = \text{Predicted public investment (before reallocation)}_t$$

$$+\overline{T} - \text{Insurance premium} + I_t(\delta) - T_t,$$

where $\overline{T}$ is the social protection endowment, $I_t(\delta)$ is the indemnity payment, and T_t is the realized drought-related social protection need, all expressed in yearly terms. The first two layers of disaster risk financing are captured by adding the reserve fund endowment and subtracting the insurance premium from predicted investment. The emergency reallocation layer is reflected by further adding the difference between indemnity payments and realized

social protection need. This equation summarizes the mechanism described above, where the government reallocates resources into or out of public investment depending on how social protection need compares to the reserve fund (net of the premium) and the insurance payouts.

Once we calculate the public investment figures after reallocation for each year, we apply the framework of Museru et al. (2014) to measure GDP losses (or gains) by comparing actual public investment after reallocation with the counterfactual public investment without reallocations. Finally, we aggregate these GDP losses over the 200 years of each simulated history. The optimal contract the government can choose corresponds to the deductible δ that minimizes expected GDP loss.

As in the fixed-budget case, we evaluate the performance of different possible insurance contracts by simulating 5,000 histories of 200 years each, using parameters based on NDMA data, the hypothetical ARC contract, and macroeconomic data from Kenya for the 2004–2019 period. Table A1 displays the simulation results under the same basis risk and markup scenarios considered in the main text. For the 0% markup case, the results closely mirror those from the fixed-budget setting: the optimal deductible rises with basis risk, from 28% at low basis risk to 52% at high basis risk. This pattern reflects the same mechanism as before—when basis risk is high, it is better to insure only the tail of the distribution where need is most extreme, while relying on reserves to tackle more moderate shocks.

Where the two frameworks diverge most clearly is in the sensitivity of the optimal insurance choice to markup. In the fixed-budget case, higher markups directly crowd out funds available for direct transfers, forcing the government to shift sharply toward higher deductibles to preserve resources. In this setup, however, all need must be met regardless, with any shortfall covered through *ex post* reallocations from public investment. Because premiums do not reduce the resources available for transfers, markups affect the problem only through their contribution to cost smoothing. As a result, optimal deductibles increase only modestly with markup. For example, they rise from 38% to 42% at medium basis risk

Table A1: Optimal Deductibles for Sovereign Index Insurance Across Basis Risk and Markup Scenarios

	Low Basis Risk	**Medium Basis Risk**	**High Basis Risk**
No Markup, Optimal Insurance			
Deductible (%)	28	38	52
Risk Layer Covered (%)	56	29	7
25% Markup, Optimal Insurance			
Deductible (%)	31	40	54
Risk Layer Covered (%)	48	25	5
50% Markup, Optimal Insurance			
Deductible (%)	33	42	55
Risk Layer Covered (%)	42	21	5

when moving from 0% to 50% markup. Insurance thus remains valuable for reducing the level and volatility of costly emergency reallocations, even when its price rises. In this sense, our results connect back to those of Clarke et al. (2017), who recommend insuring only the most extreme 7% of events under a 50% markup. Under our current formulation, when basis risk is high and markup is 50%, the optimal insured risk layer similarly amounts to 5%, bringing our findings closer to theirs. Crucially, however, our results account for the problem of contract inaccuracy.

The muted response of the optimal contract to markup highlights the key difference between the two setups: in the fixed-budget case, insurance competes with direct transfers, while in this case, it primarily functions as a fiscal smoothing device. This distinction underscores how the policy implications of sovereign parametric insurance depend critically on whether governments operate under rigid social protection budgets or can reallocate flexibly across spending categories.

Appendix B Numerical Methods for Solving Optimization Problem

This appendix details the simulation procedure we used to evaluate the optimality of alternative sovereign index insurance contracts under varying levels of basis risk and markup. We also explain and justify the distributional assumptions underpinning the stochastic modeling of drought-related social protection need.

To simulate the number of individuals requiring social protection due to drought, we model the true number of affected individuals, p_t, using a Weibull distribution calibrated to historical data from Kenya's National Drought Management Authority (NDMA), which includes biannual estimates of drought-affected populations between 2004 and 2020. This empirical distribution has a mean of approximately 1.55 million and a standard deviation of about 1.11 million.

We chose to perform the simulations using a Weibull distribution for three main reasons. First, its support is positive. Unlike a normal distribution, the Weibull distribution naturally excludes negative values, aligning with the physical constraint that population counts cannot be negative. Second, it is flexible enough to capture the skewness observed in the historical data. Drought impact data are right-skewed, with occasional high-need outliers. The Weibull distribution can accommodate this skew and heavier right tails more effectively than a truncated normal. Third, it is straightforward to calibrate. We selected Weibull parameters (shape = 2.4727, scale = 0.3497) and scaled the resulting draws by a factor of $N = 5,000,000$ to match the empirical first and second moments of the NDMA data. Specifically, we generated 100,000 realizations of $p_t \sim \text{Weibull}(k = 2.4727, \lambda = 0.3497)$, scaled such that $E[p_t] = 1.55$ million and $\text{SD}(p_t) = 1.11$ million.

The performance of the sovereign index insurance contract depends on how accurately the

index approximates actual drought-related social protection need. We model the estimated need under the contract as $\tilde{p}_t = p_t + \omega_t$, where $\omega_t \sim \mathcal{N}(0, \sigma^2)$ is an additive error term capturing basis risk. The variance of ω_t determines the level of basis risk. We set $\sigma = 600,000$ for low basis risk, $\sigma = 800,000$ for medium basis risk, and $\sigma = 1,000,000$ for high basis risk. This additive formulation allows for both under-estimations and over-estimations of true need, generating false positives and false negatives. We chose the additive specification for its tractability and its symmetric treatment of misclassification. To ensure non-negative estimates of $\tilde{p}_t$, we truncate values outside the admissible range $[0, N]$, where $N = 5,000,000$ represents the total number of vulnerable individuals. We performed the truncation by iteratively redrawing predicted need values that fell outside the bounds until all estimates were within range.

We simulate 500 independent contract "histories," each consisting of 200 periods. For each simulation, we generate a draw of true need p_t and compute the corresponding index-based estimate $\tilde{p}_t$. We then evaluate the performance of insurance contracts over a grid of deductible levels δ, which define the minimum threshold at which the index triggers a payout. For each δ, and for each scenario defined by a combination of basis risk level $\sigma \in \{600,000,\ 800,000,\ 1,000,000\}$ and insurance markup rate $m \in \{0,\ 0.25,\ 0.5\}$, we compute the resulting insurance payout, social protection payments delivered, premium cost to the government, utility of the affected, and the share of the risk layer not covered.

For each combination of σ, m, and δ, we calculate the average expected utility of the affected across the 500 simulated histories. The optimal contract is defined as the deductible level δ^* that maximizes expected social welfare for the affected, conditional on the assumed basis risk level and markup rate. This exhaustive grid search approach allows us to trace how the structure of optimal contract responds to varying economic conditions—specifically, how the trade-off between cost-efficiency (captured by the markup) and accuracy (captured by basis risk) shapes the design of sovereign index insurance instruments. By comparing

welfare frontiers across scenarios, we examine not only which contracts are optimal but also how performance deteriorates under less favorable insurance conditions.