CEIS TOR VERGATA Research Papers

Università di Roma Tor Vergata

ISSN 2610-931X

CEIS Tor Vergata

RESEARCH PAPER SERIES

Vol. 24, Issue 4, No. 629 – November 2026

How Many Vulnerability Lines Do We Need? Measuring Vulnerability as Expected Poverty

Nicola Amendola, Federico Belotti,

Domenico Moramarco, Giovanni Vecchi

How Many Vulnerability Lines Do We Need? Measuring Vulnerability as Expected Poverty*

Nicola Amendola[a], Federico Belotti[a],
Domenico Moramarco[b], Giovanni Vecchi[a]

[a] *Tor Vergata University of Rome*
[b] *University of Bari*

Abstract

A single threshold above the poverty line offers an appealing way to classify households as poor, vulnerable, or economically secure. This paper shows that, when vulnerability is defined as expected poverty, applying a common vulnerability threshold to observed consumption faces two problems. First, a vulnerability threshold may lie below the poverty line or may not exist. Second, households with the same current consumption may face very different risks of future poverty, implying household-specific vulnerability thresholds. We establish necessary and sufficient conditions for the existence of these thresholds and show that applying appropriate risk deflators to household consumption allows a common vulnerability threshold to be used on the risk-adjusted consumption scale. We then discuss three approaches to setting this common threshold and examine the consequences of leaving consumption unadjusted for risk. In our Malawi illustration for 2024–25, failing to adjust consumption for household-specific risks leads to the misclassification of 16–22 percent of the population.

Keywords: Vulnerability to poverty, Vulnerability lines, Risk-adjusted consumption, Household risk heterogeneity, Poverty measurement

JEL classification: I32, C43, D81.

*The authors are grateful to Giulia Mancini and Marco Ranzani for their useful comments on earlier versions of the manuscript. Any remaining errors are the authors' responsibility.

1. Introduction

Poverty measurement uses a monetary line to answer a static question: does a household command enough resources to meet a specified standard of living? Vulnerability analysis asks a forward-looking question: how likely is that household to be poor in the future? The distinction matters because households observed at the same consumption level may face very different earnings volatility, exposure to shocks, insurance, savings, and informal support. This forward-looking problem has generated a large literature that differs both in the welfare concept used to define vulnerability and in the data and models used to estimate it; see Gallardo (2018) for a critical survey.

A widely used approach is to define vulnerability as expected poverty. In its simplest form, household h is assigned the probability

$$v_h = \Pr(C_{h,t+1} < z), \tag{1}$$

where z is the poverty line and $C_{h,t+1}$ is uncertain future consumption over a given time horizon. The household is identified as vulnerable when $v_h \geq \tau$, for a probability threshold $\tau \in (0, 1)$. This formulation is closely associated with Chaudhuri et al. (2002) and Chaudhuri (2003), although this probabilistic approach has earlier roots in the work of Ravallion (1988), Christiaensen and Boisvert (2000), and Pritchett et al. (2000). Recent work by Dutta et al. (2025) provides an axiomatic foundation for threshold-based identification, under which households whose poverty risk is at least as high as a specified threshold are classified as vulnerable.

For policy communication, the probability rule is often translated into a monetary vulnerability line. Such a line is attractive because it can be placed above the poverty line and used to partition the observed welfare distribution into the poor, the vulnerable non-poor, and the economically secure middle class. López-Calva and Ortiz-Juárez (2014) and Dang and Lanjouw (2017), for example, use panel or synthetic panel data to estimate income or consumption levels associated with specified poverty-transition risks. Other applications have proposed a scalar line by averaging predicted welfare among households whose estimated poverty risk lies near a selected probability cutoff (Atamanov et al., 2022). These procedures are useful summaries, but they do not resolve a basic identification problem: if risks are heterogeneous, a single cutoff may classify as non-vulnerable households with high risk of poverty, and vice versa. More in general, when household face heterogeneous risks, a common cutoff

in real consumption may fail to exist even if every household's risk function is known without error.

To see the issue, consider two households with the same current consumption. One has a stable livelihood and effective insurance; the other faces volatile income and little protection. The second may have a much higher probability of future poverty. No rule based only on their common current consumption can classify one as vulnerable and the other as secure. More generally, the consumption buffer needed to keep poverty risk below τ is household-specific. This paper studies what a vulnerability line can mean in this setting. Its first contribution is to characterize the minimal requirement for the existence of a single monetary vulnerability line for each household.

The second contribution is an impossibility result. Even when every household has a well-defined individual line, a common vulnerability line in real consumption exists only if those individual lines coincide – an exceptional case rather than the general rule. With heterogeneous risks, there may be as many distinct vulnerability lines as there are households.

Third, we provide a solution to the above impossibility. By rescaling each household's real consumption to reflect the household-specific buffer required to limit poverty risk, we obtain a common vulnerability line in the space of risk-adjusted consumption. Let $z_h^v(\tau)$ be household h's vulnerability line. For any positive reference value $z^v(\tau)$, we define a household-specific risk deflator

$$RD_h(\tau) = \frac{z_h^v(\tau)}{z^v(\tau)}.$$

The probability rule $v_h \geq \tau$ is then exactly equivalent to comparing c_h/RD_h with the common line z^v. A single line therefore exists in the space of risk-adjusted consumption, not generally in the original consumption space. This construction resembles the use of true-cost-of-living indices to map heterogeneous household poverty lines into a common real-welfare metric (Amendola et al., 2024a). Here the additional deflator captures heterogeneity in the consumption buffer required to limit future poverty risk.

The common vulnerability line is not uniquely determined by the classification rule: different reference values yield the same classification when consumption is adjusted accordingly. Our fourth contribution is a comparison of three alternative approaches to drawing a common vulnerability line. The first one selects the arithmetic mean of household-specific vulnerability lines as the common one. The second approach anchors the risk deflator to the value of the poverty line. Both approaches lead to correct classification of

households into vulnerability but change the metric on which consumption is measured. The third approach we consider measures consumption using the original metric, but choses the common vulnerability line that minimizes the misclassification caused but the heterogeneity in household risks.

The theoretical results are independent of the available data and of the method used to estimate household risks. Panel models, structural models, distributional regressions, or machine-learning estimators may all supply the required risk functions, as may models estimated from a single cross-section under suitable assumptions about how future consumption depends on current consumption. To show how the framework can be implemented when only one cross-section is available, we use the lognormal approach of Chaudhuri et al. (2002).

This empirical illustration reveals an assumption that is often left implicit. Estimating a household's probability of future poverty tells us whether it is vulnerable at its current consumption level. To determine its vulnerability line, however, we also need to know how that probability would change if the household's current consumption were higher or lower. This requires additional assumptions about how future consumption depends on current consumption. Our illustration is based on the simple assumption that a change in current consumption from c_h to c multiplies every possible future consumption by the same factor c/c_h.

The paper is related to several strands of the vulnerability literature. Measures based on expected utility or expected poverty evaluate the welfare consequences of uncertain future consumption (Ligon and Schechter, 2003; Calvo and Dercon, 2013; Dutta et al., 2025). A large empirical literature estimates poverty risk with limited data and studies alternative sources of uncertainty (Günther and Harttgen, 2009; Gallardo, 2018). Other work calibrates monetary thresholds intended to identify households at risk of downward mobility or to define the lower boundary of the middle class (López-Calva and Ortiz-Juárez, 2014; Dang and Lanjouw, 2017). Our object is different. We ask when a probability-based identification rule admits an exact monetary representation and how a common reference can be defined without suppressing heterogeneous risks.

The rest of the paper is organized as follows. Section 2 develops the theoretical framework, characterizes household-specific lines, and establishes the implication of adjusting household consumption for risk. Section 3 characterizes an arithmetic reference, presents the joint classification of poverty and vulnerability using a common poverty benchmark, and assesses the errors

from using a common cutoff in real consumption. Section 4 develops the empirical illustration based on cross-sectional data. Section 5 concludes. Proofs and additional derivations are collected in the Appendix.

2. Theoretical framework

Consider a finite population of households $h \in \mathcal{H} = \{1, \ldots, H\}$. Let $c > 0$ denote current real consumption, expressed in the same units as a fixed poverty line $z > 0$. The observed value for household h is c_h. We assume that the vulnerability analysis refers to a specified forecast horizon, which is held fixed in the notation.

For household h, let future consumption at counterfactual current consumption c be

$$C_{h,t+1}(c) = g_h(c, Y_h), \tag{2}$$

where Y_h collects the uncertain states over the considered time horizon. The function g_h in equation (2) is a theoretical object which we assume to be known to the evaluator. In practice, both g_h and Y_h will need to be estimated. It is important to observe that equation (2) is not merely a forecast at the observed c_h; it also describes how the distribution of future consumption would move if current consumption were set to another level.

Under the vulnerability-as-expected-poverty approach, household h's vulnerability at any given level of current consumption c is

$$v_h(c) = \Pr(C_{h,t+1}(c) < z)\,. \tag{3}$$

The function $v_h(c)$ describes how the household's probability of future poverty varies with its current consumption. Evaluated at the observed consumption level c_h, it gives the household's vulnerability $v_h(c_h)$. We assume that the vulnerability function $v_h : (0, \infty) \to [0, 1]$ is continuous in current consumption. This holds, for example, when future consumption changes continuously with current consumption and its conditional distribution has no mass at the poverty line.

Household h is classified as vulnerable at consumption level c when

$$v_h(c) \geq \tau, \tag{4}$$

where $\tau \in (0, 1)$ is the minimum probability of future poverty at which a household is classified as vulnerable. The threshold τ is chosen by the evaluator and treated as given throughout the analysis.

This classification concerns future poverty and does not distinguish between households that are poor today and those that are not. To identify the vulnerable non-poor, rule (4) must therefore be combined with $c \geq z$: these households are currently above or at the poverty line but face a probability of future poverty of at least τ. Current poverty and vulnerability must be assessed separately, since a household may be poor today without being vulnerable at the selected threshold. As shown below, a monetary vulnerability line need not lie above z.

The first step is to establish under what conditions the probability-based rule in (4) can be expressed equivalently by comparing current consumption with a household-specific monetary vulnerability line.

Intuitively, if v_h is weakly decreasing in current consumption, a vulnerability line at τ exists provided that the household is vulnerable at some consumption level $c_L > 0$ and no longer vulnerable at a higher, finite level c_U: $v_h(c_L) \geq \tau$ and $v_h(c_U) < \tau$. The next section formalizes the conditions for the existence of $z_h^v(\tau)$.

2.1. Household-specific vulnerability lines

For a fixed τ, define the set of current consumption levels at which household h is vulnerable:

$$\Omega_h(\tau) = \{c > 0 : v_h(c) \geq \tau\}. \tag{5}$$

Definition 1 (Exact household vulnerability line)**.** *Household h has an exact vulnerability line at τ if there exists a finite $z_h^v(\tau) > 0$ such that*

$$v_h(c) \geq \tau \quad \Longleftrightarrow \quad c \leq z_h^v(\tau) \tag{6}$$

for every $c > 0$.

The line is therefore the largest current consumption level at which the household is still classified as vulnerable. It marks the boundary between vulnerability and non-vulnerability across possible consumption levels, rather than describing the household's status only at its observed consumption.

Proposition 1 (Threshold representation)**.** *Suppose that v_h is continuous and $\Omega_h(\tau)$ is nonempty and bounded above. An exact household vulnerability line exists if and only if $\Omega_h(\tau)$ is a lower set, that is, whenever $c_2 \in \Omega_h(\tau)$ and $0 < c_1 < c_2$, then $c_1 \in \Omega_h(\tau)$. When it exists, the line is unique and*

$$z_h^v(\tau) = \sup \Omega_h(\tau). \tag{7}$$

Proposition 1 gives the minimal requirement for the existence of a vulnerability line at a given τ. The vulnerability function v_h may be non-monotonic. It may rise and fall with current consumption because the composition of income, insurance, labor supply, or shock exposure changes. What matters is that it crosses the selected probability threshold only once as in panels (a) and (b) of Figure 1. In panel (c), a household may cease to be vulnerable as consumption rises and become vulnerable again at a higher consumption level, so no single consumption cutoff can separate vulnerability from non-vulnerability.

If v_h is continuous and strictly decreasing, the line is simply the unique solution to

$$v_h(z_h^v(\tau)) = \tau.$$

Proposition 1 is more general because it also permits non-monotone risk functions, provided that the selected threshold is crossed only once.

Corollary 1 (Lines at every probability threshold)**.** *Assume that v_h is continuous and that, for every $\tau \in (0,1)$, $\Omega_h(\tau)$ is nonempty and bounded above. Then v_h admits an exact line for every $\tau \in (0,1)$ if and only if $v_h(c)$ is weakly decreasing in c.*

Global monotonicity is therefore unnecessary for one policy threshold but becomes essentially necessary if the analyst wants a coherent family of lines for all thresholds. Figure 1 illustrates the distinction. Panel (a) shows the monotone case where to each τ correspond a unique z_h^v. Panel (b) shows a non-monotone risk function that nevertheless crosses the selected τ only once; a line still exists for this risk threshold but may not exist for others. In Panel (c), the vulnerable set contains separated intervals, so no scalar cutoff can reproduce the probability rule.

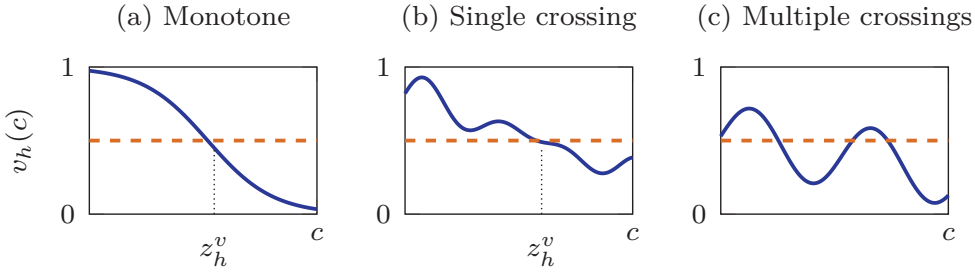


Figure 1: Threshold representability of household poverty risk. The dashed line is the probability threshold $\tau = 0.5$. A vulnerability line exists in panels (a) and (b), but not in panel (c).

Because the risk threshold τ is chosen by the analyst, we require the framework to accommodate any $\tau \in (0, 1)$. We therefore assume throughout the remainder of the theoretical analysis that the conditions of Corollary 1 hold for every household. This guarantees a finite, positive household-specific vulnerability line $z_h^v(\tau)$ for every admissible risk threshold.

These conditions ensure that each household has a well-defined vulnerability line. These lines may nevertheless differ across households, reflecting differences in their vulnerability functions. The next question is whether a single monetary line can reproduce the same classification for all households.

2.2. The problem of a common line

The conditions established in Section 2.1 ensure the existence of $z_h^v(\tau)$ for any household. However, because households face different future risks, we have a collection $\{z_1^v(\tau), \ldots, z_H^v(\tau)\}$ of household vulnerability lines.

In poverty analysis, households are classified by comparing their real consumption with a single poverty line z. Here, real consumption means consumption adjusted for spatial price differences and inflation during the survey period (Amendola et al., 2024b). Can the same approach be used for vulnerability, with a single monetary line reproducing the classification in (6)? Proposition 2 shows that this is, in general, impossible.

Proposition 2 (Impossibility of a common line in unadjusted consumption)**.** *Suppose that every household h has a well-defined vulnerability line $z_h^v(\tau)$ satisfying (6). A scalar $z_{\text{real}}^v(\tau)$ reproduces the probability-based classification for every household h and every $c > 0$, that is,*

$$v_h(c) \geq \tau \quad \Longleftrightarrow \quad c \leq z_{\text{real}}^v(\tau),$$

if and only if $z_h^v(\tau) = z_{\text{real}}^v(\tau)$ *for all* h*. Consequently, if at least two households have different vulnerability lines, no such common scalar exists.*

A common vulnerability line can therefore be used only if all household vulnerability lines coincide at the chosen probability threshold τ – and identical vulnerability functions are sufficient, but not necessary, for this coincidence. Given the differences across households in exposure to shocks, access to insurance, and coping mechanisms, coincident lines are an exceptional case rather than a plausible assumption. Proposition 2 is, de facto, an impossibility result: in general, no common line in real consumption can correctly classify every household as vulnerable or non-vulnerable at all consumption levels. This result provides normative reasons against which assessing the consistency of the methods proposed by López-Calva and Ortiz-Juárez (2014) or Dang and Lanjouw (2017) to set a common real-consumption vulnerability line.

Section 2.3 develops a simple solution: deflating each household's real consumption to account for differences in risk. A common vulnerability line can then be used in these risk-adjusted consumption units.

2.3. An exact representation in risk-adjusted consumption

This section shows how adjusting consumption allows households with different vulnerability lines to be assessed against a common reference. Each household is assumed to have a well-defined, positive vulnerability line $z_h^v(\tau)$. For any positive common reference $z^v(\tau)$, define

$$RD_h(\tau) = \frac{z_h^v(\tau)}{z^v(\tau)} \qquad \text{and} \qquad \widetilde{c}_h(c;\tau) = \frac{c}{RD_h(\tau)}. \tag{8}$$

The risk deflator $RD_h(\tau)$ expresses household h's vulnerability line relative to the chosen reference. $\widetilde{c}_h(c;\tau)$ is instead the risk-adjusted consumption which maps each household's vulnerability line to the common one: $\widetilde{c}_h(z_h^v(\tau);\tau) = z^v(\tau)$.

This construction is formally analogous to the use of equivalence scales in poverty analysis: household consumption is divided by a household-specific factor before being compared with a common reference. Here, however, the factor reflects the household's vulnerability cutoff relative to the reference, rather than differences in household needs. This simple rescaling allows us to represent each household's vulnerability classification using the same vulnerability line in risk-adjusted consumption. Proposition 3 shows that this rescaling preserves the vulnerability classification exactly and that, once the

common vulnerability line is fixed, the household-specific deflator is uniquely determined.

Proposition 3 (Exact representation and conditional uniqueness)**.** *Suppose that every household h has a positive vulnerability line $z_h^v(\tau)$, and fix any positive common reference $z^v(\tau)$. Let $RD_h(\tau)$ be defined as in (8). Then:*

1. *For every household h and every $c > 0$,*

$$v_h(c) \geq \tau \quad \Longleftrightarrow \quad c \leq z_h^v(\tau) \quad \Longleftrightarrow \quad \frac{c}{RD_h(\tau)} \leq z^v(\tau). \tag{9}$$

2. *If a positive factor $d_h(\tau)$ also satisfies*

$$v_h(c) \geq \tau \quad \Longleftrightarrow \quad \frac{c}{d_h(\tau)} \leq z^v(\tau) \qquad \text{for every } c > 0,$$

then $d_h(\tau) = RD_h(\tau)$.

The first part of Proposition 3 shows how the adjustment works. Before adjustment, household h is classified as vulnerable when $c \leq z_h^v(\tau)$. Dividing both consumption and the household's vulnerability line by $RD_h(\tau)$ leaves this comparison unchanged. Since $z_h^v(\tau)/RD_h(\tau) = z^v(\tau)$, the same household is classified as vulnerable when $c/RD_h(\tau) \leq z^v(\tau)$. The second part shows that, once $z^v(\tau)$ is chosen, the adjustment factor for each household is uniquely determined: $RD_h(\tau) = z_h^v(\tau)/z^v(\tau)$. The value of $z^v(\tau)$ itself, however, remains a choice. A different positive value produces different adjustment factors but exactly the same vulnerability classification. An additional criterion is therefore needed to select the common line.

The adjustment does not remove differences in household vulnerability lines; it incorporates them into $RD_h(\tau)$, allowing risk-adjusted consumption to be compared with a common vulnerability line $z^v(\tau)$. A household with $RD_h(\tau) > 1$ has an individual vulnerability line $z_h^v(\tau)$ higher than the common line. Its consumption is therefore divided by a factor greater than one, reflecting the higher consumption level it needs to bring its probability of future poverty below τ.

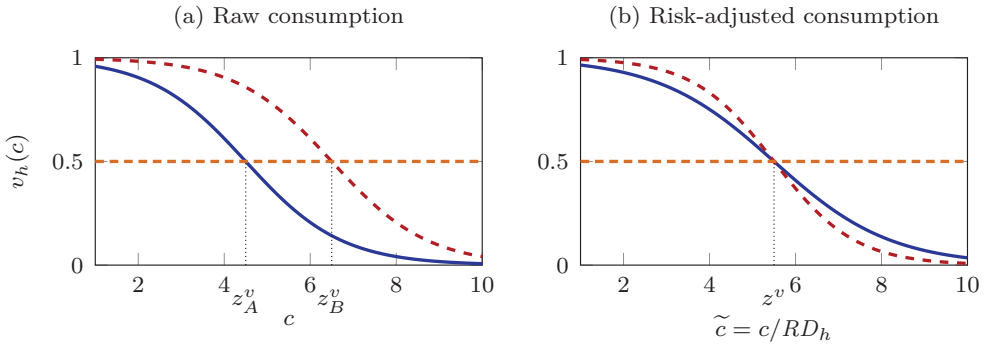
As in poverty measurement, nominal consumption may first be adjusted for differences in the cost of living and household needs (Amendola et al., 2024b; Mancini and Vecchi, 2022). If x_h denotes nominal consumption and $c_h = x_h/TCLI_h$ denotes the resulting comparable consumption, then

$$v_h(c_h) \geq \tau \quad \Longleftrightarrow \quad \frac{x_h}{TCLI_h RD_h(\tau)} \leq z^v(\tau). \tag{10}$$

The two adjustments serve different purposes: $TCLI_h$ accounts for differences in the cost of living and needs, whereas $RD_h(\tau)$ expresses the household's vulnerability cutoff relative to the common reference.

Figure 2 illustrates the equivalence established by Proposition 3 using two households, A and B. In the left panel, their vulnerability functions cross the probability threshold τ at different real consumption levels: $z_A^v = 4.5$ and $z_B^v = 6.5$. Each household therefore requires its own vulnerability line. In the right panel, consumption is divided by the corresponding risk deflator, $RD_A = 4.5/5.5$ for household A and $RD_B = 6.5/5.5$ for household B. Both crossings then occur at the chosen common line $z^v = 5.5$. Comparing adjusted consumption with this common line gives exactly the same classification as comparing real consumption with each household's individual line. The adjustment aligns the vulnerability cutoffs; it does not make the two vulnerability functions identical. The value 5.5 is chosen here for illustration: any positive common line would preserve the classification, provided that the deflators were adjusted accordingly.

Figure 2: Household vulnerability lines before and after risk adjustment.



Note: The solid blue curve represents household A and the dashed red curve household B; the horizontal dashed line marks $\tau = 0.5$. In panel (a), their risk functions reach this threshold at different consumption levels, $z_A^v = 4.5$ and $z_B^v = 6.5$. In panel (b), consumption is divided by $RD_h = z_h^v/z^v$, with $z^v = 5.5$. Both cutoffs therefore map to the same reference value, while the risk functions retain different shapes.

The vulnerability classification alone does not determine the monetary value of the common line. To see why, replace z^v by az^v and RD_h by RD_h/a, for any $a > 0$. Both risk-adjusted consumption and the common line are then multiplied by a, leaving each household on the same side of the line as before. The classification and the ranking of households by risk-adjusted

consumption therefore remain unchanged. Selecting a particular value of z^v requires an additional criterion. Section 3 compares two normalizations of the exact representation and then examines an approximation based on a common cutoff in real consumption.

The main results of this section can be summarized as follows. First, under the conditions of Corollary 1, each household has a finite, positive vulnerability line $z_h^v(\tau)$ in real consumption for every $\tau \in (0, 1)$. Second, dividing consumption by the household-specific risk deflator makes it possible to use a common vulnerability line $z^v(\tau)$ without changing which households are classified as vulnerable. Third, the classification rule does not uniquely determine the value of this common line: any positive $z^v(\tau)$ yields the same classifications when the deflators are adjusted accordingly. What remains is a practical criterion for choosing the common vulnerability line.

3. A common vulnerability line

This section addresses a practical dilemma in classifying households by poverty and vulnerability. In real consumption units, a common poverty line can be retained, but vulnerability generally requires household-specific lines. Adjusting consumption for risk allows a common vulnerability line, but makes the poverty line household-specific on that same scale. Neither scale therefore generally supports the familiar partition based on two common monetary cutoffs. This section provides an axiomatic basis for choosing the common vulnerability line and shows how combining current consumption with household-specific vulnerability lines yields a coherent partition into poor, vulnerable non-poor, and non-vulnerable non-poor households. It also considers anchoring the common reference to the poverty line, so that current poverty and vulnerability can be assessed against the same benchmark using two distinct consumption indicators. Finally, it provides a practical tool for measuring the classification errors that arise when a single vulnerability line is nevertheless used alongside the poverty line in real consumption units, and for selecting the vulnerability cutoff that minimizes those errors.

3.1. The average vulnerability line

Let $\omega_h > 0$ be normalized household weights, with $\sum_h \omega_h = 1$. Let $A(\mathbf{z}, \boldsymbol{\omega})$ be an aggregation rule that maps the positive household-specific lines $\mathbf{z} = (z_1^v, \ldots, z_H^v)$ and their weights into a positive common reference z^v.

Proposition 3 does not select a particular vulnerability line: any positive value can be used if the deflators are adjusted accordingly. In what follows we suggest two simple criteria to chose z^v.

The first criterion is a simple unanimity principle according to which, if all households have the same vulnerability line, the national one must coincide with this value. Let $\mathbf{1}_H$ denote a vector of H ones, so that $a\mathbf{1}_H$ represents a population in which every household has vulnerability line a.

Unanimity. For every $a > 0$, every population size H, and every normalized positive weight vector $\boldsymbol{\omega}$,

$$A(a\mathbf{1}_H, \boldsymbol{\omega}) = a. \tag{11}$$

Because policy analysis are often concerned with the measurement and comparison of vulnerability within population subgroups, we find also desirable for a common vulnerability line z^v to be *additively decomposable*. This allows a the national vulnerability line value to be calculated from subgroup vulnerability lines and their population weights, so that each group's contribution to the national reference is explicit. We express decomposability as follows.

- *Additive subgroup consistency.* For every partition $\mathcal{G}$ of the population into nonempty groups,

$$A(\mathbf{z}, \boldsymbol{\omega}) = \sum_{g \in \mathcal{G}} W_g A(\mathbf{z}_g, \boldsymbol{\omega}_g). \tag{12}$$

  where $W_g = \sum_{h \in g} \omega_h$ and $\omega_{h|g} = \omega_h / W_g$.

This property has a practical interpretation. Applied vulnerability analysis often compares population subgroups: Atamanov et al. (2022), for example, reports vulnerability rates separately for rural and urban areas, regions, and subregions in Uganda. When vulnerability lines are also calculated for such groups, the national line should be the same whether it is calculated directly from household lines or as a weighted average of subgroup lines. This parallels the decomposition of the poverty measures proposed by Foster, Greer and Thorbecke (1984), whereby national poverty is a population-weighted average of subgroup poverty. Here, the same aggregation principle applies to monetary vulnerability lines rather than poverty indices, allowing subgroup results to be combined into a consistent national summary.

Together, unanimity and additive subgroup consistency uniquely determine the common vulnerability line. Proposition 4 shows that, for a given τ,

$z^v(\tau)$ must equal the weighted arithmetic mean of the household-specific vulnerability lines.

Proposition 4 (Characterization of the common reference)**.** *An aggregation rule A satisfies unanimity and additive subgroup consistency for every finite weighted population if and only if*

$$z^v(\tau) = A(\mathbf{z}(\tau), \boldsymbol{\omega}) = \sum_{h=1}^{H} \omega_h z_h^v(\tau). \tag{13}$$

The result is simple both analytically and in practice: once the household-specific vulnerability lines have been obtained, calculating the common line requires only their weighted arithmetic mean. Unanimity and additive subgroup consistency thus turn an otherwise arbitrary choice into a straightforward calculation. Section 4 illustrates how these calculations can be implemented and combined with current poverty status to partition the population into poor, vulnerable non-poor, and non-vulnerable non-poor households.

3.2. Identifying the poor and the vulnerable

The relationship between the poverty line z and the common vulnerability line $z^v(\tau)$ requires clarification. The two serve different purposes. The poverty line specifies the standard of current consumption below which a household is classified as poor. It is expressed in comparable real-consumption units and held fixed when evaluating vulnerability. The common vulnerability line, by contrast, summarizes the household-specific consumption levels at which the probability of future poverty reaches the selected threshold τ. Its monetary value depends on z, the forecast horizon, the vulnerability functions, the population weights, and the normalization characterized in Proposition 4. It provides a single monetary summary of household-specific vulnerability lines but, unlike the poverty line z, which is compared with real consumption, must be compared with real consumption further adjusted for household-specific risk.

The relationship between the poverty line and the vulnerability line deserves closer attention. The common vulnerability line may lie below the poverty line, $z^v(\tau) < z$. This possibility also arises at the household level: a household's vulnerability line may satisfy $z_h^v(\tau) < z$, even when its probability of future poverty decreases as current consumption rises. This potentially counterintuitive result calls for an explanation. To understand this possibility, consider the poverty risk that a household would face if its

current consumption were exactly equal to the poverty line, $v_h(z)$. This risk need not be as high as τ: a household whose current consumption just meets the poverty standard may have sufficiently favourable prospects that its probability of future poverty is already below the selected threshold. If the household has an exact vulnerability line, this means that its vulnerability line lies below the poverty line:

$$z_h^v(\tau) < z \quad \Longleftrightarrow \quad v_h(z) < \tau.$$

In this case, current consumption would have to fall below the poverty line, to $z_h^v(\tau)$ or lower, for the household to meet the vulnerability criterion. At consumption levels between $z_h^v(\tau)$ and z, the household is therefore poor today but has a probability of future poverty below τ. This is not a contradiction: current poverty describes a household's present living standard, whereas vulnerability concerns its prospects over the specified forecast horizon. A household may be unable to meet its needs today yet have sufficiently favourable prospects of escaping poverty. Being classified as non-vulnerable does not mean that this outcome is certain, only that the probability of future poverty is below the selected threshold.

Thus, a household's vulnerability line is at or above the poverty line if and only if the household is vulnerable when its current consumption equals z. Some individual lines may lie below z while their weighted average lies above it, or vice versa.[1] An individual line below z is consistent with the definition of vulnerability: even at the poverty line, the household's probability of future poverty is less than τ. For that household, no current consumption level is simultaneously non-poor and vulnerable at the selected threshold.

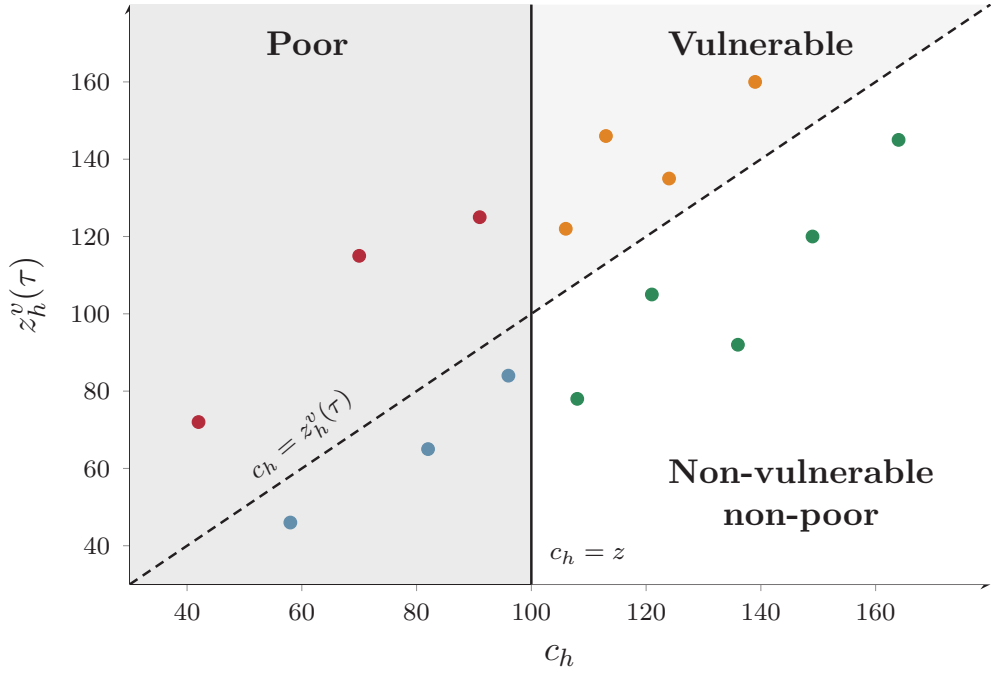
Finally, although z and $z^v(\tau)$ are both monetary values, they cannot generally be used as two common cutoffs on a single distribution of risk-adjusted consumption. The reason is that the two lines must be compared with consumption measured on different scales. The poverty line z applies to real consumption, $x_h/TCLI_h$, whereas the common vulnerability line $z^v(\tau)$ applies to consumption further adjusted for household-specific risk, $x_h/(TCLI_h RD_h(\tau))$. Using either scale therefore prevents the two lines from serving as common boundaries that divide a consumption distribution into the poor, the vulnerable non-poor, and the middle class, defined here as

[1] In particular, $v_h(z) \geq \tau$ for every household is sufficient for $z^v(\tau) \geq z$, but is not required for the weighted average to exceed z.

households that are neither poor nor vulnerable.

The desired population partition is nevertheless possible. Our proposed solution is to represent households in two dimensions, using their current real consumption and their household-specific vulnerability line. Both quantities are expressed in the same real-consumption units. This representation retains the information needed to assess current poverty and future poverty risk separately, while making the three population groups visible in a single diagram. Figure 3 illustrates this approach.

Figure 3: Mapping poverty and vulnerability



Note: Each point represents an household. Consumption and vulnerability lines are measured in the same monetary units; the poverty line is set to $z = 100$ for illustration. Currently poor households satisfy $c_h < z$, whereas vulnerable households satisfy $c_h \leq z_h^v(\tau)$. Red points identify poor and vulnerable households; blue points, poor and non-vulnerable households; orange points, vulnerable non-poor households; and green points, non-vulnerable non-poor households.

Figure 3 represents each household by its current real consumption, $c_h = x_h / TCLI_h$, and its household-specific vulnerability line, $z_h^v(\tau)$. Both axes use the same real-consumption units, consistent with the consumption-based welfare framework discussed by Deaton and Zaidi (2002), before the

additional risk adjustment. Households with identical current consumption may have different vulnerability lines because their risks and future consumption prospects differ. The diagram preserves this heterogeneity, which generally prevents a common cutoff in real consumption from reproducing the probability-based classification for every household at every consumption level, as established by Proposition 2.

The vertical line at $c_h = z$ separates currently poor from non-poor households. The diagonal $c_h = z_h^v(\tau)$ separates vulnerable households, on or above it, from non-vulnerable households below it, since $c_h \leq z_h^v(\tau)$ is equivalent to $v_h(c_h) \geq \tau$. These boundaries identify four groups. Among the currently poor, households on or above the diagonal are vulnerable, while those below it satisfy $z_h^v(\tau) < c_h < z$: they are poor today but face a probability of future poverty below τ, without a guaranteed exit from poverty. Among the non-poor, vulnerable households satisfy $z \leq c_h \leq z_h^v(\tau)$, while non-vulnerable households satisfy $c_h \geq z$ and $c_h > z_h^v(\tau)$. Combining the two currently poor groups yields the three reporting categories: poor, vulnerable non-poor, and non-vulnerable non-poor.

Figure 3 shows how current consumption and household-specific vulnerability lines jointly determine poverty and vulnerability status. The same classification can be made easier to read by expressing each household's consumption relative to its two thresholds: the poverty line and its own vulnerability line. Both comparisons then use a benchmark of one, making it immediately visible whether the household falls short of or exceeds each threshold. Choosing the poverty line as the common reference for risk adjustment gives this representation a simple interpretation: current and risk-adjusted consumption are assessed against the same poverty standard.

Given the probability threshold τ, and assuming that each household has a vulnerability line $z_h^v(\tau) > 0$, Proposition 3 permits any positive common vulnerability line. Since $z > 0$, we can therefore choose the poverty line itself and define

$$RD_h^z(\tau) = \frac{z_h^v(\tau)}{z}, \qquad \widetilde{c}_h^z(c;\tau) = \frac{c}{RD_h^z(\tau)} = \frac{zc}{z_h^v(\tau)}. \tag{14}$$

The superscript z makes explicit the choice of the poverty line as the reference for the normalization. Naturally, for every $c > 0$,

$$v_h(c) \geq \tau \quad \Longleftrightarrow \quad c \leq z_h^v(\tau) \quad \Longleftrightarrow \quad \widetilde{c}_h^z(c;\tau) \leq z,$$

which is a special case of equation (9).

By implementing equation (14) vulnerability is assessed by comparing $\widetilde{c}_h^z(c_h;\tau)$ with the same benchmark used for poverty analysis. In other words, current poverty is assessed by comparing c_h with z, and vulnerability is assessed by comparing $\widetilde{c}_h^z(c_h;\tau)$ with z as well. The two consumption indicators – c_h and $\widetilde{c}_h^z(c_h;\tau)$ – retain their distinct meanings: one records current consumption, while the other incorporates the household-specific vulnerability threshold.

To compare this construction with the arithmetic reference, retain $z^v(\tau) = \sum_h \omega_h z_h^v(\tau)$ and the corresponding adjusted consumption $\widetilde{c}_h(c;\tau)$ defined earlier. Direct substitution gives

$$\widetilde{c}_h^z(c;\tau) = \frac{z}{z^v(\tau)}\widetilde{c}_h(c;\tau) \iff \frac{\widetilde{c}_h^z(c;\tau)}{z} = \frac{\widetilde{c}_h(c;\tau)}{z^v(\tau)} = \frac{c}{z_h^v(\tau)}. \tag{15}$$

Changing the reference thus rescales every household's adjusted consumption by the same positive factor. It preserves its ordering across households, the vulnerability classification, and the ratio of adjusted consumption to the chosen benchmark.

The two references serve different purposes. The arithmetic reference is determined by the household lines and population weights. Proposition 4 selects it through unanimity and additive subgroup consistency, and its associated deflators have weighted mean one. The poverty-anchored reference is fixed by the poverty standard. Its deflators instead satisfy

$$\sum_h \omega_h RD_h^z(\tau) = \frac{\sum_h \omega_h z_h^v(\tau)}{z} = \frac{z^v(\tau)}{z}. \tag{16}$$

The two normalizations coincide precisely when $z^v(\tau) = z$.

Viewed as a rule for aggregating household lines, the fixed reference $A_z(\mathbf{z}, \boldsymbol{\omega}) = z$ generally fails unanimity: if all individual lines equal $a \neq z$, it returns z rather than a. It satisfies additive subgroup consistency when the same poverty line is used in every subgroup, since $\sum_g W_g z = z$, but this identity contains no information about subgroup vulnerability thresholds.

Choosing z as the reference allows current and risk-adjusted consumption to be assessed against the same monetary benchmark. Define the current-consumption welfare ratio and its vulnerability-adjusted counterpart as

$$w_h = \frac{c_h}{z}, \qquad \widetilde{w}_h(\tau) = \frac{\widetilde{c}_h^z(c_h;\tau)}{z} = \frac{c_h}{z_h^v(\tau)}. \tag{17}$$

These dimensionless indicators yield

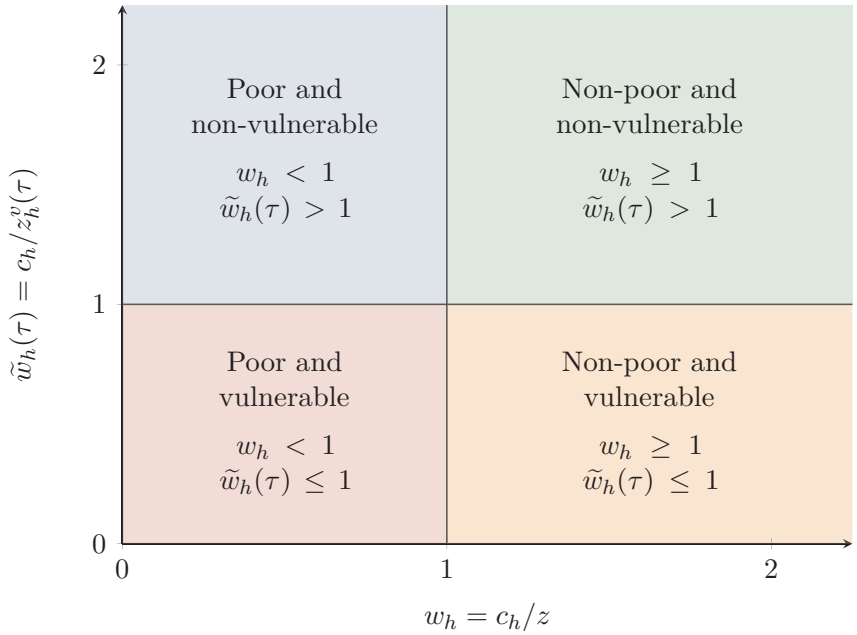
$$c_h < z \quad \Longleftrightarrow \quad w_h < 1, \qquad v_h(c_h) \geq \tau \quad \Longleftrightarrow \quad \widetilde{w}_h(\tau) \leq 1. \tag{18}$$

Figure 4 places the two ratios on separate axes. The vertical line identifies current poverty and the horizontal line identifies vulnerability. Their intersection separates the four combinations of current poverty and future poverty risk. This figure retains the information in the representation in Figure 3 based on $(c_h, z_h^v(\tau))$. For a given z, the transformation is invertible:

$$c_h = zw_h, \qquad z_h^v(\tau) = \frac{zw_h}{\widetilde{w}_h(\tau)}.$$

The two diagrams, however, emphasize different aspects of the same classification: the monetary representation (Figure 3) displays heterogeneity in required consumption thresholds, while the welfare ratios (Figure 4) display each household's position relative to its two standards.

Figure 4: Current poverty and vulnerability relative to a common unit benchmark



Note: The horizontal axis records current consumption relative to the poverty line. The vertical axis records current consumption relative to the household-specific vulnerability line.

3.3. Single lines and misclassification

Consider a simple experiment. Suppose that, for practical reasons, an analyst chooses to use a single vulnerability cutoff in real consumption, setting aside the household-specific vulnerability lines. How often would this change a household's vulnerability status? We answer this question by comparing the classification produced by the common cutoff with that obtained using each household's own vulnerability line.

We then ask a further question: which single cutoff produces the fewest such errors? If even this best-performing cutoff misclassifies a substantial share of households, the cost of ignoring heterogeneity is high. If the disagreement is sufficiently small for the intended application, however, the cutoff may offer a useful practical approximation: a compromise between retaining household-specific risk information and using a single threshold in real consumption.

Let $\zeta > 0$ denote a candidate common vulnerability line, applied to real consumption without adjustment for household-specific risk. Household h is then classified as vulnerable if $c_h \leq \zeta$, whereas the probability-based rule classifies it as vulnerable if $v_h(c_h) \geq \tau$. Let $\mathbf{1}\{\cdot\}$ equal one when the condition in braces holds and zero otherwise. Using the normalized household weights ω_h, define

$$M(\zeta;\tau) = \sum_{h=1}^{H} \omega_h \left|\mathbf{1}\{c_h \leq \zeta\} - \mathbf{1}\{v_h(c_h) \geq \tau\}\right|. \tag{19}$$

For each household, the absolute difference between the two indicators equals one if the single line and the household-specific line assign different vulnerability statuses, and zero otherwise. Their weighted sum, $M(\zeta;\tau)$, is therefore the misclassification rate: the population-weighted share of households whose vulnerability status changes when their individual lines are replaced by a single line ζ applied to real consumption. A value of zero means that the two rules classify every household identically at its observed consumption level. This agreement does not imply that the single line would reproduce each household's classification at every possible consumption level: different individual lines can yield the same classification at the consumption levels observed in the data.

The analyst can then select a line that minimizes this disagreement:

$$\zeta^*(\tau) \in \arg\min_{\zeta>0} M(\zeta;\tau). \tag{20}$$

The minimum, $M(\zeta^*(\tau);\tau)$, measures the smallest misclassification achievable with a single line in real consumption for the population under study. If

this disagreement is sufficiently small for the intended application, the selected line may provide a useful practical approximation. Its usefulness also depends on which households are misclassified: a low overall disagreement rate can conceal a failure to identify a substantial share of vulnerable households when that group is small. The empirical section illustrates the use of this measure.

4. An empirical illustration

This section illustrates the framework using data from the Malawi Sixth Integrated Household Survey 2024–2025. Following Chaudhuri et al. (2002, CJS), we use a single cross-section to estimate each household's probability of future poverty, corresponding to the vulnerability measure in equation (3), and derive household-specific vulnerability lines. We then assess how closely a single line applied to observed real consumption reproduces this risk-based classification, examining the extent and geographical pattern of disagreement.

Let $c_h > 0$ denote current real consumption and let $\mathbf{X}_h$ be a $K \times 1$ column vector of household and location characteristics, including a constant. The CJS specification is

$$\begin{aligned} \log c_h &= \mathbf{X}_h'\boldsymbol{\beta} + \varepsilon_h, \qquad \mathbb{E}(\varepsilon_h \mid \mathbf{X}_h) = 0, \\ \sigma_h^2 &= \mathrm{Var}(\varepsilon_h \mid \mathbf{X}_h) = \mathbf{X}_h'\boldsymbol{\theta} > 0, \end{aligned} \tag{21}$$

where the conditional variance is modelled as a function of the same K-dimensional covariate vector $\mathbf{X}_h$, with $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ denoting the corresponding $K \times 1$ coefficient vectors. This heteroskedastic specification is important because households with similar expected log consumption may face different levels of consumption uncertainty and, consequently, different risks of falling into poverty.

CJS propose to estimate model (21) through a feasible generalized least squares procedure yielding the fitted conditional mean $\hat{\mu}_h$ and standard deviation $\hat{\sigma}_h$.[2] The estimation of future poverty risk (vulnerability) requires

[2]An initial regression of log consumption on $\mathbf{X}_h$ produces residuals, whose squares are regressed on the variance covariates to obtain a preliminary variance function. The variance equation is then estimated again with precision weights inversely proportional to the square of this preliminary fitted variance. Finally, the consumption equation is estimated again with weights inversely proportional to the resulting fitted variance. Survey weights enter separately from these precision weights: the former determine population representation, whereas the latter account for heteroskedasticity.

the additional assumption that the conditional consumption distribution is stable over the forecast horizon and lognormal

$$\log C_{h,t+1} \mid \mathbf{X}_h \sim N(\mu_h, \sigma_h^2), \qquad \mu_h = \mathbf{X}_h'\boldsymbol{\beta}. \tag{22}$$

Replacing the unknown parameters by their estimates gives a draw $\widehat{Y}_h$ from the fitted distribution of future consumption satisfying

$$\log \widehat{Y}_h \mid \mathbf{X}_h \sim N(\hat{\mu}_h, \hat{\sigma}_h^2), \qquad \widehat{\sigma}_h = \sqrt{\hat{\sigma}_h^2} > 0. \tag{23}$$

For a poverty line $z > 0$, the fitted probability of future poverty at observed consumption is

$$\widehat{v}_h(c_h) = \Pr(\widehat{Y}_h < z \mid \mathbf{X}_h) = \Phi\left(\frac{\log z - \hat{\mu}_h}{\widehat{\sigma}_h}\right), \tag{24}$$

where Φ is the standard normal distribution function. Comparing this probability with a pre-specified threshold τ identifies vulnerable households.[3]

4.1. Constructing household-specific vulnerability lines

To find a household's vulnerability line, we need to know how its probability of future poverty would change at different levels of current consumption. We assume that a given percentage change in current consumption produces the same percentage change in every possible future consumption outcome. We call this assumption proportional scaling. If current consumption changes from the observed level c_h to $c > 0$, future consumption is therefore represented by $(c/c_h)\widehat{Y}_h$. At $c = c_h$, this gives the fitted distribution of future consumption. At other consumption levels, it changes the scale of that distribution while keeping all outcomes positive.

Let $\widehat{Q}_h(\tau)$ be the level of future consumption below which household h is predicted to fall with probability τ, evaluated at its observed current consumption c_h. Under the distributional assumptions above, this quantile is

$$\widehat{Q}_h(\tau) = \exp\left[\hat{\mu}_h + \widehat{\sigma}_h \Phi^{-1}(\tau)\right], \qquad 0 < \tau < 1. \tag{25}$$

This quantile allows us to determine the level of current consumption at which the household's probability of future poverty equals τ. A change in

[3]Notice that, in practice, cross-sectional dispersion is used to proxy household uncertainty over time, potentially confounding persistent unobserved heterogeneity and measurement error with future uncertainty.

current consumption from c_h to $c > 0$ is assumed to multiply every possible future consumption outcome by the same factor c/c_h. The lower τ-quantile therefore becomes

$$\frac{c}{c_h}\widehat{Q}_h(\tau).$$

The vulnerability boundary is reached when this quantile coincides with the poverty line z. Setting the two equal and solving for c gives

$$\frac{c}{c_h}\widehat{Q}_h(\tau) = z \quad \Longleftrightarrow \quad \widehat{z}_h^v(\tau) = \frac{z c_h}{\widehat{Q}_h(\tau)}. \tag{26}$$

Hence, $\widehat{z}_h^v(\tau)$ is the level of current consumption at which household h faces a probability exactly equal to τ of falling below the poverty line in the future. Equivalently, $z/\widehat{Q}_h(\tau)$ is the proportional adjustment to current consumption required to bring the household exactly to this risk boundary.

This expression also provides a direct interpretation of the household-specific line. If $\widehat{Q}_h(\tau) < z$ at observed consumption, the household already faces a poverty probability greater than τ, and therefore $\widehat{z}_h^v(\tau) > c_h$: current consumption would have to increase to bring its risk down to the threshold. Conversely, if $\widehat{Q}_h(\tau) > z$, then $\widehat{z}_h^v(\tau) < c_h$: current consumption could fall before the household reached the vulnerability threshold.

More formally, under proportional scaling,

$$\Pr\left(\frac{c}{c_h}\widehat{Y}_h < z \,\middle|\, \mathbf{X}_h\right) \geq \tau \quad \Longleftrightarrow \quad \frac{c}{c_h}\widehat{Q}_h(\tau) \leq z \quad \Longleftrightarrow \quad c \leq \widehat{z}_h^v(\tau). \tag{27}$$

The first equivalence follows from the continuity and strict monotonicity of the fitted lognormal distribution. Since c_h, z, and $\widehat{Q}_h(\tau)$ are positive and finite, $\widehat{z}_h^v(\tau)$ is also positive and finite. Importantly, when evaluated at observed consumption $c = c_h$, this rule yields exactly the same vulnerability classification as comparing the probability in (24) with τ.

Following Proposition 4, household-specific vulnerability lines can then be summarized by a common line,

$$z^v(\tau) = \sum_h \omega_h \widehat{z}_h^v(\tau), \qquad \omega_h > 0, \qquad \sum_h \omega_h = 1, \tag{28}$$

where ω_h are normalized population weights in the estimation sample.

4.2. Results

We use consumption data from the Malawi Sixth Integrated Household Survey 2024–2025 to assess how closely a common vulnerability line in real consumption reproduces the classification based on household-specific poverty risks. Our working sample consists of 12,501 households.[4]

We use real monthly per capita consumption as the welfare indicator and the official poverty line of $z = 47,528.20$ MWK, expressed in July–August 2024 prices.[5] Estimation and population shares are based on survey person weights, with ω_h denoting the corresponding normalized weight. In the CJS model specification, the conditional mean and variance equations include the same set of regressors: the household head's sex, age and age squared, education, and primary economic activity during the preceding twelve months, along with indicators for each combination of region and urban or rural residence.

We construct the household lines using (26), with $\tau = 0.29$, and obtain $z^v(\tau)$ from their mean computed with population weights.[6] Figure 5 is the empirical counterpart of Figure 3. The vertical line separates the currently poor from the non-poor, while households on or above the diagonal are classified as vulnerable because their consumption does not exceed their household-specific vulnerability line. The two classifications differ: 36.90% of the population is non-poor but vulnerable, while 2.79% is currently poor but has an estimated probability of future poverty below $\tau = 0.29$. Note that, for this latter group, the classification reflects an expected lower probability of poverty, not an observed exit from poverty. The dispersion of vulnerability lines at

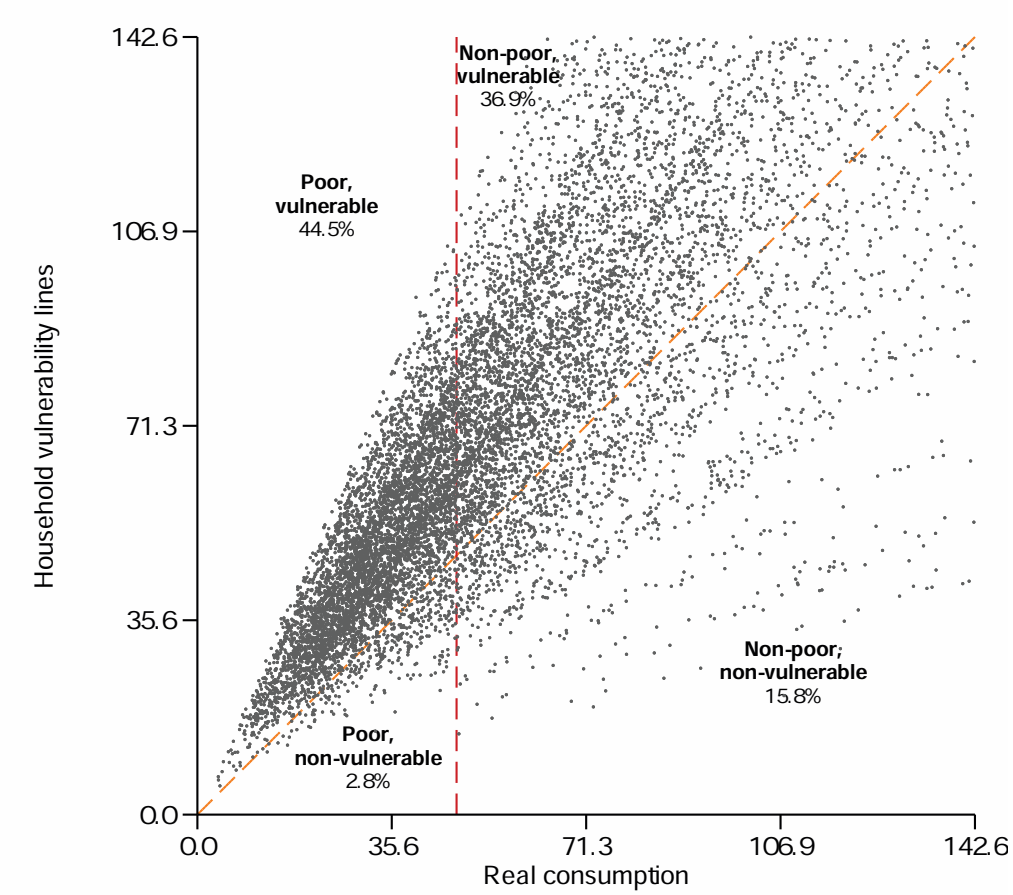
[4]We exclude only one observation due to missing information on the household head's schooling attendance and qualification. Household heads who never attended school are distinguished from those who attended school but obtained no qualification. We recoded the 374 documented questionnaire skips in annual activity into a separate category, "No listed annual activity reported", rather than treating these observations as economically inactive. The thirteen observed annual activity statuses are harmonized into five broad groups, with cooperative or unclassifiable status retained as an explicit residual group.

[5]We adopt the official national poverty line supplied with the other survey information, which is based on the cost-of-basic-needs approach. The poverty line combines the cost of a food basket providing 2,215 kilocalories per person per day with an allowance for essential non-food consumption (National Statistical Office, 2026). The report uses a July–August 2024 price reference. We retain the supplied real consumption and poverty line without further deflation.

[6]Under equal annual poverty risks and temporal independence, a 50 percent probability of poverty at least once within two years corresponds to an annual probability of $1 - 1/\sqrt{2}$, approximately 0.29.

similar consumption levels illustrates why a single line applied to observed consumption cannot generally reproduce the probability-based classification.

Figure 5: Poverty and vulnerability in Malawi, 2024–2025 ($\tau = 0.29$).



Notes: Each point represents a household. Both axes are in thousands of real MWK per capita per month, at July–August 2024 prices. Percentages are shares calculated using person weights for all 12,501 households in the estimation sample. The display shows 10,627 households with consumption and vulnerability lines no greater than three times the poverty line. Current poverty uses $c_h < z$ and vulnerability uses $c_h \leq \widehat{z}_h^v(0.29)$.

The monetary representation shows how much current consumption each household would need to reach the vulnerability boundary. A complementary question is how its living standard compares with the poverty standard and

with its own vulnerability threshold. Using the poverty line z as the common reference allows both assessments to use a unit benchmark, without treating current consumption and risk-adjusted consumption as the same indicator. Define

$$w_h = \frac{c_h}{z}, \qquad \widetilde{w}_h(\tau) = \frac{c_h}{\widehat{z}_h^v(\tau)} = \frac{\widehat{Q}_h(\tau)}{z}. \tag{29}$$

The first ratio measures current consumption relative to the poverty line: a value below one identifies current poverty, while a value above one records consumption in excess of that standard. The second measures current consumption relative to the household's vulnerability line. At one the household is exactly at the selected risk boundary; above one its current consumption exceeds the level required to bring fitted poverty risk down to τ, and below one it falls short. Under proportional scaling, the same ratio equals the fitted lower τ-quantile of future consumption relative to z. Neither ratio is a probability, and equal distances from one do not imply equal differences in poverty risk across households.

This representation retains the arithmetic reference in equation (28) as a summary of household lines. With $RD_h(\tau) = \widehat{z}_h^v(\tau)/z^v(\tau)$, its corresponding adjusted consumption is $\widetilde{c}_h = c_h/RD_h(\tau)$. Choosing z instead gives $RD_h^z(\tau) = \widehat{z}_h^v(\tau)/z$ and $\widetilde{c}_h^z = c_h/RD_h^z(\tau)$, so

$$\widetilde{c}_h^z = \frac{z}{z^v(\tau)}\widetilde{c}_h, \qquad \frac{\widetilde{c}_h^z}{z} = \frac{\widetilde{c}_h}{z^v(\tau)} = \widetilde{w}_h(\tau). \tag{30}$$

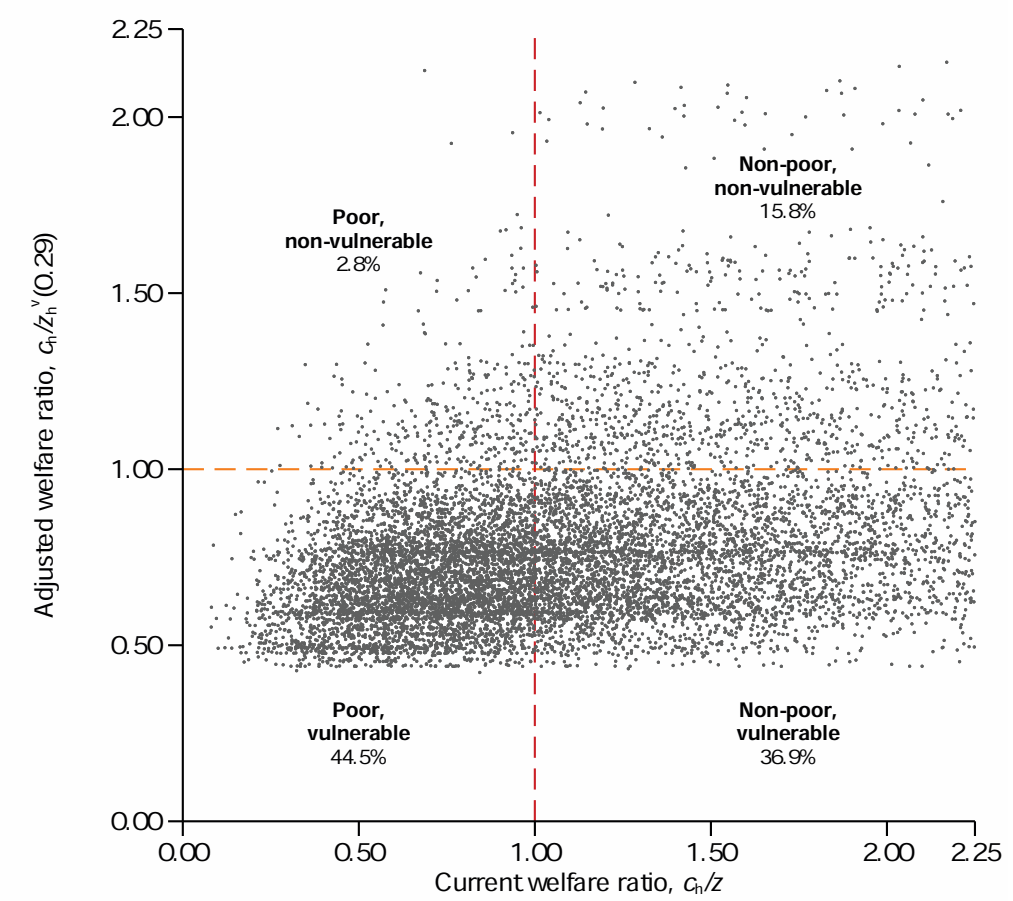
The normalization determines the scale of adjusted consumption; the welfare ratio records position relative to the chosen reference; the classification compares that ratio with one. Changing the normalization therefore leaves the classification unchanged, as established by the theory.[7]

Figure 6 represents these two welfare ratios. The vertical line at $w_h = 1$ identifies current poverty and the horizontal line at $\widetilde{w}_h(0.29) = 1$ identifies vulnerability. The lower-left region contains the poor and vulnerable, accounting for 44.48% of the population. The lower-right region contains the vulnerable non-poor (36.90%): these households meet the current poverty standard but remain below their own vulnerability threshold. The upper-left region

[7] As an aggregation rule, fixing the reference at z would not satisfy unanimity when all household lines equal some $a \neq z$. The arithmetic reference retains that property and its population-weighted subgroup decomposition.

contains the poor and non-vulnerable (2.79%), while the upper-right region contains the non-vulnerable non-poor (15.83%). Current poverty therefore affects 47.27% of the represented population, whereas 81.38% is classified as vulnerable by the model. The two diagrams give the same partition, but emphasize different aspects of it: monetary thresholds in Figure 5, and living standards relative to the two relevant boundaries in Figure 6.

Figure 6: Current poverty and vulnerability relative to a common unit benchmark, Malawi, 2024–2025 ($\tau = 0.29$).



Notes: The horizontal axis records $w_h = c_h/z$ and the vertical axis $\widetilde{w}_h(0.29) = c_h/\widehat{z}_h^v(0.29)$. Current poverty lies to the left of one and vulnerability on or below one. Percentages use person weights for all 12,501 households in the estimation sample. The plot shows 10,405 households with both ratios no greater than 2.25.

We next ask how well the probability-based classification can be reproduced when the household-specific vulnerability lines are replaced by a single threshold applied directly to observed consumption. Table 1 compares two such thresholds, $z^v(\tau)$ and $\zeta^*(\tau)$.

The common vulnerability line $z^v(\tau)$ is calculated as the weighted average of household-specific lines, as defined in equation (28), and is constructed for use with risk-adjusted consumption. Applying its value of 76.86 thousand MWK directly to observed consumption produces a disagreement rate of 21.67% relative to the probability-based classification. Thus, applying this common line to observed consumption assigns more than one fifth of the population to a different vulnerability category than the household-specific lines would. This disagreement can be reduced by choosing a line specifically for use with observed consumption. The line $\zeta^*(\tau)$, defined by the minimization problem in equation (20), equals 127.83 thousand MWK and yields a disagreement rate of 15.74%. The selected line lies above the poverty line, an outcome of minimizing over all $\zeta > 0$ without imposing $\zeta \geq z$.

Table 1: Common vulnerability lines and the share of the population classified differently from household-specific lines, Malawi, 2024–2025 ($\tau = 0.29$).

Line	Value (thousand MWK)	Disagreement (%)
$z^v(\tau)$	76.86	21.67
$\zeta^*(\tau)$	127.83	15.74

Notes: Values are thousands of real MWK per capita per month, at July–August 2024 prices. Disagreement is the person-weighted percentage whose vulnerability status differs from the fitted probability rule on the complete estimation sample. The minimizing interval is approximately $[127,804.26,\ 127,848.03)$ MWK, and the table reports its midpoint. Rounded endpoints do not define exact inclusion boundaries.

A simple benchmark helps assess how well $\zeta^*(\tau)$ performs. If every household were classified as non-vulnerable, the classification would differ from that based on household-specific vulnerability lines for 81.38% of the population: precisely the share identified as vulnerable by the model. Since vulnerability is widespread in this application, classifying everyone as vulnerable provides the more demanding constant-classification benchmark. Its disagreement is 18.62%, compared with 15.74% under the optimized cutoff, an improvement of 2.88 percentage points. Thus, even choosing the best possible line for these

data leaves about one sixth of the population classified differently from the household-specific vulnerability lines.

This finding does not contradict the theoretical result that vulnerability can be assessed using a common line. That result requires comparing the common line $z^v(\tau)$ with risk-adjusted consumption. The adjustment is specific to each household and preserves the information contained in its vulnerability line, so the resulting classification is exactly the same. When a common line is applied directly to observed consumption, this information is lost. The remaining disagreement of 15.74% shows that, in this application, even the best single line cannot fully compensate for omitting the household-specific adjustment. This is disagreement with the model-defined classification, not a measured error in forecasting observed future poverty.

We next examine whether this loss of information has a systematic geographical pattern. Partitioning the population by region and urban/rural status, we compute subgroup-specific vulnerability lines $z_g^v(\tau)$. By construction, the national line can be written as

$$z^v(\tau) = \sum_g W_g z_g^v(\tau),$$

where W_g denotes subgroup g's share of the represented population. The weighted subgroup lines reproduce exactly the national value of 76.86 thousand MWK. Table 2 then compares the geographical composition of the vulnerable population under the fitted probability rule with that obtained by applying the nationally optimized threshold $\zeta^*(0.29)$ directly to observed consumption.

Table 2: Geographical distribution of the vulnerable population under household-specific and common vulnerability lines, Malawi, 2024–2025 ($\tau = 0.29$).

	Probability rule $\widehat{v}_h(c_h) \geq 0.29$			Optimized common line $c_h \leq \zeta^*(0.29)$		
Region	Rural	Urban	Total	Rural	Urban	Total
North	11.24	0.71	11.95	10.72	2.01	12.73
Central	39.00	2.33	41.33	36.40	5.91	42.31
South	44.73	1.99	46.72	39.50	5.46	44.96
Total	94.98	5.02	100.00	86.62	13.38	100.00

Notes: Entries show how the population classified as vulnerable under each rule is distributed across regions and urban/rural areas. The two rules identify groups of different sizes: 81.38% of the total population under the probability rule and 91.05% under the optimized common line. These overall shares are not entries in the table. The common-line classification compares observed consumption directly with $\zeta^*(0.29)$, without household-specific risk adjustment or territorial reoptimization. All figures use person weights for the complete estimation sample.

Using a common vulnerability line changes not only the estimated prevalence of vulnerability but also its geographical pattern. In urban Central Malawi, for example, the share classified as vulnerable rises from 25.39% under household-specific lines to 72.18% under the optimized common line. Ignoring household risk heterogeneity can therefore substantially alter the geographical priorities suggested by the analysis.

5. Conclusion

How many vulnerability lines do we need? As it is often the case in many economic contexts, our answer is "it depends."

If the only goal of our analysis is to identify the vulnerable population by tracing a single line, then we should adjust household consumption to account for risk. When the metric is risk-adjusted household consumption, one common vulnerability line is sufficient. Our theoretical discussion and results have shown the conditions under which a vulnerability line exists for each household, and suggested a pragmatic approach to uniquely identify a national vulnerability line. However, because risk-adjustment factors are household-specific, in this metric we have one vulnerability line but household-specific poverty lines.

Conversely, when household-specific vulnerability lines differ, a common vulnerability cutoff in real consumption cannot reproduce the probability-based classification for every household at every consumption level. A common cutoff may nevertheless reproduce the classifications at the consumption levels observed in a particular sample. Its empirical adequacy must therefore be assessed rather than assumed.

The practical implication is that current consumption alone generally cannot identify both poverty and vulnerability. Figure 3 provides a simple alternative: representing each household by its consumption and its own vulnerability line. Figure 4 expresses the same information as consumption relative to the two thresholds, with a benchmark of one on each axis. Both representations identify the poor, the vulnerable non-poor, and the non-vulnerable non-poor, while retaining the distinction between vulnerable and non-vulnerable households among the currently poor.

A single line applied directly to real consumption may nevertheless provide a useful approximation in a particular application. Its accuracy must be assessed rather than assumed. In our Malawi illustration for 2024–25, the two common lines considered assign approximately 16–22 percent of the population a different vulnerability status from that obtained using household-specific lines. Even the line chosen to minimize disagreement leaves 15.74 percent classified differently. These discrepancies also affect the geographical pattern of vulnerability: in urban Central Malawi, the share classified as vulnerable rises from 25.39 percent under household-specific lines to 72.18 percent under the optimized common line.

For this reason, we suggest moving beyond the familiar partition of the real consumption distribution using a national poverty line and a national vulnerability line. Instead, we propose using a simple graphical instrument as in Figure 3, in which each household is represented by its real consumption and its vulnerability line, leading to a clear distinction of three relevant groups: the poor, the vulnerable non-poor and the non-vulnerable non-poor. Figure 4 provides an equivalent representation using current consumption relative to the poverty line and to the household-specific vulnerability line, with both ratios assessed against a unit benchmark.

A final corollary, so far left implicit, concerns the measurement of the middle class. When a vulnerability line defines its lower bound, as in Dang and Lanjouw (2017), households above that line are regarded as economically secure. Yet some of these households may still face a probability of future poverty above the selected threshold. If membership of the middle class

requires being both non-poor and non-vulnerable, consumption above a common line is therefore not sufficient: household-specific poverty risks must also be considered.

Data Availability

The data used in this study come from the Sixth Integrated Household Survey 2024–2025 (IHS6), conducted by the National Statistical Office (NSO) of Malawi and available through the World Bank Microdata Library (catalog 8507). The analysis uses five files from the official survey release: `ihs6_consumptionaggregates.dta`, `hh_mod_a_filt.dta`, `hh_mod_b.dta`, `hh_mod_c.dta`, and `hh_mod_e.dta`.

Researchers can obtain the microdata from https://microdata.worldbank.org/catalog/8507, subject to the Library's registration and access requirements. Neither the original microdata nor derived household-level datasets are redistributed with the replication code. As of 22 October 2026, the catalog lists release `MWI_2024-2025_IHS-VI_v01_M`, dated 21 July 2026.

The analysis also relies on the Stata package `vulnpov` available upon request from the authors.

References

Amendola, N., Belotti, F., Edo, M., Ranzani, M., Vecchi, G., 2024a. Poverty Lines and Spatial Differences in the Cost of Living. Policy Research Working Paper 10978. Washington, DC: World Bank.

Amendola, N., Mancini, G., Vecchi, G., 2024b. Setting Poverty Lines for Consumption-Based Welfare Measures: A Research Agenda. Policy Research Working Paper 10813. Washington, DC: World Bank.

Atamanov, A., Mukiza, C.N., Ssennono, V.F., 2022. Quantifying Vulnerability to Poverty in Uganda. Policy Research Working Paper 10007. Washington, DC: World Bank.

Calvo, C., Dercon, S., 2013. Vulnerability to Individual and Aggregate Poverty. Social Choice and Welfare 41, 721–740.

Chaudhuri, S., 2003. Assessing Vulnerability to Poverty: Concepts, Empirical Methods and Illustrative Examples. Department of Economics, Columbia University.

Chaudhuri, S., Jalan, J., Suryahadi, A., 2002. Assessing Household Vulnerability to Poverty from Cross-Sectional Data: A Methodology and Estimates from Indonesia. Department of Economics Discussion Paper 0102-52. New York: Columbia University.

Christiaensen, L.J., Boisvert, R.N., 2000. On Measuring Household Food Vulnerability: Case Evidence from Northern Mali. Working Paper 2000-05. Ithaca, NY: Cornell University.

Dang, H.A.H., Lanjouw, P.F., 2017. Welfare Dynamics Measurement: Two Definitions of a Vulnerability Line and Their Empirical Application. Review of Income and Wealth 63, 633–660.

Deaton, A., Zaidi, S., 2002. Guidelines for Constructing Consumption Aggregates for Welfare Analysis. Living Standards Measurement Study Working Paper 135. Washington, DC: World Bank.

Dutta, I., Foster, J., Mishra, A., Nandeibam, S., 2025. Identifying the Vulnerable: Concepts and Measurement. Review of Income and Wealth 71, e12722.

Foster, J., Greer, J., Thorbecke, E., 1984. A Class of Decomposable Poverty Measures. Econometrica 52, 761–766.

Gallardo, M., 2018. Identifying Vulnerability to Poverty: A Critical Survey. Journal of Economic Surveys 32, 1074–1105.

Günther, I., Harttgen, K., 2009. Estimating Households' Vulnerability to Idiosyncratic and Covariate Shocks: A Novel Method Applied in Madagascar. World Development 37, 1222–1234.

Ligon, E., Schechter, L., 2003. Measuring Vulnerability. Economic Journal 113, C95–C102.

López-Calva, L.F., Ortiz-Juárez, E., 2014. A Vulnerability Approach to the Definition of the Middle Class. Journal of Economic Inequality 12, 23–47.

Mancini, G., Vecchi, G., 2022. On the Construction of a Consumption Aggregate for Inequality and Poverty Analysis. Washington, DC: World Bank.

National Statistical Office, 2026. Poverty and Inequality Report: Evidence from the Sixth Integrated Household Survey (IHS6), 2024/25. Technical Report. Government of Malawi. URL: https://cms.nsomalawi.mw/api/download/834/Malawi_Poverty_and_Inequality_Report_IHS6_Final.pdf.

Pritchett, L., Suryahadi, A., Sumarto, S., 2000. Quantifying Vulnerability to Poverty: A Proposed Measure, Applied to Indonesia. Policy Research Working Paper 2437. Washington, DC: World Bank.

Ravallion, M., 1988. Expected Poverty under Risk-Induced Welfare Variability. Economic Journal 98, 1171–1182.

Appendix A. Proofs

Appendix A.1. Proof of Proposition 1

Proof. Suppose first that an exact household vulnerability line z_h^v exists. Then, by Definition 1, $v_h(c) \geq \tau$ is equivalent to $c \leq z_h^v$. Therefore,

$$\Omega_h(\tau) = \{c > 0 : v_h(c) \geq \tau\} = \{c > 0 : c \leq z_h^v\} = (0, z_h^v].$$

where the first equality is the definition of Ω_h in equation (5). In other words the set of current consumption levels at which household h is vulnerable tkes the form of an interval. It is thus a lower set, since if $c_2 \in \Omega_h(\tau)$ and $0 < c_1 < c_2$, then $c_1 \leq z_h^v$ and hence $c_1 \in \Omega_h(\tau)$, and its supremum is z_h^v.

Conversely, suppose that $\Omega_h(\tau)$ is nonempty, bounded above, and a lower set. Let

$$s = \sup \Omega_h(\tau).$$

Because the set is nonempty and its elements are positive, $s > 0$; because it is bounded above, $s < \infty$. We first show that every $c \in (0, s)$ belongs to $\Omega_h(\tau)$. By the definition of a supremum, for every such c there is an $x \in \Omega_h(\tau)$ satisfying $c < x \leq s$. The lower-set property therefore implies $c \in \Omega_h(\tau)$.

It remains to establish that the boundary belongs to the vulnerable set. Choose a sequence $c_n \in \Omega_h(\tau)$ with $c_n \uparrow s$. Continuity of v_h gives

$$v_h(s) = \lim_{n\to\infty} v_h(c_n) \geq \tau,$$

so $s \in \Omega_h(\tau)$. No $c > s$ can belong to the set, by the definition of s. Hence $\Omega_h(\tau) = (0, s]$, and s is an exact line. Uniqueness follows because two cutoffs generating the same lower interval must be equal. □

Appendix A.2. Proof of Corollary 1

Proof. Suppose first that v_h is weakly decreasing. For every τ, if $c_2 \in \Omega_h(\tau)$ and $0 < c_1 < c_2$, then

$$v_h(c_1) \geq v_h(c_2) \geq \tau.$$

Thus $\Omega_h(\tau)$ is a lower set. Whenever it is nonempty and bounded above, Proposition 1 yields an exact line.

For the converse, suppose that an exact line exists for every $\tau \in (0, 1)$ but that v_h is not weakly decreasing. Then there are $0 < c_1 < c_2$ such that $v_h(c_1) < v_h(c_2)$. Choose

$$\tau \in \Big(v_h(c_1), v_h(c_2)\Big).$$

The interval contains a value in $(0, 1)$ because both endpoints are probabilities and the inequality is strict. We then have $c_2 \in \Omega_h(\tau)$ but $c_1 \notin \Omega_h(\tau)$, contradicting the lower-set property that is necessary for a line at τ. Hence v_h must be weakly decreasing. □

Appendix A.3. Proof of Proposition 2

Proof. If a common line z^v_{real} exists, then for each household h

$$\Omega_h = \{c > 0 : v_h(c) \geq \tau\} = (0, z^v_{\text{real}}].$$

By uniqueness of the household-specific cutoff in Proposition 1, $z^v_h = z^v_{\text{real}}$ for every h. Conversely, if $z^v_h = z^v_{\text{real}}$ for every household, the defining equivalence for each household-specific line immediately gives the common-line equivalence. □

Appendix A.4. Proof of Proposition 3

Proof. The first equivalence in equation (9) is the definition of the household line. Since $z^v_h > 0$ and $z^v > 0$, the risk deflator is positive and, for every $c > 0$,

$$\begin{aligned} c \leq z^v_h \quad &\iff \quad c \leq RD_h z^v \\ &\iff \quad \frac{c}{RD_h} \leq z^v. \end{aligned}$$

This proves exactness.

For the converse, suppose a positive adjustment $d_h(\tau)$ and a positive common line $z^v(\tau)$ satisfy

$$v_h(c) \geq \tau \quad \iff \quad \frac{c}{d_h(\tau)} \leq z^v(\tau)$$

for every $c > 0$. The right-hand side is equivalent to $c \leq d_h(\tau) z^v(\tau)$, so $d_h(\tau) z^v(\tau)$ is a household-specific cutoff. Uniqueness of that cutoff gives $d_h(\tau) z^v(\tau) = z^v_h(\tau)$. □

Appendix A.5. Proof of Proposition 4

Proof. Suppose that A satisfies unanimity and additive subgroup consistency for every finite weighted population. Apply additive subgroup consistency to the partition into singleton households, $\mathcal{G} = \{\{1\}, \ldots, \{H\}\}$. For singleton

$g = \{h\}$, the group weight is $W_g = \omega_h$, its normalized within-group weight equals one, and unanimity gives $A((z_h^v), (1)) = z_h^v$. Hence

$$A(\mathbf{z}, \boldsymbol{\omega}) = \sum_{h=1}^{H} \omega_h z_h^v.$$

Conversely, let $A(\mathbf{z}, \boldsymbol{\omega}) = \sum_h \omega_h z_h^v$. If all individual lines equal a, then $A(a\mathbf{1}_H, \boldsymbol{\omega}) = a \sum_h \omega_h = a$, so unanimity holds. For any partition $\mathcal{G}$,

$$\begin{aligned}\sum_{g \in \mathcal{G}} W_g A(\mathbf{z}_g, \boldsymbol{\omega}_g) &= \sum_{g \in \mathcal{G}} W_g \sum_{h \in g} \frac{\omega_h}{W_g} z_h^v \\ &= \sum_{h=1}^{H} \omega_h z_h^v = A(\mathbf{z}, \boldsymbol{\omega}),\end{aligned}$$

so additive subgroup consistency holds.

Finally, because $z^v > 0$,

$$\sum_h \omega_h RD_h = 1 \quad \Longleftrightarrow \quad \frac{\sum_h \omega_h z_h^v}{z^v} = 1 \quad \Longleftrightarrow \quad z^v = \sum_h \omega_h z_h^v.$$

Thus the arithmetic reference and the mean-one normalization are equivalent. □

Appendix B. Additional empirical derivations

Appendix B.1. Lognormal poverty probabilities and quantiles

Under equation (22),

$$\begin{aligned}\Pr(C_{h,t+1} < z \mid X_h) &= \Pr(\log C_{h,t+1} < \log z \mid X_h) \\ &= \Phi\left(\frac{\log z - \mu_h}{\sigma_h}\right),\end{aligned}$$

which gives equation (24) after replacing the parameters by their estimates. The population lower τ-quantile $Q_h(\tau)$ solves

$$\Phi\left(\frac{\log Q_h(\tau) - \mu_h}{\sigma_h}\right) = \tau.$$

Solving gives

$$Q_h(\tau) = \exp\left[\mu_h + \sigma_h \Phi^{-1}(\tau)\right].$$

Replacing the parameters by their estimates yields equation (25).

Appendix B.2. Additive transition

For $C^A_{h,t+1}(c) = c - c_h + Y_h$,

$$\Pr(C^A_{h,t+1}(c) < z) = \Pr(Y_h < z - c + c_h).$$

At the cutoff, this probability equals τ. Since $\widehat{Q}_h(\tau)$ is the lower τ-quantile of Y_h,

$$z - c + c_h = \widehat{Q}_h(\tau),$$

and therefore

$$\widehat{z}^v_{h,A}(\tau) = z + c_h - \widehat{Q}_h(\tau).$$

Appendix B.3. Proportional transition

For $C^P_{h,t+1}(c) = (c/c_h)Y_h$ and $c, c_h > 0$,

$$\Pr(C^P_{h,t+1}(c) < z) = \Pr\left(Y_h < \frac{zc_h}{c}\right).$$

Equating this probability to τ gives

$$\widehat{Q}_h(\tau) = \frac{zc_h}{c},$$

and hence

$$\widehat{z}^v_{h,P}(\tau) = \frac{zc_h}{\widehat{Q}_h(\tau)}.$$

Appendix B.4. Alternative derivation of vulnerability lines

The proportional formulation in the main text yields positive vulnerability lines. An alternative is to hold absolute, rather than proportional, consumption changes fixed, but this can produce nonpositive boundaries. In this appendix, $\widehat{z}^v_h(\tau)$ denotes the boundary under this alternative specification. Let $\widehat{Y}_h$ have the same fitted lognormal distribution as in the main text and define the net loss implied by the model as $\hat{L}_h = c_h - \widehat{Y}_h$. Negative losses represent consumption gains. If the distribution of this net loss is held fixed when current consumption is evaluated at $c > 0$, future consumption is represented by

$$\widehat{C}_{h,t+1}(c) = c - \hat{L}_h = c - c_h + \widehat{Y}_h. \qquad \text{(Appendix B.1)}$$

The associated poverty probability is $\Pr(\widehat{Y}_h < z - c + c_h \mid \mathbf{X}_h)$. Since the fitted distribution is continuous and its lower quantile of order τ is $\hat{Q}_h(\tau) > 0$, the probability reaches τ when

$$z - c + c_h = \hat{Q}_h(\tau).$$

Solving for c gives the algebraic boundary

$$\widehat{z}_h^v(\tau) = z + c_h - \hat{Q}_h(\tau). \qquad \text{(Appendix B.2)}$$

When this boundary is positive, it is the household's vulnerability line under the additive formulation: among positive consumption levels, the household is vulnerable exactly when $c \leq \widehat{z}_h^v(\tau)$.

Positivity is not guaranteed. If $z + c_h - \hat{Q}_h(\tau) \leq 0$, there is no positive consumption level at which the household meets the vulnerability criterion. The algebraic boundary must then not be treated as a positive vulnerability line or used to construct a positive multiplicative risk deflator. The additive formulation can also assign positive probability to negative future consumption when $c < c_h$. These are domain limitations, not reasons to clip the boundary or silently discard the household.

Both formulations reproduce $\widehat{Y}_h$ at the observed c_h. Their classifications at that position consequently coincide:

$$\begin{aligned} \widehat{v}_h(c_h) \geq \tau \quad &\Longleftrightarrow \quad \hat{Q}_h(\tau) \leq z \\ &\Longleftrightarrow \quad c_h \leq \widehat{z}_h^v(\tau). \end{aligned} \qquad \text{(Appendix B.3)}$$

The equivalence at observed consumption remains an algebraic identity when the additive boundary is nonpositive, in which case both rules classify that household as non-vulnerable. Agreement at the observed position does not make the monetary lines identical: one construction holds absolute changes fixed, while the other holds proportional changes fixed. Only the proportional lines are used in the empirical comparison in the main text.

RECENT PUBLICATIONS BY *CEIS Tor Vergata*

Investment Mandates and the Value of Distribution Timing in Currency Portfolios Choice
Mattia Alfero, Stefano Grassiosimo Petracchi and Francesco Violante
CEIS *Research Paper*, 628 October 2026

Exchange-Rate Regimes and the Behaviour of Exporters
Cosimo Petracchi, Luca Riva and Marco S. Petterson
CEIS *Research Paper*, 627 August 2026

Pricing options on illiquid assets using liquid market benchmarks: an application to energy markets
Federico Aluigi, Lucia Caramellino, Paolo Pigato and Edoardo Scrima
CEIS *Research Paper*, 626 August 2026

Image Considerations in Second-Hand Adoption: A Vignette Experiment
Luca Congiu, Luisa Lorè and Mariangela Zoli
CEIS *Research Paper*, 625 July 2026

Communication Failure: The Hidden Face of The Tragedy of the Commons
Philippe Colo and Guillaume Pommey
CEIS *Research Paper*, 624 July 2026

Female Mayors and Public Procurement
Erminia Florio and Giancarlo Spagnolo
CEIS *Research Paper*, 623 July 2026

Optimal Structure of Penalties with Judgment-Proof Injurers
Guillaume Pommey
CEIS *Research Paper*, 622 June 2026

Flexibility Versus Security in Agency Contracts with Moral Hazard
Lorenzo Bozzoli and Guillaume Pommey
CEIS *Research Paper*, 621 June 2026

The Misallocation Costs of Inflation: A Sufficient Statistics Approach
Klaus Adam, Andrey Alexandrov and Henning Weber
CEIS *Research Paper*, 620 April 2026

Sensitivity of the Euro OIS Term Structure to ECB Policy Rate Surprises
Stefano Herzel and Marco Nicolosi
CEIS *Research Paper*, 619 January 2026

DISTRIBUTION

Our publications are available online at www.ceistorvergata.it

DISCLAIMER

The opinions expressed in these publications are the authors' alone and therefore do not necessarily reflect the opinions of the supporters, staff, or boards of CEIS Tor Vergata.

COPYRIGHT

Copyright © 2026 by authors. All rights reserved. No part of this publication may be reproduced in any manner whatsoever without written permission except in the case of brief passages quoted in critical articles and reviews.

MEDIA INQUIRIES AND INFORMATION

For media inquiries, please contact Barbara Piazzi at +39 06 72595584 or by e-mail at barbara.piazzi@ceis.uniroma2.it. Our web site, www.ceistorvergata.it, contains more information about Center's events, publications, and staff.

DEVELOPMENT AND SUPPORT

For information about contributing to CEIS Tor Vergata, please contact at +39 06 72595587 or by e-mail at segr.ceis@economia.uniroma2.it