

ON THE DEFINITION AND ESTIMATION OF ECONOMIC RESILIENCE USING COUNTERFACTUALS^{*}

M. Alloush[†]

Michael R. Carter[‡]

October 7, 2026

Abstract

This paper develops a measure of resilience based on the accumulated economic losses that a household experiences following a shock. Implementing this measure requires a clean counterfactual for what economic well-being would have been without the shock. Using data simulated from a dynamic optimization model that allow us to set aside identification concerns, we implement this measure and show it can be used to judge the efficacy of resilience-promoting interventions. The simulated data also allow us to explore the measure under a variety of scenarios, interventions and data generation processes. We show how the measure performs with and without poverty traps. In the presence of poverty traps, we find that when time series are short, consumption may be an unreliable indicator of resilience. Finally, we implement the measure using the World Bank's Living Standard Measurement panel data from three countries, as well as with data from two randomized controlled trials of interventions intended to improve the economic well-being of poor families. These applications shows that the financial returns to promoting resilience ex ante can be high and can save public resources compared to ex post, reactive social protection.

Keywords: Economic Resilience, Well-being Dynamics, Poverty Traps

JEL Codes: O12, J24, C63

^{*}Financial support from the BASIS Markets, Risk Resilience Innovation Lab and the Gates Foundation is gratefully acknowledged. We thank seminar audiences at Cornell University, Universidad Politécnica de Madrid, the University of Göttingen, the University of Washington and the World Bank for many helpful critiques and comments. We thank those who generously shared their work on resilience with which the authors were not previously familiar.

[†]Colorado State University, Department of Economics

[‡]University of California, Davis, NBER & University of Cape Town. Corresponding author, Email: mrcarter@ucdavis.edu.

1 Introduction

Economic resilience improves as the losses of current and future economic well-being that occur in the wake of a shock become smaller.¹ This paper develops a measure of economic resilience that is consistent with this loss-based conceptualization. Our resilience measure is a standardized and continuous measure that is based on the cumulative losses that a shock-exposed household experiences relative to a counterfactual measure of what its economic well-being would have been absent the shock.

The demand for a conceptually grounded measure of economic resilience has increased with the frequency and severity of climate shocks and other adversities. Governments and development agencies have launched a variety of policies intended to stabilize livelihoods in the face of shocks and to promote economic resilience. But absent a reliable measure of resilience, it is hard to gauge the efficacy of these policies and to evaluate whether or not the pursuit of resilience is wise public policy. The current situation is akin to pursuing poverty reduction policies without cogent measures of poverty or its depth. The measure we develop here is meant to fill this lacuna in the discussion around resilience. Importantly, our measure allows for the calculation of benefit-cost ratios that can be used to judge efficacy of programs and policies that build resilience.

Our analysis builds on an estimable measure of the cumulative economic losses that a unit (such as a household) experiences over time in the wake of a shock.² We denote cumulative losses over a time period T as $\mathcal{L}^T$, and define this term as the discounted cumulative difference between a household's shock and post-shock levels of economic well-being (y_t^s) compared to a counterfactual measure of what the household's well-being would have been without the shock (y_t^c):

$$\mathcal{L}^T = \sum_{t=0}^T \beta^t (y_t^c - y_t^s) \quad (1)$$

¹Such a definition has been common in policy discussions. The former US Agency for International Development often used such a definition ([US Agency for International Development, 2018](#)). Resilience has been similarly defined as the ability of countries, communities and households to manage change, by maintaining or transforming living standards in the face of shocks or stresses—such as earthquakes, drought or violent conflict—without compromising their long-term prospects ([Department for International Development, 2011](#); [World Bank, 2013](#)). In contrast, much academic literature has followed [Barrett and Constanas \(2014\)](#) and defines resilience in terms of the capacity to avoid poverty in the face of shocks and stresses, an approach which stirs together resilience with economic mobility.

²The measure and approach discussed here could be applied to units such as individuals or larger units of analysis, including communities, value chains, *etc.* Various measures of economic well-being, can be used, including assets, income, or consumption. As we later discuss, the choice between, say, income and consumption as the measure of economic well-being can be consequential in some cases whereas it is less so in others.

where β is the discount factor. We assume that the shock occurs in time period 0 and that we observe the household for T periods thereafter.³ Normalizing $\mathcal{L}^T$ by the counterfactual stream of income the household would have had absent the shock ($\sum_{t=0}^T \beta^t y_t^c$) provides the basis for an interpretable resilience measure that ranges between 0 and 1. A resilience measure of 0.6 would mean that the shock cost the household 40% of the discounted income stream that it otherwise would have enjoyed.

To illustrate the use of this loss-based metric, we first utilize simulation data created by a known data generation process, namely a dynamic stochastic programming model that we use to create data for households who choose consumption and asset levels in order to maximize their discounted stream of expected utility. These simulated data allow us to cleanly explore how the measure operates. The model also allows us to cleanly and ethically destroy productive assets of a random sample of households. The households that do not suffer the asset shock serve as reliable counterfactuals for what would have happened to the shocked households had they not suffered the shock.

Using these generated data, we implement an econometric approach that allows us to measure cumulative losses and implement the proposed resilience measure. Using the same data generation process, we then show how the resilience measure can be used to evaluate the economic efficacy of an asset replacement policy by assigning a subset of shocked households access to the policy and comparing their recovery trajectories to shocked households without access to the policy.

We then build on these key economic resilience concepts and extend our analysis to consider resilience in the presence of multi-equilibrium poverty traps. We do so by altering the data generation process in our simulation model to open it up to multiple equilibrium poverty traps (in the sense of [Barrett and Carter, 2013](#) and as analyzed by [Ikegami et al., 2019](#)). For reasons discussed below, when a subset of households fall into a poverty trap following a shock, population economic resilience declines only modestly, while the benefit-cost ratio attached to an asset replacement policy increases sharply.

We also show that in the presence of poverty traps, it matters whether we use income or consumption outcomes as the basis for calculating resilience, especially if the measurement time horizon is modest. In the short-term, households that fall into a poverty trap will have higher consumption but lower income than households that do not. Thus, households falling into a poverty trap may look more resilient using a consumption-based measure

³In practice, the counterfactual y_t^c will be estimated. Also, to account for the fact that shock households, especially those who have benefited from a resilience intervention, may build back better and stronger, we can more generally write equation 1 as $\mathcal{L}^T = \sum_{t=0}^T \beta^t \max[0, (\hat{y}_t^c - y_t^s)]$. Section 6.2.2 below discusses how to augment this measure to separately account for the building back better or resilience dividend phenomenon.

and less resilient with an income-based measure. This perhaps counterintuitive pattern emerges because households that fall into a poverty trap will optimally de-accumulate assets and temporarily boost their consumption on their optimal path to their new low level equilibrium position. This finding implies that the choice of outcome variable and the frequency and length of observation are inter-connected issues.

Finally, we present a number of empirical applications using our resilience measure. First, using a harmonized panel dataset that covers three sub-Saharan African countries, we estimate resilience following a drought shock and also show suggestive evidence that education and intercropping are correlated with higher levels of resilience. We then turn to two randomized controlled trials (RCTs) that allow cleaner identification of the impact of interventions on resilience. Using data from [Bandiera et al. \(2017\)](#), we show that a multi-faceted poverty-alleviation program increased resilience among poor women in Bangladesh. Building on [Balboni et al. \(2022\)](#), who document evidence of a poverty trap among these same women, we show that measuring resilience using consumption rather than income or assets can be misleading in the short run, mirroring our simulated results. We then use RCT data from a program that bundled financial and genetic technologies to build small farmer resilience in Mozambique and Tanzania. We find that the program led to a large increase in resilience and was highly cost-effective, even if the resilience gain is valued only in terms of humanitarian assistance avoided.

The remainder of this paper is structured as follows. Section 2 graphically presents the intuition behind our proposed measure of economic resilience. Section 3 then reviews the rich economic literature rooted in the permanent income hypothesis that explores the sensitivity of household well-being indicators to climate and other shocks. This literature lies at the heart of our proposed resilience measure. We also review other measures of resilience, including a growing literature that defines resilience without explicit reference to shocks. Section 4 uses simulation data from a known dynamic stochastic optimization model to estimate the resilience metric. This section also considers the impact of a publicly-provided asset replacement policy on resilience, and also shows how the resilience metric can be used to provide a benefit-cost measure of this policy intervention. Section 5 explores the behavior of our resilience measure in the presence of poverty traps, showing how traps can create divergent patterns of individual resilience and make the choice of economic well-being indicator crucial. Section 6 presents multiple empirical applications of our method, while Section 7 concludes.

2 A Cumulative Loss Measure of Economic Resilience

This section defines our cumulative avoided loss measure of economic resilience with simple graphical illustrations. The red solid line in Figure 1a plots the income trajectory of a household that experiences a shock in period 0 and only slowly recovers after that time. The blue dashed line illustrates the counterfactual trajectory that the shocked household would have followed without a shock. For illustrative purposes, we assume that this counterfactual trajectory is upward sloping as we might expect in a growing economy in which well-being improves over time. It could, of course, also be downward sloping or even flat.

The area $\mathcal{L}$ between the no-shock counterfactual trajectory and the shocked income path measures the overall cumulative income losses from the shock as defined by Equation 1 above. To define a clear quantitative measure of resilience, we normalize the discounted cumulative loss of well-being by the discounted stream of well-being that the individual would have had without the shock:⁴


$$\mathcal{R}^T = 1 - \left[\frac{\mathcal{L}^T}{\sum_{t=0}^T \beta^t (y_t^c)} \right] \quad (2)$$

where $\mathcal{L}^T$ is the cumulative loss measure defined in expression 1 above and β is the discount factor.⁵ The term in square brackets captures the share of loss in current and future well-being induced by the shock. When those losses approach zero, the resilience measure $\mathcal{R}^T$ approaches a value of 1, which would indicate full resilience. Lower values of $\mathcal{R}^T$ would signal larger cumulative losses with this measure bounded below by 0. The resilience measure $\mathcal{R}^T$ would only take the value of zero if individual economic well-being collapsed to zero and never recovered. Our measure is a ratio based on a counterfactual measure of economic well-being without a shock, and improvements in resilience come from increasing the Cumulative Avoided Loss; we will refer to the measure presented in equation (2) as the *CAL-Resilience Measure*.


Economists may recognize the similarity between $\mathcal{L}$ and the macroeconomic gap between actual and potential GNP, where an estimate of potential GNP serves as a measure

⁴This normalization requires that well-being is measured in levels, not in changes. In our applications, we use monetary measures that are identified in levels. If we specified our measure differently, then this normalization could be problematic.

⁵Note that in the spirit of Genicot and Ligon (2014)—see note 17 below—the measure in equation 2 could be straightforwardly transformed into an expected utility-based metric. Such a transformation would make losses more consequential at lower levels of well-being. State-dependent discounting would achieve the same thing. While these transformations would have some merit, it would lose the simple money metric comparability that we use when comparing resilience benefits with the cost of creating resilience. This weakness in turn could be ameliorated by transforming expected utilities into certainty equivalents.



(A) Elements for Resilience Measure



(B) Enhanced Resilience Two-ways

FIGURE 1: Cumulative Loss-based Resilience Measurement

of the counterfactual state.⁶ To the best of our knowledge, this earlier literature never labeled a smaller GNP-gap as signifying greater resilience, although the policy focus was very much on "automatic stabilizers" as policies that shrink the cumulative losses represented by the GNP gap. In a more recent discussion of macroeconomic resilience to weather shocks, Hallegatte (2014) formulates a measure akin to the GNP gap, labeling cumulative losses as a measure of resilience.⁷ His work includes a far-ranging discussion

⁶Okun (1963) appears to be the source for early discussions of the GNP gap. Black and Russell (1969) displays the Okun graph used to illustrate the gap between actual and potential GDP.

⁷In contrast to both the GNP gap literature and our approach here, Hallegatte (2014) measures cumulative

of the factors that can amplify a given shock to the economy’s capital stock and reduce resilience.

The *CAL*-Resilience measure, $\mathcal{R}^T$, also bears similarities to the ecological concept of resilience defined as speed of recovery to a known equilibrium position (Perrings, 2006). There have been a few economic studies that apply this ecological concept where the equilibrium position is proxied by the pre-shock level of the economic or welfare measure of interest. For example, Alfani et al. (2015) define a quantitative measure of resilience based on households who are hit by shocks but their pre-shock welfare is not very different from their post-shock welfare. Smith and Frankenberger (2022) similarly conceptualize “realized resilience” as the ability to maintain or improve food security after a shock, where again the implicit counterfactual is the household’s pre-shock level of food security. Zaharia et al. (2021) conceptualize resilience based on a household’s speed of recovery to its pre-shock position.

Although it may be appropriate in the context of ecological studies to consider the equilibrium to be static and known, such an assumption seems less appropriate when studying economic phenomena. As can be seen in Figure 1a, cumulative losses would be understated if they were calculated with respect to the pre-shock income level of the shocked. In our case, we do not think of the equilibrium position as static. Instead, we define resilience with respect to a counterfactual well-being time path. Only in the special case in which the counterfactual trajectory is flat will these before-after comparisons accurately measure the cumulative losses experienced by the shocked household.⁸

Interest in resilience measurement stems in part from a desire to gauge the efficacy of programs and interventions intended to reduce cumulative losses created by shocks. Imagine two groups of households, i and j , that suffer the same shock. The households are identical except that j benefits from access to a program or a market hypothesized to promote resilience, while i does not. We will say that unit j is more resilient than i if:

$$\mathcal{R}_j^T - \mathcal{R}_i^T > 0, \quad (3)$$

meaning that j has avoided some of the cumulative losses experienced by i .

losses relative to the pre-shock income level rather than to a counterfactual measure of what income would have been absent the shock.

⁸The logic here is of course the same as that which motivates the use of a control group in impact evaluation studies as opposed to simply doing a before-after comparison of the well-being of treated households. Similar to our work, and reflecting this same logic, Knippenberg and Hoddinott (2019) suggest that the resilience impact of an intervention (specifically Ethiopia’s PSNP social protection scheme) can be gauged by comparing the post-shock time paths of PSNP beneficiaries and non-beneficiaries.

While we will later show how this difference in resilience can be used as the basis for a benefit-cost measure of the resilience-promoting intervention, here we graphically illustrate the two mechanisms by which an intervention may bolster resilience. First, a household can become more resilient if it more quickly approaches its counterfactual state after having suffered the same initial impact from the shock as the reference household. This mechanism is illustrated by the orange, dot-dash, line in Figure 1b. The gap between the red and orange line is a measure of the cumulative avoided losses allowed by recovering more quickly.⁹ Insurance is an example of a market or intervention that might operate through this mechanism and we label this type of resilience as Type 1 resilience.

A household can also become more resilient if it is better insulated to begin with against the contemporaneous impact of a shock, as illustrated by the green, dotted line in Figure 1b. The household experiences less of a decline in well-being conditional on the same shock size. Use of stress tolerant seed varieties among farmers when facing a shock such as a drought is one example of a program that might improve resilience through this Type 2 resilience. These two types of resilience promotion are not mutually exclusive.¹⁰

The challenge is ultimately to estimate the ratio $\mathcal{R}^T$ and test the hypotheses of the form suggested by expression 3. After a brief review of the current resilience literature, the next section steps back and considers insights from the literature on the permanent income hypothesis, and the empirical literature it spawned, on how to estimate $\mathcal{R}^T$.

3 Resilience, the Permanent Income Hypothesis and Estimating Shock Responsiveness

The interest of governments and international organizations in policies that promote resilience gave rise to a literature on measuring “resilience capacities,” to use the term suggested by Barrett et al. (2021).¹¹ The Resilience Indicators for Measurement and Analysis (RIMA) put out by the UN Food and Agricultural Organization was the first such measure (Alinovi, Mane and Romano, 2008), and it has been further modified over time (d’Errico,

⁹Hallegatte (2014) makes similar observations about the mechanisms that promote resilience. He labels the two mechanisms dynamic and instantaneous resilience.

¹⁰Empirically, several recent randomized controlled trials have taken something akin to this approach by measuring a reduced-form impact of an intervention on household sensitivity to shocks (e.g., Macours, Premand and Vakis (2022) and Premand and Stoeffler (2022)). While these empirical studies do not have a standard that can be used to measure resilience, their basic empirical approach offers important insights into how panel data can be used to create a counterfactual against which resilience can be measured in experimental and quasi-experimental situations.

¹¹For more on the history of the term resilience in economics and other academic disciplines, see Béné et al. (2012), Alexander (2013), and Barrett and Constan (2014).

Garbero and Conostas, 2016). This approach treats resilience as a latent variable and measures characteristics or capacities hypothesized to lessen the impact of shocks and hence improve resilience as we have defined it above. For example, the Alinovi, Mane and Romano (2008) study uses this approach and measure resilience as a latent variable defined according to four main measures of well-being: income and food access, household assets, access to public services, and social safety nets. Smith and Frankenberger (2018) define resilience capacity as a set of characteristics that enable households to achieve resilience in the face of shocks. Their approach combines measures of three types of hypothesized resilience capacities (absorptive, adaptive, and transformative) into a single capacity measure using factor analysis.

While informative, these resilience capacity approaches neither consider the impact of shocks nor do they directly examine any measure of economic well-being. Cissé and Barrett (2018) push the resilience literature towards a focus on economic outcomes, defining resilience as the ability of households to avoid poverty over time as they weather whatever good and bad events come their way. Specifically, their measure is based on the probability of avoiding poverty, identifying which households are resilient in the sense that they avoid poverty in the future. But this measure does not directly consider the impact of shocks; instead, it relies on the estimated distribution of the error term from a well-being regression.¹²

While the poverty-centricity of Cissé and Barrett (2018) assures that their resilience measure is a pro-poor concept, it risks conflating resilience with poverty dynamics. Imagine two households that are both poor prior to being hit by a shock. One recovers quickly with a cumulative loss-based resilience metric approaching 1, while the other household fails to recover at all and becomes mired in even deeper poverty with an $\mathcal{R}$ measure well below 1. The Cissé and Barrett (2018) measure would say that neither is resilient as they both are poor in the long-run and yet from a conventional welfare perspective, the gains by the first household are large. While the measurement of poverty dynamics is important in its own right, we argue that it is important to distinguish resilience *per se* from poverty dynamics. Similarly, if we are not interested in the resilience of non-poor households, then we suggest focusing the measure of resilience only on poor and near-poor households

¹²Upton, Constenla-Villoslada and Barrett (2022) criticize both the resilience capacities and the Cissé-Barrett approaches for failing to predict out of sample well-being outcomes relative to past measures of well-being. Moreover, the approaches are inconsistent with one another in whom it identifies as resilient. In recent work, Lee, Abay and Barrett (2023) show that the method introduced by Cissé and Barrett (2018) is sensitive to the choice of well-being measure. More specifically, they show that measures of resilience using each of consumption, dietary diversity, or livestock assets as well-being indicators are only weakly correlated.

rather than shifting the meaning of resilience. Section 6.2.2 gives an empirical example of focusing the resilience measure only on the poor.

In contrast to this resilience literature, our measure of economic resilience defined above is based on observing shocks and the economic damage they cause, immediately and over time.¹³ The starting point for our approach is the very rich literature, spawned by the permanent income hypothesis (PIH) which explores the impact of shocks on household consumption and savings decisions. Unfortunately, the burgeoning resilience literature has thus far remained separate from the PIH shock-sensitivity literature. In this paper, we argue that by building on the PIH literature, we can arrive at a richer resilience metric that is descriptively useful and is a powerful tool for analyzing the impact of policy intended to promote resilience of households or other economic units.

Friedman (1957) put forward the permanent income hypothesis, positing that the consumption of credit-unconstrained households would respond very little to transitory shocks that lowered households' current income.¹⁴ Instead, Friedman argued that the household would spread out the impact of the shock over time (i.e., smooth consumption) by dissaving or borrowing as needed. In terms of the trajectories illustrated in Figure 1 above, the PIH suggests that consumption should barely respond to shocks and that income would show a contemporaneous but no lasting impact. Measured either with respect to consumption or income, cumulative losses induced by shocks would be minimal.

In an important addition to that literature, Deaton (1991) analyzed how the sensitivity of consumption to transitory shocks changes when households cannot freely borrow on a credit market in the wake of a shock. In the Deaton analysis, "impatient" households (those whose discount rate is no smaller than the rate of interest on savings) build up buffer savings stocks and draw down those stocks as needed to neutralize the negative impact of a shock. Even in those cases where savings are inadequate to fully buffer the shock, the impact is short-lived as Deaton assumes that the households earn wages and that the household returns to business as usual after one period.¹⁵ Again the cumulative losses induced by shocks would be small.

The striking implications of the PIH spawned a large empirical literature using data

¹³In recent work, Kilic et al. (2026) build on earlier approaches such as Cissé and Barrett (2018) and use machine learning methods to predict heterogeneity in shock effects, constructing a resilience measure that outperforms those analyzed in Upton, Constenla-Villoslada and Barrett (2022) at out-of-sample prediction of future food security. Unlike our approach, which is built around observed shock exposure and an explicit counterfactual well-being trajectory, their measure is designed to flexibly capture which households are systematically more or less able to absorb a shock, without specifying the source of heterogeneity in advance.

¹⁴A transitory shock is one that reduces current but not long-run or permanent income.

¹⁵Assuming shocks are not auto-correlated.

from both high and low income countries. [Paxson \(1992\)](#) is an early important example of the latter as she uses data on rainfall anomalies to estimate transitory income fluctuations. To directly test the PIH, she explores the impact of shocks on contemporaneous farmer net savings assets. Using data on farming households in Thailand, Paxson does not find evidence against the PIH and concludes that shocks are unimportant in terms of farmer well-being.¹⁶

Subsequent efforts to test the PIH in low-income countries are less sanguine about its veracity in environments with thin or missing markets. [Jalan and Ravallion \(1999\)](#), [Fafchamps, Udry and Czukas \(1996\)](#), and [Kazianga and Udry \(2006\)](#) all find some behavior consistent with the PIH. They find that at least for better off households, savings and assets are at least partially drawn down in the wake of shocks. But these studies also find substantial dissonant evidence in that many households (even those with positive amounts of assets) suffer larger consumption drops than the PIH predicts. Cumulative losses following a shock would thus seem to be much larger than predicted by the PIH, making resilience a more relevant consideration than Deaton's analysis would suggest.

An explanation for these dissonant empirical findings is that for agricultural and other self-employed households, assets are not simply accumulated to buffer shocks but also to generate income. Drawing down assets to smooth consumption this year thus has implications for the ability to generate income the following year. In Friedman's original formulation in which households can freely borrow against future earnings, the productive nature of assets held by poor households would matter little. It is less obvious that the same would be true in a credit-constrained setting and without the restrictive assumption that all households have wage-based income streams. In an early poverty trap model, [Zimmerman and Carter \(2003\)](#) show that consumption will not be smooth for poorer households when households can accumulate both buffer and productive assets that are necessary to produce income.

Building on these insights, [Carter and Lybbert \(2012\)](#) show that this "imperfect" consumption smoothing results from the fact that assets are necessary and productive to generate income (implying a different dynamic calculus than that deployed by [Deaton \(1991\)](#)) and that the non-convex production sets found in poverty trap theory (*e.g.*, [Ikegami et al., 2019](#)) will lead to dynamically optimal behavior that departs sharply from that predicted by the PIH. Using threshold econometric methods, [Carter and Lybbert \(2012\)](#) reanalyze

¹⁶([Paxson, 1992](#)) point estimate is 0.7, below the value of 1.0 predicted by the PIH. However, she cannot reject the hypothesis that the estimate is different 1, perhaps substantial heterogeneity in coping strategies as we discuss below.

the same data used by [Fafchamps, Udry and Czukas \(1996\)](#) and [Kazianga and Udry \(2006\)](#) and find that richer households tend to show behavior closer to what is predicted by the PIH, while households with scarcer productive capital are significantly more affected by shocks.¹⁷

Stepping back, the empirical analysis of the economic impact of shocks uncovers behaviors that suggest that shocks have much more negative and potentially longer-lasting impacts than those suggested by the PIH, at least for less-well off rural households in low-income economies. While most of that literature focuses on the contemporaneous impacts of shocks, this finding suggests that the longer-term effects (and the longer-term cumulative losses experienced by households) are likely to be substantially more complex and important than hypothesized by the standard PIH. In the next section, we develop a way to build on the empirical literature spawned by the PIH to investigate the full impact of shocks over time.

4 Estimating the *CAL* Measure of Economic Resilience

In order to clearly illustrate our resilience measure, we create noisy data generated by a known dynamic stochastic optimization model.¹⁸ [Appendix A](#) details this dynamic optimization model of occupational choice in which individuals (endowed with different levels of assets and entrepreneurial skill) choose between a low-income casual wage labor occupation and a higher-earning entrepreneurial occupation that requires the accumulation of productive capital. The model is general enough to encapsulate both poverty trap and non-poverty trap (e.g., [\(Roy, 1951\)](#) type) models. As explained further in [Appendix A](#) and [Appendix C](#), the model is solved numerically using the parametrization presented in [Janzen, Carter and Ikegami \(2021\)](#). Their parameterization is itself based on empirically informed parameters meant to represent East African pastoralist economies.

¹⁷ Building on the [Ligon and Schechter \(2003\)](#) concept of vulnerability, [Genicot and Ligon \(2014\)](#) suggest a measure of what might be termed consumption resilience, which they label insurance-adjusted risk. Their measure is the difference between the utility of permanent income and the expected utility of a stochastic consumption stream. Note that this difference would be trivial under the PIH, and even under Deaton’s analysis of consumption smoothing in the face of capital constraints. In contrast, [Genicot and Ligon \(2014\)](#) argue this difference between these two objects is substantial and an important source of lost human well-being. The [Genicot and Ligon \(2014\)](#) concept bears some resemblance to the *CAL*-Resilience measure with the utility of permanent income playing the role of counterfactual well-being in equation (1) above.

¹⁸The data are noisy in the sense that we introduce classical measurement error to measures of assets, income, and consumption that we extract from the results from the simulation. For each measure, we add normally distributed error with mean zero and a standard deviation that is 0.2 of the standard deviation of the variable in question. In reality, these measures may have different levels of measurement error as we discuss momentarily.

We add a modest dose of classical measurement error to the optimal choices generated by the dynamic model (see [Appendix C](#)). We do not adjust the magnitude of the measurement error based on whether we are looking at consumption, income or assets outcomes. In real data for poor populations, it is usually the expectation that income has the highest measurement error, followed by consumption and then with productive assets having the least measurement error. While in many cases these three measures would be expected to track each other, [Section 5](#) discusses the case of multiple equilibrium poverty traps in which these measures might diverge over a short to medium term.¹⁹ In the empirical applications of our measure in [Section 6](#), we deploy all three measures depending on the data and context.

Individuals in the model are credit-constrained (a la [Deaton, 1991](#)) and are subject to shocks that destroy capital, reducing cash-on-hand. Credit constraints bind consumption and investment choices, limiting their choices to be less than or equal to available cash-on-hand. The model assumes that individuals understand the probability that shocks occur and take that risk into consideration as they plot their expected utility-maximizing trajectory of consumption and savings. Depending on the shock received, individuals optimally adjust their consumption and savings trajectories moving forward. In this section, despite variation in skill levels among individuals, we fix parameter values such that it is always optimal for all individuals of all skill levels to accumulate capital and try to pursue the entrepreneurial occupation (see [Appendix B](#)). In this world, consumption, assets and income closely track each other as the recovery of the latter two well-being indicators occurs at a modest rate in keeping with individuals' preference for smooth consumption. In [Section 5](#), we consider an alternative world in which it may no longer be optimal for some households to pursue the higher income entrepreneurial occupation, depending on their endowments and history of shocks. In this world, the choice between welfare indicators is more consequential, as discussed below and illustrated empirically in [section 6.2.1](#).

By avoiding the messiness of real world data, this approach allows us to create well defined and ethical experiments by exposing randomly selected households to shocks and to policy interventions intended to promote resilience. Analysis of actual empirical data (see [Section 6](#) below) may be messier and requires more attention to assure that shocked or treated households are well-matched with their counterfactual counterparts. While RCTs can assure that exposure to an intervention is randomized and hence orthogonal to other

¹⁹The model we use is general enough to admit both ([Roy, 1951](#)) and multiple equilibrium poverty trap models (a la [Buera \(2009\)](#) and [Ikegami et al. \(2019\)](#)), as discussed in [section 5.1](#). We make no claim that poverty traps are everywhere, although even poverty trap skeptics (e.g., [Kraay and McKenzie \(2014\)](#)) find it likely that poverty traps do exist in locations such as the pastoralist regions of East Africa.

TABLE 1: Experimental Design

<i>Resilience Intervention</i>		<i>Shock Exposure</i>	
		Not Exposed	Exposed
	Not Treated	P_{00}	P_{01}
	Treated	P_{10}	P_{11}

characteristics, weather and other shocks are not subject to experimental control, making it incumbent on the analyst to probe the balance between shocked and non-shocked households, especially with regards to risk.²⁰

Table 1 shows the research design we employ in this section to develop our counterfactual based measures of economic resilience. In Section 4.1, we will focus on the experiment defined in the first row of the table, comparing the sub-population P_{01} exposed to a severe shock to an otherwise comparable sub-population P_{00} not exposed to the shock. The latter sub-population identifies the counterfactual economic trajectory that the shock-exposed population P_{01} would have experienced had they not received a shock.

Section 4.2 will consider the full 2×2 design shown in Table 1. Under this design, some units are experimentally offered an intervention that may increase resilience through either the Type 1 or Type 2 mechanisms illustrated in Figure 1b. This 2×2 design presents several candidates measures of the counterfactual, no-shock state. The resilience of the population of the treated, shocked P_{11} can be gauged via comparison to either the P_{10} or the P_{00} population counterfactuals. In later empirical applications, we will consider instances in which each comparison can be useful. Note that a purely resilience-building intervention (such as the contingent social protection program considered in Section 4.2 below) will not provide direct benefits to the treated, not-shocked population.²¹ In this case, the average trajectories of groups P_{00} and P_{10} will mirror each other, rendering the choice of counterfactual irrelevant.

4.1 The Cumulative Loss-based Measure of Population Resilience

This section develops a simple regression model that can be applied to panel data that includes measures of household well-being (income, consumption, or assets) that span a shock event that affects a subset of the households. In our artificial data, we are able

²⁰Balance on the risk of shocks is a one way to probe balance between shocked and non-shocked households.

²¹As will be discussed later, credibly increasing resilience may induce behavioral change—e.g., increased investment—that ultimately increases average income trajectory.

to apply the shock to a well-balanced random subset of households, so that the shock treatment is orthogonal to all variables, including those like skill that might be latent in real world data. In real data, fixed effects or other control variables might be required as discussed in the real world application in Section 6.

In the spirit of the literature on the permanent income hypothesis, we first define S_i as the binary treatment variable that takes on the value of 1 if household i is subjected to a severe shock in period 0. We assume data are available for τ periods before the period 0 shock and for T periods after the shock. Letting y_{it} represent an economic well-being measure for household i in time period t , we write the basic resilience regression model as:

$$y_{it} = \sum_{t=-\tau}^T \left(\beta_t^c d_t + \delta_t^s (S_i \times d_t) \right) + \varepsilon_{it}, \quad (4)$$

where there are $\tau + T + 1$ time periods in the panel data set, d_t is a vector of time-period binary variables and β_t^c is a vector of coefficients that will trace out the average well-being path for control households over time and δ_t^s is a vector of coefficients that will show the average difference over time for shocked households. While initial empirical investigations of the PIH such as Paxson (1992) examined only the immediate impact of shocks on say consumption or dis-savings, evidence that shocks are more consequential suggests the need for Equation 4's T -period extension of the short time horizons used in the PIH literature.

One feature of Equation 4 is that δ_t for $t < 0$ provide a simple balance test as those δ_t 's will be indistinguishable from zero if the shocked and non-shocked households are well balanced. Our sub-population estimates needed for cumulative loss measure Equation 1 are:

$$E[y_{it}|S_i = 0] = \hat{y}_t^c = \hat{\beta}_t^c d_t$$

and

$$E[y_{it}|S_i = 1] = \hat{y}_t^s = \hat{\beta}_t^c d_t + \hat{\delta}_t^s d_t.$$

To illustrate the use of this machinery, we now turn to data simulated using the dynamic stochastic optimization model outlined in Appendix A. We create data for two pre-shock periods and 9 post-shock periods. Figure 2 plots the estimates of $\hat{y}_t^c$ and $\hat{y}_t^s$ time paths from regression equation 4. The solid (blue) line traces out the average trajectory for non-shocked households (our counterfactual), whereas the dashed (red) line shows the same for households that suffered a severe drought in year 0, but did not experience any further losses. As can be seen in this controlled simulation experiment, the shocked and




FIGURE 2: Cumulative Income Loss from Shock

non-shocked households are well-balanced as their pre-shock time paths appear identical (*i.e.*, we cannot reject that hypotheses that $\delta_t = 0 \forall t < 0$).

Using equation 1 above, the estimated cumulative loss measure that captures current and future losses from the shock is $\hat{\mathcal{L}}^9 = \sum_{t=0}^9 \hat{\delta}_t$. Interpretation of $\hat{\mathcal{L}}^9$ as the causal impact of the shock on the current and future well-being on households of course depends on the usual orthogonality conditions between shocks and the error term. If shocks occur randomly and households are not spatially sorted by shock vulnerability (e.g., poorer households do not disproportionately live in, say, flood plains), then the non-shocked households are in fact a good counterfactual for the shocked households. While the structure of our controlled data generation process guarantees that these conditions are met, real world data of course requires greater caution, as we discuss in Section 6 below.

The total cumulative loss in economic well-being is represented by the cross-hatched area $\mathcal{L}^9$. Under this data generation process, total discounted cumulative losses are $\mathcal{L}^9 = \$2084$ per-household, where the superscript 9 indicates that in this case resilience is measured over nine post-shock time periods.²² Note that if households had been completely protected (say, by elaborate social protection schemes), then the area $\mathcal{L}$ would shrink to nothing. On the other hand, the less resilient are households, the larger the area $\mathcal{L}$ becomes.

Using the normalization in expression 2 results in a *CAL*-Resilience measure for this

²²This figure assumes that the income numbers given in the figure are measured in hundreds of dollars and that the discount rate is 5%.

population of 0.77, as reported in the first column of Table 2. As mentioned earlier, our approach to resilience has much in common with the ideas explored in Alfani et al. (2015), with the important exception that we offer a dynamic counterfactual that evolves over time (y_t^c) instead of assuming that the counterfactual for future time periods is the unit’s pre-shock level of economic well-being. As can be seen from the figure, projecting forward the pre-shock income level would substantially understate the cost of the shock, at least in the case of the data generated by our dynamic economic model in which income is growing.²³

4.2 Using the CAL-Resilience Measure to Evaluate Shock-responsive Social Protection

The prior sub-section shows how the CAL-Resilience measure can be used to characterize the average resilience of a population. Much of the interest in resilience measurement stems from the desire to evaluate policies designed to make households better-positioned to withstand shocks and avoid long-lasting negative impacts on household economic well-being that are visible in Figure 2. In this section, we “experimentally” introduce an unanticipated asset replacement policy that rebuilds assets for households following a severe shock. This policy can be considered as either a private insurance contract or a shock-responsive social protection scheme. To keep matters simple, we will assume that the policy is an unanticipated social protection scheme (for which the government purchases insurance) so as to rule out *ex ante* behavioral responses to the policy. The presence of this policy introduces the additional sub-populations defined in the second row of Table 1. We examine the impact of such a policy on resilience using our controlled data generation process.

In this analysis, we assume that the policy provides every household insurance against catastrophes that has the following characteristics:

- The policy pays nothing for shocks that destroy less than 40% of household assets;
- The policy pays half the value of any asset losses over and beyond 40%; and,
- Indemnities in the form of replacement assets are transferred one time period after the shock.

Returning to Figure 1b, this kind of policy would be expected to accelerate recovery (reducing future cumulative losses), while leaving unaltered the depth of the contempo-

²³In the simulated data, growth occurs because of continued capital accumulation by the non-shocked households who are below their steady state level of capital.




FIGURE 3: The Impact of Asset Insurance on Resilience

aneous fall from the shock itself.²⁴ As discussed above, under these assumptions there is only a single, no shock trajectory that can be used to gauge resilience.

Regression equation 4 can be extended to consider the full 2 by 2 experimental design shown in Table 1:

$$y_{it} = \sum_{t=1}^T \left(\beta_t^c d_t + \delta_t^S (S_i \times d_t) \right) + I_{it} \left[\beta_I + \sum_{t=1}^T \left(\delta_t^I d_t + \delta_t^{IS} (S_i \times d_t) \right) \right] + \varepsilon_{it}, \quad (5)$$

where I_{it} is a binary indicator variable that takes the value of 1 when sampled unit i is given the asset replacement policy.

To allow for a fair evaluation of the benefits and costs of this policy, we assume that the government has been buying a sovereign insurance contract to cover its contingent social protection liabilities for the entire population for a decade. Given that the severe loss events happen about 5% of the time in our simulation, this gives a fair representation of the cost of the insurance program relative to its benefits (with half the population receiving a shock once in 10 years). The present value of those public expenditures over the decade long-time span then stand as the measure of the cost of the program.

²⁴We here ignore the behavioral consequences of insurance discussed by Janzen, Carter and Ikegami (2021), which as they show can add substantially to the resilience-promoting impacts of this kind of insurance through a behavioral, investment incentive effect. We consider such impacts in Section 6.2.2 when we turn to analyze real world data in which such impacts can occur.

TABLE 2: Is Asset Insurance Good Policy?

	Standard DGP (No Poverty Trap)		DGP with Poverty Trap	
	No Insurance	Insurance	No Insurance	Insurance
<i>Mean CAL Resilience, R</i>	0.77	0.88	0.69	0.76
<i>Benefit Cost Ratio</i>		1.8		2.2

Figure 3 shows the impact of insurance on the income trajectory of households. As can be seen, it takes a season for the policy to rebuild household assets and assist the recovery of income. The shaded area marked $\mathcal{G}$ measures the resilience gain from the policy. That is, $\mathcal{G}$ is the reduction in cumulative losses induced by the policy. The immediate impact on income is quite substantial, but in later time periods, the uninsured households begin to catch back up in this single equilibrium convergence data set. As reported in Table 2, the resilience for the sub-population covered with the asset insurance policy rises from 0.77 to 0.88.²⁵

In order to evaluate the economic efficacy of this insurance policy, we calculate the present value of the resilience gain (the area $\mathcal{G}$ in Figure 3) and then compare it to the present value of the full public expenditure on insurance for the entire population (shocked or not). We further augment the cost by assuming that the government had been purchasing the insurance for the ten, shock-free time periods that preceded the shock. Using the probabilities in our underlying model, we can calculate the actuarially fair price of this insurance policy. We assume that the sovereign insurance policy is sold to the government at a 25% mark-up over the actuarially fair price.

Based on these assumptions, Table 2 shows that the benefit cost ratio for the insurance policy is 1.8, meaning that every public dollar spent on insurance-based social protection reduces cumulative losses by 1.8 dollars. Another way to think about this exercise is to imagine that the government has a shock-responsive social protection policy that returns shocked people to their counterfactual level of well-being. The benefit-cost measure indicates that every dollar spent by the government on an *ex ante* asset insurance policy would save \$1.8 in post-shock social protection expenditures.²⁶ This way of conceptualizing the

²⁵In real-world applications, complexity of response to the combination of treatment and a shock may make isolating the effect of an intervention on resilience more difficult. We show in one of our empirical examples in Section 6 that treated households exposed to a shock change their production choices in response that make them even better off in the long run.

²⁶We define resilience without regard to the absolute income level of the shocked household unit. Barrett and Constan (2014) and Cissé and Barrett (2018) have argued that conceptually resilience is primarily of interest

benefit-cost ratio makes clear the similarity between improvements in *CAL-Resilience* and the cash-benchmarking approach to evaluating the effectiveness of complex development interventions (e.g., [McIntosh and Zeitlin \(2024\)](#)). The resilience gain (area $\mathcal{G}$ in Figure 3) reveals the gross impact of insurance relative to an unconditional cash transfer intervention. With a benefit-cost ratio well above 1 (Table 2), these figures imply that the complex intervention is more cost-effective than a simple, *ex post* transfer intended to address the income shortfall created by the shock.

5 Poverty Traps and Resilience Measurement

This section explores how the *CAL-Resilience* measure behaves when some households are subject to what [Barrett and Carter \(2013\)](#) call multiple equilibrium poverty traps. As discussed in the literature, the presence of such traps can create permanent consequences as households that suffer sufficiently severe shocks will switch from an accumulation path towards the entrepreneurial occupation to a path to the lower-income (chronically poor), casual wage labor occupation (see [Ikegami et al., 2019](#) for example). This observation suggests that shocks will drive at least some households into a downward spiral such that the gap between the shocked and non-shocked will increase, instead of dissipating, over time. Using simulated data generated by a dynamic model that admits multiple equilibrium poverty traps, we first explore the implications of traps on a population-level measure of resilience and the returns to an asset insurance policy. Because the simulated data allow easy creation of a specific counterfactual for each shocked individual, we explore in more detail how poverty traps create divergent patterns of resilience within a population and how different measures of economic well-being (consumption versus income) can create different resilience outcomes over the short to medium term. While the creation of individual-level counterfactuals with real data confronts numerous statistical challenges, later analysis in Section 6 illustrates how insights from this analysis of individual resilience

when focused on low-income or poor households. Transforming the components of our resilience measure into their expected utility counterparts would automatically make the *CAL-Resilience* measure more pro-poor by giving greater weight to the resilience of poorer households (see note 5). However, some policy discussions around resilience ask if policies promoting resilience can pay for themselves by reducing the cost of social or humanitarian assistance that results from a resilience-promoting intervention. To accommodate an interest only in the resilience of the poor, we can focus the analysis on poor households without changing the approach to resilience measurement. When it comes to assistance, households with income below some normative value can be deemed eligible for humanitarian or other social assistance. A benefit-cost ratio based solely on humanitarian assistance avoided could thus be constructed using the present value of avoided losses below the poverty line. In section 6.2.2, we illustrate using the *CAL-Resilience* measure to measure benefits and costs when only humanitarian assistance avoided is valued as a benefit.

can guide analysis with real data when poverty traps are present.

5.1 Population Resilience in the Face of Multiple Equilibrium Poverty Traps

To explore the impact of multiple equilibrium poverty traps on resilience, we modify a key parameter of the model generating the simulated data such that the fixed time cost parameter of being an entrepreneur is strictly positive. As discussed in [Appendix B](#), this modest change in specification exposes a subset of middle ability individuals to multiple equilibrium poverty traps, meaning that if these individuals suffer a shock that makes them too poor, they will optimally settle into a lower income, non-entrepreneurial equilibrium.

This change in specification also means that low-skill individuals will never become entrepreneurs and will settle into a lower income wage labor occupation. While income for this subset of the population would thus be expected to be lower under this scenario (and their poverty higher), we measure resilience as a concept distinct from poverty dynamics.²⁷ In general, because the low ability subset of the population will operate with a much lower capital stock under this modified data generation process (*i.e.*, they are closer to the wage process imagined by [Deaton \(1991\)](#)), we might anticipate their resilience when measured against an appropriate counterfactual to be higher than under the data generation process considered above in [Section 4](#). At the same time, middle ability individuals now face the risk of falling into a poverty trap such that the impacts of the shock are long-lasting and irreversible, suggesting that average population resilience may decline. Finally note that individuals with high entrepreneurial skill will be largely unaffected by the change in specification between the no poverty trap and poverty trap cases.

Given this heterogeneity, it is unclear whether average resilience will be higher or lower in the presence of a poverty trap mechanism. The right-hand panel of [Table 2](#) shows that the presence of multiple equilibrium poverty traps modestly decreases population resilience from 0.77 (the no poverty traps case) to 0.69. While this 10% change appears to be small, repeating the analysis of the impact of an asset insurance policy for the poverty traps shows that benefit cost ratio of this resilience-promoting intervention rises 22% from 1.8 to 2.2. In present value terms, every dollar spent on promoting resilience with insurance saves \$2.2 in terms of reduced social protection expenditures. This jump in the value of insurance intuitively makes sense as the insurance can help middle ability households

²⁷ As discussed in [Section 3](#), some discussions of economic resilience that build on [Barrett and Constanas \(2014\)](#) risk conflating resilience with escape from poverty. Here we think it best to keep these two dynamic processes separate, especially as we need to be able to clearly evaluate what policies dedicated to promoting resilience do versus what they do not do. In other words, a policy may improve resilience but not resolve poverty (which may require asset transfers).

recover to asset levels that help them avoid falling into the low level equilibrium. In our data, nearly a third of agents in the simulation have an ability level that allows them to be prone to multiple equilibria. The benefit-cost ratio grows as this share increases. We now explore this further by deriving and analyzing individual-specific resilience measures.

5.2 Divergent Individual Resilience Trajectories in the Presence of Poverty Traps

To explore divergent resilience trajectories, we need to identify a well-matched counterfactual for each shocked individual. While individual counterfactuals are difficult to implement empirically (see [Appendix D](#)), our simulated data from a known model allows us to straightforwardly recover them. Doing so allows us to unpack further insight into how poverty traps generate divergent resilience trajectories.

Define $\hat{y}_{it}^C(y_{i-1}, \alpha_i)$ as the matched counterfactual time path for person i , where α_i is the individual's entrepreneurial ability and y_{i-1} is their immediate pre-shock level of well-being. In our simulated data, we are able to use exact matching based on pre-shock well-being and entrepreneurial ability. We can then define the cumulative individual loss for household i as:

$$\mathcal{L}_i^9 = \sum_{t=0}^9 \left(\hat{y}_{it}^C(y_{i0}, \alpha_i) - y_{it}^S(y_{i0}, \alpha_i) \right),$$

which is simply the cumulative losses for shocked household i relative to its specific matched counterfactual.

Converting this to the proposed resilience measure again requires an appropriate normalization. As with our population measure above, we normalize by the sum of non-shock income which is captured by the counterfactual income path for household i , yielding the individual-specific *CAL*-Resilience measure:

$$\mathcal{R}_i^9 = 1 - \left(\frac{\mathcal{L}_i^9}{\sum_{t=0}^9 \hat{y}_{it}^C(y_{i0}, \alpha_i)} \right).$$

Figure [4a](#) displays a histogram of the individual resilience metrics using our simulated data generated by a model that does not admit multiple equilibrium poverty traps. While the average resilience is 0.77 the individual measures range from 0.45 to 1.

In contrast, Figure [4b](#) shows the resulting distribution of individual resilience in the presence of multi-equilibrium poverty traps. Comparing Figures [4a](#) and [4b](#), we see the clear presence of poverty traps in the latter figure. A not inconsequential number of households exhibit near zero resilience as they fail to ever return to their matched counterfactual





FIGURE 4: The Distribution of Individual Resilience

trajectory. It is these households who have switched their equilibrium occupation and steady state income. At the upper end of the distribution, we also see many households with resilience measures at or near 1 in the poverty trap data generation process. These new highly resilient types are low ability individuals who are essentially wage workers who do not need assets to generate income and hence are able to recover their counterfactual living standard quickly. While their resilience is clearly a welfare improvement compared to no resilience, it should not of course be taken to mean that these households have escaped low incomes and poverty.


5.3 Consumption versus Income

In the presence of poverty traps, middle ability individuals are liable to fall into a position of economic non-viability following a severe shock (see [Appendix A](#)). As just shown, the resilience of these individuals approaches zero when gauging their economic well-being with incomes (or assets). But what if we were instead to measure resilience using consumption?

Because it is no longer always economically optimal for shocked middle ability individuals to aspire to accumulate and return to their original (see [Appendix B](#)), high equilibrium time path, they may find themselves with post-shock asset holdings that are above the steady state level for their now optimal, non-entrepreneurial equilibrium. Given these post-shock excess asset holdings, these individuals will optimally de-accumulate the excess assets. This asset de-accumulation allows them to temporarily boost their consumption as they descend toward their new long run equilibrium position. This logic suggests



(A) Income/Assets



(B) Consumption

FIGURE 5: Measuring resilience using consumption versus income may lead to different conclusions in the presence of poverty traps.

that in the short run these individuals may appear to be quite resilient if we look at their consumption levels, whereas their income will show the opposite pattern.

Figure 5 shows that this is indeed the case in our simulated data in the poverty traps case. Graph 5a shows that those who have fallen into a poverty trap (whose trajectory is graphed as the dotted, orange curve) have an immediate income decline as they shift towards their new, lower income level. In contrast, Figure 5b shows that if we were to examine the consumption of these same households, they would appear more resilient than their counterparts who did not fall into a poverty trap (at least for the first few post-shocks season). The latter are suppressing consumption in order to rebuild their asset base

and return to their higher income long-term economic equilibrium. However, past the fourth season, the true long-term lack of resilience of middle skills, trapped households becomes apparent.²⁸

While our controlled data generation process allows us to clearly see the divergent experiences and resilience of households, it raises a cautionary note about the economic outcome used to measure resilience and the time period over which we can measure that outcome. For shorter duration post-shock time series, this analysis suggests that income or assets will be the more reliable measure of resilience than consumption, at least in the presence of multiple equilibrium poverty traps. We return to this important issue in the empirical application in section 6.2.1.

6 Empirical Applications of the *CAL*-Resilience Measure

The preceding analysis has used simulated data from known data generation processes. While this approach allowed focus on formulating a conceptually clean measure of resilience, this section illustrates the empirical application of the *CAL*-Resilience measure using real-world data that are available at lower temporal frequency and/or with shorter duration time series than the simulation data. Our first empirical applications utilize longitudinal Living Standards Measurement Surveys (LSMS-ISA) data from three African countries. For these data, the time series is relatively long, but the temporal frequency is low (every two to three years). We then use data from randomized controlled trials (RCTs), which typically have higher frequency (seasonal) but shorter time series duration. For both types of data it is possible that shocks are not well-balanced as they are not subject to experimental control (see especially the discussion in section 6.2.2). While the observational LSMS data allow us to explore the association between resilience and household characteristics, the RCT data also allow us to causally identify the resilience impacts of specific interventions. In all cases, we adapt the econometric approach utilized with the simulation data to the specific empirical applications.

²⁸We can imagine other circumstances in which consumption and income trajectories may not track each other very well. One such circumstance (suggested by an anonymous referee) is when informal risk sharing offers help that can be used only for consumption support, but not for asset accumulation. This could occur when support is offered in kind and the giving neighbors would object if they observe the receiving households selling off food gifts to purchase productive assets.




FIGURE 6: Resilience among drought shocked farming households in Tanzania, Malawi, and Ethiopia

6.1 Measuring Resilience with LSMS-ISA Panel Data

Resilience is inherently a dynamic concept, and measuring it requires both shocks and longitudinal, follow-up data. Panel datasets are becoming more common around the world, we can increasingly link them to other data that allow us to measure specific shocks such as droughts. As a first empirical application, we use the World Bank's "harmonized" LSMS-ISA panel data from Tanzania, Malawi, and Ethiopia that are collected at two to three year frequencies (Bentze and Wollburg, 2025).²⁹ Because the harmonized LSMS surveys interview the same households across multiple waves and were designed with cross-country comparability in mind, they let us pool observations across three settings while still tracking each household's own trajectory before and after the shock.³⁰

For households that suffered a drought shock in the second survey wave, we use nearest neighbor matching methods to construct a control sample of un-shocked households whose trajectories can serve as a counterfactual estimate of the well-being trajectories for

²⁹The Ethiopian LSMS covers the years 2012, 2014, 2016, 2018 and 2021. The years covered by the Malawi data are 2010, 2013, 2016 and 2019. For Tanzania, the survey years are 2009, 2011, 2013, 2015 and 2019. In all cases, we carry out the analysis using shocks that occurred in the second survey wave, leaving one pre-shock wave to establish baseline and 3-4 post-shock waves (6-8 years) to study recovery.

³⁰In an approach that parallels the one taken here, Hansaki, Carter and Alloush (2026) use remote sensing data to measure shocks that occurred between LSMS survey rounds. They then regress the outcome each year on the set of lagged shocks experienced by each household, rather than focusing on only one cohort of individuals treated by a shock in a particular time period.

shocked households. Matching employed pre-shock agricultural characteristics such as plot size, main crop, crop yield, and demographic characteristics such as household size, age and sex of the household head and if they have had any formal education. As with any observational application of our measure, identification rests on the assumption that, conditional on the matching variables, exposure to the drought shock is as good as random—an assumption we view as reasonable here given that drought is a covariate shock driven by rainfall rather than by household-level characteristics correlated with recovery.³¹

Figure 6 shows that households experiencing a shock in wave 2 nearly recover by wave 5 (a 6-8 year time period), with an estimated *CAL*-Resilience score of 0.64. Beyond this average measure, the harmonized panel also lets us explore heterogeneity in resilience along dimensions that the resilience-capacities literature has hypothesized to matter—education and the adoption of intercropping. We find a resilience score of 0.79 for farmers with formal education versus 0.47 for those without, and 0.75 for farmers who intercrop versus 0.58 for those who do not; Figure A2 in Appendix E shows the corresponding figures. These are of course descriptive correlations and cannot be interpreted as an estimate of the characteristics’ causal impacts on resilience. Neither education, nor intercropping was randomly assigned, and both are likely correlated with other unobserved determinants of resilience, such as distance to a market or access to credit.³² What our measure offers here is not identification of these mechanisms, but a more direct and theoretically-grounded gauge of the outcome—actual resilience—against which such candidate correlates can eventually be assessed.

6.2 Measuring Causal Resilience Impacts using Data from Randomized Controlled Trials

Using LSMS panel data, section 6.1 showed that resilience differences can be measured for populations and for distinct sub-populations such as less versus more educated. While we could not claim that more education causes greater resilience (for the usual reasons), this section takes the next step and applies the *CAL*-resilience measure to two RCTs where such causal claims about the resilience impact of an intervention can be plausibly made. The first RCT application analyzes data from a multi-faceted poverty-reduction interven-

³¹This assumption is key to the application of this measure. Moreover, when considering how interventions (with potentially quasi-experimental assignment of treatment) may affect resilience, it is important to consider if selection into shock operates similarly across both groups.

³²It would of course be possible to deploy methods such as inverse probability weighting with regression adjustment (IPWRA—see [Słoczyński, Uysal and Wooldridge \(2022\)](#)) to try to address the bias potentially caused by the endogeneity of these characteristics.

tion. That intervention combines business asset transfers with training and support for extremely poor households in Bangladesh. Somewhat unusual for an RCT, the data for this study span 10 years, although the control group was treated after 4 years.³³

The second RCT application uses data from a bundled intervention of drought-tolerant maize and index insurance among smallholder farmers in Mozambique and Tanzania. A severe drought during the 2016/17 agricultural year affected both treatment and control households. A more conventional RCT, the data from this study cover only 3 years: a baseline and two post-treatment seasons.

6.2.1 The Resilience Impact of a Multi-Faceted Poverty-reduction RCT in Bangladesh

We use data from a randomized multi-faceted intervention, which gave randomly selected women livestock assets and skills training aimed at improving both their income levels and perhaps changing their resilience to shocks.³⁴ The data included a 2007 baseline, and standard follow-ups in 2009 and 2011. After 2011, an important component of the control group was treated. Additional surveys were carried out in 2014 and 2018. When comparing the resilience of treatment and controls groups, we employ only the first three rounds of the data.³⁵ As suggested in [Bandiera et al. \(2017\)](#) and later in [Balboni et al. \(2022\)](#), many treated individuals exhibit behavior consistent with a poverty trap, de-accumulating their assets and falling back toward a low equilibrium.³⁶ This gives us an opportunity to study the divergence of consumption-versus-income measures of resilience that we illustrated with simulated data in Section 5.3. For our analysis of trapped versus non-trapped members of the treatment group, we employ the data from all rounds.

First, we illustrate the impact of the graduation program on resilience. In Figure 7, we denote the baseline as wave 1. Following [Balboni et al. \(2022\)](#), we approximate the im-

³³RCT studies typically have a shorter time span as research ethics compel offering the treatment to the control group once the efficacy of the intervention has been established. As detailed in [Bandiera et al. \(2017\)](#) treatment was in fact offered in control areas following the second endline survey. To adjust for this, the authors deploy a number of assumptions to offset this contamination of the control group.

³⁴[Bandiera et al. \(2017\)](#) and [Balboni et al. \(2022\)](#) both utilize this data. The public version of the dataset does not have identifiable geographical location available to create a similar drought shock measure as in Section 6.1. We therefore rely on an imperfect measure of a shock which indicates a shock for those in the sample who experienced a decline in productive assets coming into Wave 2. [Balboni et al. \(2026\)](#) have connected this data to external measures of floods and droughts but the locational information needed to implement this approach not publicly available.

³⁵In their study, [Bandiera et al. \(2017\)](#) utilize a variety of methods to allow them to incorporate the 2014 into their analysis.

³⁶Using this same data, [Alloush and Carter \(2026\)](#) use an estimation developed to distinguish multiple equilibrium from single equilibrium poverty and find evidence that subjects of the [Bandiera et al. \(2017\)](#) confront the sort of multiple equilibrium trap considered here.

mediate post-treatment status as the sum of baseline assets plus the value of transferred assets. While imperfect, this approximation allows us to see the longer-term impact of the shocks that occurred between waves 1 and 2. As can be seen in the figure, the data indicate substantial gains in productive assets for the graduation program beneficiaries.³⁷ While such increases were of course the intention of the graduation program, we might also ask whether this shift of individuals from a wage labor to an entrepreneurial occupation increased or decreased resilience. In contrast to the pure resilience program analyzed in section 4.2, the relevant question of interest here requires that we compare the trajectory of treated, shocked households with the treated, no-shock households. As before, we construct a nearest neighbor match for those who experience a shock between the first and second wave.³⁸

As can be seen in Figure 7, untreated households—those without the assets and skills training—suffered more than the graduation program households when each is compared to their respective treated and untreated no-shock counterfactual. Untreated households suffer more from the shock and take longer to recover, while treated households are better cushioned from the outset and return more quickly to their counterfactual path. The estimated resilience measures are 0.47 for the untreated group and 0.58 for the treated group, indicating that the intervention improved households’ resilience as well as increasing assets. If we calculate resilience using log asset value, which could be seen as a utility transformation³⁹ that makes the resilience measure more pro-poor, the resilience of the treated group is 0.78 and that of the untreated group is only 0.25.

Balboni et al. (2022) document a poverty trap among households that fell below a certain asset threshold. Earlier, Bandiera et al. (2017) suggest that for over 30% of the treated sample, the treatment had no long-term effect on household assets.⁴⁰ This finding of extreme impact heterogeneity is consistent with the poverty-trap theory we discuss above. In Section 5.3, we show that in the presence of poverty traps, using consumption instead of assets or income can lead to different conclusions about a program’s resilience impacts in the short run. Households sliding toward a lower equilibrium may see consumption rise temporarily along the time path to their new asset position. To illustrate the importance of this consideration, we here crudely identify such households among the treated sample by

³⁷Bandiera et al. (2017) analyze program impacts on average levels of income and consumption as well as assets.

³⁸We match on available household and individual characteristics including asset levels in the baseline wave prior to wave one.

³⁹Footnote 5 above discusses the merits of using a concave utility transformation of our loss measure.

⁴⁰In a follow-up analysis, (Alloush and Carter, 2026) show that there are multiple equilibrium poverty traps in this data, not just the kind of bifurcation suggested by Roy (1951) type models.





FIGURE 7: Resilience among treated and untreated groups

tagging those whose asset levels are low in the last observed wave. All treated households received assets in wave 1; if a household's asset level was still low by wave four or five, it either suffered repeated shocks or sold off its assets along the way.


Figure 8 first shows asset trends over time for households that experienced a shock relative to a no-shock counterfactual. Figure 8a shows that both shocked groups saw assets decline immediately after the shock; from there, the groups diverge—one recovers, while the other continues, on average, to lose assets. Measuring resilience using asset values correctly identifies the recovered group as the more resilient one. Figure 8b shows however, that food expenditure among the group that failed to recover in the long run actually rose in the short run mirroring the simulated result from Section 5.3. In the presence of this kind of post-shock bifurcation, the estimated medium term resilience gain from an intervention will vary substantially depending on whether well-being is measured with assets or consumption.

6.2.2 The Resilience Impacts of an Insurance RCT in Mozambique and Tanzania

Boucher et al. (2024) implemented a randomized evaluation of an intervention that offered farmers a bundle of financial and genetic insurance technologies meant to bolster resilience. Against the backdrop of this controlled experiment, nature provided a natural experiment with drought and other shocks that affected farmers in both the control and



(A) Assets



(B) Consumption

FIGURE 8: Measuring Resilience with Consumption versus Income in the Presence of Poverty Traps

treatment groups, particularly during the middle year of the study (2016/17). In their Table 1, [Boucher et al. \(2024\)](#) show that the natural experiment is well-balanced across treatment and control groups, not only in terms of shock realizations, but more importantly in terms of long-term risk. The RCT in combination with this natural experiment fully populates the 2×2 research design in Table 1. Similar to most RCTs, the data are shorter duration than the simulated data used in section 4.2. Despite this limitation, we show that

substantial insights can be garnered by applying the CAL-Resilience metric.⁴¹

Using a variant of equation 4 above, Boucher et al. (2024) recover estimates of the contemporaneous and lingering future effects of severe shocks on control group farmers.⁴² Figure 9 graphs the estimated effect of the shock on the control group, where the counterfactual is the maize income level for control farmers who did not suffer the shock.⁴³ The solid red line shows the pre- and post-shock estimated average income of control households who received a severe drought shock, while the blue dashed line shows the expected income of counterfactual group of control households who neither experienced the drought nor received the insurance treatment. For those affected, the shock was large, resulting in a nearly 40% reduction in income and recovery was slow over the two-year time period we observe.⁴⁴ As analyzed in detail by Boucher et al. (2024), this slow recovery results from coping strategies that de-capitalized farms in the wake of the shock. For this population, the 3-year CAL-Resilience metric is $\mathcal{R}^3 = 0.64$, meaning that shocked farmers had a discounted stream of income over 3 years that is only 0.64 of the income enjoyed by non-shocked households. The total discounted cumulative losses of these households total \$993. These large accumulated losses and modest resilience measures are similar to those seen in the application in section 6.2.1.

In contrast to the solid red line for the experimental control group, the dotted green line in Figure 9 shows the estimated maize income path for households who suffered the shock but had been experimentally exposed to the resilience-building treatment of drought tolerant maize bundled with index insurance. This intervention is similar in intent to the asset replacement policy analyzed above in section 4.2. Empirically, the intervention had a small, but statistically insignificant impact on maize yields and income for the treated, no shock group. Interestingly, as discussed in detail by Boucher et al. (2024) these treated, no-shock households apparently concluded that they had wasted their money on the new technology. They subsequently dis-adopted the technology and behave almost identically to no-treatment, no-shock households and are estimated in period 3 to produce an almost identical maize income to these latter households. To keep matters simple and to avoid

⁴¹Full details on the technologies and the study are reported in Boucher et al. (2024).

⁴²Boucher et al. (2024) estimate yields and area cultivated separately. Here we use their data and estimation methods, but simply estimate total maize income—the product of yield and area—as a single variable. Our basic estimating equation is a straightforward application of equation (1) in that paper.

⁴³Because the study only featured data for three years (a baseline and two follow-ups), we add a fictive 4th year in which we imagine that shocked control group households had begun to recover. We take this liberty only to better illustrate the operation of the resilience measure.

⁴⁴The data from this study only allow good estimation of income one year after the shock and we here assume a partial recovery in the second post-shock year.




FIGURE 9: Evaluating the Impact of a Resilience Intervention for Maize Farmers

cluttering the figure, Figure 9 shows the cumulative loss area for treated, shocked households relative to the not-treated, not-shocked households (area $\mathcal{L}^T$). This reduction in loss (the resilience gain, $\mathcal{G}^T$) results because treated households did not fall as far given the same shock (which we labeled as Type 2 resilience in Section 2) and because they recovered immediately after the shock (Type 1 resilience). Indeed, as can be seen, treated farmers more than recover from the shock as their experience with resilience-building technologies leads to a subsequent intensification of the adoption of these technologies at both extensive and intensive margins (see Boucher et al. (2024) and the literature review in Lane (2024)).

This behavioral change by shocked, treated farmers illustrates the difficulty of isolating the first order impacts of an intervention that has ancillary or secondary effects. While in some applications these additional effects might cloud identification, in this case we see these additional effects as an intrinsic part of resilience-building interventions for an underinsured population that engages in income smoothing behavior by underinvesting in profitable inputs. We label the income impacts of this behavioral change the *resilience dividend*. To capture this dividend and separate it out from the pure resilience effect, we expand expression 1 into a total economic impact measure by comparing the economic well-being of treated, shocked households offered the resilience-building intervention ($\hat{y}_t^{sT}$) to the well-being of non-shocked households not offered the intervention ($\hat{y}_t^{nC}$):

$$\mathcal{T}^T = \sum_{t=0}^T \max \left[0, \left(\hat{y}_t^{sC} - \hat{y}_t^{sT} \right) \right] + \sum_{t=0}^T \max \left[0, \left(\hat{y}_t^{sT} - \hat{y}_t^{nC} \right) \right] = \mathcal{G}^T + \mathcal{D}^T, \quad (6)$$

where the first summation is simply the resilience gain of the treatment ($\mathcal{G}^T$) as defined earlier and the new term ($\mathcal{D}^T$) measures the resilience dividend for households offered the intervention.

Table 3 summarizes the key results of this analysis. Considering only the losses and ignoring the resilience dividend, the $\mathcal{R}^3$ resilience measure for the population of treated farmers jumps from 0.64 to 0.91, meaning that treated households suffered much more modest losses than control households that suffered a similar shock.

This CAL measure is based on the intention-to-treat estimates from the [Boucher et al. \(2024\)](#) study. Because only 50% of the farmers offered the insurance bundle adopted it, these numbers understate the resilience gains experienced by those farmers who adopted the bundle. The time paths graphed in Figure 9 are the estimates for those offered the treatments, not those who adopted it. The table also shows the benefit cost ratio for this intervention. While the resilience analysis covers only three years, considering the cost of the technology only over those 3 years would understate its cost as the kind of shock used for the analysis occurs approximately once every eight years. We thus assume that farmer had been purchasing the technology for 5 years prior to the shock, the shock occurred in year 6 and then the farmer intensified investment to reap the resilience dividend in years 7 and 8. The cost of the technology is measured at the discounted value of the additional input expenditures treated farmers would have incurred over this eight-year period. Note that both benefits and costs are calculated on an intention to treat basis. Following this procedure, the benefit-cost ratio for the intervention if we value the reduction in discounted cumulative losses is 2.6 to 1, meaning that every dollar spent on the technology returned \$2.6 in reduced losses. If we further add-in the discounted value of the resilience dividend (area $\mathcal{D}^T$ in Figure 9), the benefit cost ratio increases to 11.2 to 1.

Finally, as discussed earlier, some authors have suggested focusing resilience measurement only on the benefits provided to the poor. In that spirit, we follow [Janzen, Carter and Ikegami \(2021\)](#) and calculate the total change in government social protection expenditures under a hypothetical disaster response regime in which the government commits to using income supports to close the poverty gap for households hit by the shock. For this exercise, we assume that maize income is 50% of total household income, that each household has 5 members and that poverty or humanitarian assistance line is \$PPP 1 per-person, per-day. The first two assumptions are supported by the survey data underlying the [Boucher et al. \(2024\)](#) study. As shown in Table 3, the treatment households require just over \$500 less in humanitarian or social protection assistance. If we only value this impact of the insured seed intervention, the implied benefit-cost ratio is 1.8 to 1. While lower than the

TABLE 3: Resilience and the Resilience Dividend amongst Maize Farming Households*

	Experimental Status	
	Not Offered Insurance Bundle	Offered Insurance Bundle
Discounted Losses, $\mathcal{L}^*$	\$993	\$226
Resilience Measure, $\mathcal{R}^*$	0.64	0.91
Humanitarian Assistance Received (discounted)	\$561	\$42
Discounted Resilience Dividend, $\mathcal{D}^*$	0	\$2490
Benefit Cost Ratios		
Valuing Reduced Losses only	–	2.6
Valuing Humanitarian Assistance Avoided only	–	1.8
Valuing Reduced Losses plus Dividend	–	11.2

* Just under half the farmers offered the bundle of drought tolerant & insurance.

Reported estimates are ITT estimates, so the estimated impacts on the treated are twice as large as those shown in the table.

other benefit-cost ratio, this exercise shows that promoting resilience *ex ante* can dominate a reactive public policy response that simply provides social protection *ex post* to victims of agricultural shocks. In the spirit of [McIntosh and Zeitlin \(2024\)](#), who benchmark complex interventions against cash transfers, these results imply that the benefits of ex post cash transfers can be obtained by the studied ex ante resilience intervention at a fraction of the cost.

7 Conclusion

As climate change, epidemics, and conflict increase the frequency and severity of shocks, it becomes increasingly important to measure the resilience of households and the efficacy of programs and policies that improve resilience. To answer this need, this paper derives a resilience metric based on the cumulative economic losses caused by a shock. Drawing on the rich economics literature on how households respond to shocks, we propose a resilience metric based on the comparison between a shocked household’s actual economic trajectory with a counterfactual measure of what that trajectory would have been absent the shock. Using simulated data derived from known data generation processes, we show how these metrics can be estimated and how they can be used to gauge the efficacy of policies and programs intended to bolster resilience. The simulation data also allows us to explore the impact of shocks in the presence of multiple equilibrium poverty traps and importantly warns that using consumption trajectories (versus income or asset trajectories) to measure resilience can lead to misleading results when using short- to medium-term time series data.

Building on these findings, we present a set of empirical applications of our measure using several different types of data. We first use the longitudinal LSMS-ISA data from Tanzania, Malawi, and Ethiopia. We show that with modest assumptions, our measure allows us to gauge population resilience. While these data can be used to gauge the resilience of sub-populations (*e.g.*, the better educated or those who adopt certain production techniques) inferring causality from differences between groups confronts the usual identification concerns. Our second set of applications thus shows how resilience can be measured using RCT data and used to identify the causal resilience impacts of interventions. We use data from an impact evaluation of an asset-building graduation program in Bangladesh to show that the intervention increased resilience. Those data, collected from a population subject to poverty traps, also reveal the same divergence between consumption- and income-based resilience measures shown using our simulation data. Finally, we use data from a resilience-promoting agricultural intervention in Southern and Eastern Africa to show that not only did the program work, but it was highly cost effective. Moreover, the data reveal that this program changed post-shock behavior and generated a resilience dividend in which the newly resilient households ultimately surpassed the economic position to which they would have returned had they been merely resilient. Taking the resilience dividend into account, the resilience-based benefit cost ratio rises from 2.6 to 11-to-1. While these very large returns emerge from a particular program, they do show that a proper resilience measure allows us to identify instances in which resilience promotion is a very good public investment, and substantially cheaper than responding *ex post* to shocks with the cash transfers needed to restore people to their pre-shock positions.

References

- Alexander, David E.** 2013. "Resilience and disaster risk reduction: an etymological journey." *Natural hazards and earth system sciences*, 13(11): 2707–2716.
- Alfani, Federica, Andrew Dabalen, Peter Fisker, and Vasco Molini.** 2015. "Can we measure Resilience? A proposed method and Evidence from Countries in the Sahel." World Bank Policy Research Working Paper 7170.
- Alinovi, Luca, Ergdin Mane, and Donato Romano.** 2008. "Towards the Measurement of Household Resilience to Food Insecurity: An Application to Palestinian Households." In *Deriving Food Security Information From National Household Budget Surveys. Experiences, Achievement, Challenges.*, ed. R. Sibrian, 137–152. Rome:FAO.
- Alloush, Mo, and Michael R Carter.** 2026. "Empirical detection of poverty traps and unnecessary deprivation." working paper.
- Balboni, Clare, Oriana Bandiera, Robin Burgess, Anton Heil, Clement Mazet-Sonilhac, Munshi Sulaiman, and Yifan Wang.** 2026. "Weathering Poverty." Working Paper.
- Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil.** 2022. "Why do people stay poor?" *The Quarterly Journal of Economics*, 137(2): 785–844.
- Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman.** 2017. "Labor Markets and Poverty in Village Economies." *Quarterly Journal of Economics*, 132(2): 811–870.
- Barrett, Christopher B., and Mark Conostas.** 2014. "Toward a Theory of Resilience for International Development Applications." *Proceedings of the National Academy of Sciences*, 111(40).
- Barrett, Christopher B., and Michael R. Carter.** 2013. "The economics of poverty traps and persistent poverty: Empirical and policy implications." *Journal of Development Studies*, 49(7): 976–990.
- Barrett, Christopher B., Kate Ghezzi-Kopel, John Hoddinott, Nima Homami, Elizabeth Tennant, Joanna Upton, and Tong Wu.** 2021. "A scoping review of the development resilience literature: Theory, methods and evidence." *World Development*, 146: 105612.
- Béné, Christophe, Rachel Godfrey Wood, Andrew Newsham, and Mark Davies.** 2012. "Resilience: new utopia or new tyranny? Reflection about the potentials and limits of the

- concept of resilience in relation to vulnerability reduction programmes." *IDS Working Papers*, 2012(405): 1–61.
- Bentze, Thomas, and Philip Wollburg.** 2025. "A longitudinal cross-country dataset on agricultural productivity and welfare in Sub-Saharan Africa." *Scientific Data*, 12(1): 1843.
- Black, Stanley, and R. Robert Russell.** 1969. "An Alternative Estimate of Potential GNP." *Review of Economics and Statistics*, 51(1): 70–76.
- Boucher, Stephen, Michael R. Carter, Jon Einar Flatnes, Travis J. Lybbert, Jonathan Malacarne, Paswel Marennya, and Laura Paul.** 2024. "Bundling Genetic and Financial Technologies for More Resilient and Productive Small-scale Agriculture." *Economic Journal*, ueae012.
- Buera, Francisco J.** 2009. "A dynamic model of entrepreneurship with borrowing constraints: theory and evidence." *Annals of Finance*, 5(3-4): 443–464.
- Carter, Michael R., and Travis J. Lybbert.** 2012. "Consumption versus asset smoothing: testing the implications of poverty trap theory in Burkina Faso." *Journal of Development Economics*, 99(2): 255–264.
- Cissé, Jennifer Denno, and Christopher B. Barrett.** 2018. "Estimating development resilience: A conditional moments-based approach." *Journal of Development Economics*, 135.
- Deaton, Angus.** 1991. "Saving and Liquidity Constraints." *Econometrica*, 59(5): 1221–1248.
- Department for International Development.** 2011. "Defining Disaster Resilience: A DFID Approach Paper." DFID, London.
- d’Errico, M., A. Garbero, and M. Conostas.** 2016. "Quantitative Analyses for Resilience Measurement. Guidance for constructing variables and exploring relationships among variables." Food Security Information Network Technical Series 7.
- Fafchamps, Marcel, Christopher Udry, and Katherine Czukas.** 1996. "Drought and Saving in West Africa: Are Livestock a Buffer Stock?" *Journal of Development Economics*, 55(2).
- Friedman, Milton.** 1957. *A Theory of the Consumption Function*. Princeton University Press.
- Genicot, Garance, and Ethan Ligon.** 2014. "Risk and the Extent of Insurance." World Bank Background Paper for World Development Report, 2014.

- Hallegatte, Stephane.** 2014. "Economic Resilience: Definition and Measurement." World Bank Policy Research Working Paper 6852.
- Hansaki, Takashi, Michael R Carter, and Mo Alloush.** 2026. "Measuring Resilience with LSMS Data in Malawi." University of California Davis Working Paper.
- Ikegami, Munenobu, Michael R. Carter, Christopher B. Barrett, and Sarah Janzen.** 2019. *Poverty Traps and the Social Protection Paradox*. University of Chicago Press.
- Jalan, Jyotsna, and Martin Ravallion.** 1999. "Are the poor less well insured? Evidence on vulnerability to income risk in rural China." *Journal of Development Economics*, 58(1): 61–81.
- Janzen, Sarah A., Michael R. Carter, and Munenobu Ikegami.** 2021. "Can insurance alter poverty dynamics and reduce the cost of social protection in developing countries?" *Journal of Risk and Insurance*, 88(2): 293–324.
- Kazianga, Harounan, and Christopher Udry.** 2006. "Consumption smoothing? Live-stock, insurance and drought in rural Burkina Faso." *Journal of Development Economics*, 79(2): 413–446.
- Kilic, Talip, Marco Letta, Pierluigi Montalbano, and Federica Petrucci.** 2026. "CLARE: A causal machine learning approach to resilience estimation." *World Bank Policy Research Working Paper*.
- Knippenberg, Erwin, and John Hoddinott.** 2019. "Shocks, social protection, and resilience: Evidence from Ethiopia." *Gates Open Research*, 3(702): 702.
- Knippenberg, Erwin, Nathaniel Jensen, and Mark Constan.** 2019. "Quantifying household resilience with high frequency data: Temporal dynamics and methodological options." *World Development*, 121: 1–15.
- Kraay, A., and D. McKenzie.** 2014. "Do Poverty Traps Exist? Assessing the Evidence." *Journal of Economic Perspectives*, 28(3): 127–148.
- Lane, Gregory.** 2024. "Adapting to Floods with Guaranteed Credit: Evidence from Bangladesh." *Econometrica*, 94(2): 355–386.
- Lee, Seungmin, Kibrom Abay, and Christopher Barrett.** 2023. "Estimating Multidimensional Development Resilience." Working Paper.

- Ligon, Ethan, and Laura Schechter.** 2003. "Measuring Vulnerability." *The Economic Journal*, 113(486): C95 – C102.
- Macours, Karen, Patrick Premand, and Renos Vakis.** 2022. "Transfers, Diversification and Household Risk Strategies: Can productive safety nets help households manage climatic variability?" *The Economic Journal*, 132(647): 2438–2470.
- McIntosh, Craig, and Andrew Zeitlin.** 2024. "Cash versus kind: Benchmarking a child nutrition program against unconditional cash transfers in Rwanda." *The Economic Journal*, 134(664): 3360–3389.
- Okun, Arthur.** 1963. "Potential GNP: Its Measurement and Significance." *Proceedings of the Business and Economic Statistics Section of the American Statistical Association*.
- Paxson, Christina H.** 1992. "Using weather variability to estimate the response of savings to transitory income in Thailand." *The American Economic Review*, 15–33.
- Perrings, Charles.** 2006. "Resilience and sustainable development." *Environment and Development Economics*, 11(4): 417–427.
- Premand, Patrick, and Quentin Stoeffler.** 2022. "Cash transfers, climatic shocks and resilience in the Sahel." *Journal of Environmental Economics and Management*, 116: 102744.
- Roy, A. D.** 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers*, 3(2): 135–146.
- Sloczynski, Tymon, Derya Uysal, and Jeffrey Wooldridge.** 2022. "Doubly Robust Estimation of Local Average Treatment Effects Using Inverse Probability Weighted Regression Adjustment." *SSRN Electronic Journal*.
- Smith, Lisa C., and Timothy R. Frankenberger.** 2018. "Does resilience capacity reduce the negative impact of shocks on household food security? Evidence from the 2014 floods in Northern Bangladesh." *World Development*, 102: 358–376.
- Smith, Lisa C., and Timothy R. Frankenberger.** 2022. "Recovering from severe drought in the drylands of Ethiopia: Impact of Comprehensive Resilience Programming." *World Development*, 156: 105829.
- Upton, Joanna, Susana Constenla-Villoslada, and Christopher B. Barrett.** 2022. "Caveat utilitor: A comparative assessment of resilience measurement approaches." *Journal of Development Economics*, 157: 102873.

- US Agency for International Development.** 2018. "Resilience Evidence Forum Report." USAID Technical Report.
- World Bank.** 2013. "Building Resilience: Integrating Climate and Disaster Risk into Development." World Bank, Washington, DC.
- Zaharia, Sonia, William A. Masters, Shibani Ghosh, Gerald E. Shively, Sabi Gurung, Swetha Manohar, Andrew L. Thorne-Lyman, Keith P. West, Katherine H. Appel, Lichen Liang, et al.** 2021. "Recovery without resilience? A novel way to measure nutritional resilience in Nepal, Bangladesh, and Uganda." *Global Food Security*, 31: 100573.
- Zheng, Geyi, Michael R. Carter, Nathaniel Jenson, and Laurel Krovetz.** 2023. "Psychosocial Constraints, Impact Heterogeneity and Spillovers in a Multifaceted Graduation Program in Kenya." NBER Working Paper w31611.
- Zimmerman, Frederick J., and Michael R. Carter.** 2003. "Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints." *Journal of Development Economics*, 71(2): 233–260.

Appendices (For Online Publication)

Appendix A A Dynamic Model of Optimal Occupational Choice, Consumption and Accumulation

These appendices lay out a dynamic stochastic programming model that can be used to generate data on how households manage shocks, how quickly their income and consumption recover. It also allows us to generate counterfactual data by examining the trajectories of households that manage to avoid all shocks, but live in the knowledge that shocks could occur.

Consider an economy comprised of individuals each endowed with an initial level of wealth (k_{i0}) and a latent level of entrepreneurial skill (α_i), as suggested by [Buera \(2009\)](#). In this model, individuals can devote their resources to one of two different occupations:

- Casual Wage Labor which generates income $F_{jt}^w = w_0 + f^w(k_{it})$; or,
- Entrepreneurial Occupation which generates income $F_{jt}^e = (w_0 - A) + f^e(k_{jt})$.

We assume both income-generation functions are increasing and concave in k , that $f_e(k) > f_w(k) \forall k$ and that $A \leq w_0$. The parameter A is the time that must be withdrawn from the casual labor market in order become an entrepreneur.⁴⁵ Combining these two income functions yields a non-concave set with locally increasing returns to scale: $F(\alpha, k) = \max[F^w, F^e]$.

Following [Ikegami et al. \(2019\)](#), we assume that capital is subject to shocks and evolved according to:

$$k_{jt+1} = (k_{jt} + f(k_{jt}) - c_{jt}) (\theta_{jt+1} - \delta)$$

where c_{it} is consumption, $0 \leq \theta_t \leq 1$ is a random capital depreciation shock with known probability distribution function and δ is the standard, fixed rate of capital depreciation.

To study the dynamics of occupational choice and consumption dynamics, we assume

⁴⁵In the [Bandiera et al. \(2017\)](#) study program beneficiaries shifted from casual wage labor occupations to capital-dependent entrepreneurship.

that individuals solve the following inter-temporal maximization problem:

$$\begin{aligned}
& \max_{c_{jt}} \quad E_{\theta} \sum_{t=0}^{\infty} \beta^t u(c_{jt}) \\
& \text{subject to:} \\
& \quad c_{jt} \leq k_{jt} + F(\alpha_j, k_{jt}) \\
& \quad F(\alpha, k) = \max [F^w, F^e] \\
& \quad k_{jt+1} = (k_{jt} + f(k_{jt}) - c_{jt}) (\theta_{jt+1} - \delta) \\
& \quad k_{jt} \geq 0
\end{aligned}$$

where E_{θ} is the expectation taken over the distribution of the negative shocks and β is the time discount factor. $u(c_{jt})$ is the utility function defined over consumption and has the usual properties. Note that the final constraint reflects the absence of credit markets, placing this model in the [Deaton \(1991\)](#) world.⁴⁶ [Appendix C](#) below gives numerical values for parameters and shock distribution that underlie [Figure A1](#), including the assumption that $A > 0$, meaning a discrete amount of time must be withdrawn from the labor market to operate the entrepreneurial technology. The core parameters that we utilize are based on [\(Janzen, Carter and Ikegami, 2021\)](#), who in turn utilize empirically informed parameters meant to capture East African pastoralism. While the model's risk parameterization together with the assumption that $A > 0$ can generate multiple equilibrium poverty traps, the model is sufficiently general to encapsulate standard Roy [\(Roy, 1951\)](#) models of occupational choice. Indeed, our primary analysis in [Section 4](#) relies on a single equilibrium version of the model.

In order to draw out the implications of this model, we numerically solve the model for a wide array of initial asset positions over a number of randomly drawn shock sequences. Specifically, for each of 1500 initial positions evenly distributed across the initial endowment space shown in [Figure A1](#). The infinite horizon model was solved for each asset position, generating an optimal consumption value as well an optimal asset holding. A random shock was then generated, assets were updated and infinite horizon model was again solved for each updated asset position. This procedure was repeated 60 times, yielding a single history of consumption, income and assets for each initial asset position. At the end of each 60-year history, an indicator variable was formed indicating whether or not

⁴⁶Note also that this model assumes that capital is used for production and is not a strictly buffer asset. As such it matches the kinds of graduation programs analyzed in [\(Bandiera et al., 2017\)](#) which make productive asset grants to poor women to assist them in transitioning from wage labor to more remunerative entrepreneurial occupations.




FIGURE A1: Long-term Occupational Choice as a Function of Initial Endowments
Source: Adapted from (Zheng et al., 2023).

the individual was pursuing the wage labor or the entrepreneurial livelihood in period 60.

This entire process was then repeated 1000 times, generating 1000 histories for each of the 1500 initial endowment positions. The heat map in Figure A1 displays the probability that an individual at the indicated initial asset position will end up at the higher income entrepreneurial occupation across the 1000 histories. This procedure also generated a very large data set of observations on households with different skills, initial endowments and luck.

Examining Figure A1 we can see that the endowment space identifies three types of individuals based on their entrepreneurial skill endowment:

- Type A individuals with low skill endowments who will always move toward the casual wage-labor occupation and a poor standard of living irrespective of their initial endowment and shock history;
- Type C individuals with high levels of entrepreneurial skill will almost surely end up with sufficient capital to undertake the entrepreneurial occupation, even if they are born with zero initial capital;
- Type B individuals with intermediate skills levels whose long-term fate depends on

their initial capital endowments and history of shocks. If they are born too poor (below what [Ikegami et al. \(2019\)](#) call the Micawber Frontier), they will remain in the wage labor occupation. If they are begin with capital endowments above that frontier, they will attempt to become entrepreneurs, but may fail because of bad shocks, falling below the frontier and optimally remaining in the wage labor occupation.

Foreshadowing later discussion, note that only Type B individuals are subject to what [Barrett and Carter \(2013\)](#) call multiple equilibrium poverty traps.

Appendix B Generating Data with and without Multiple Equilibrium Poverty Traps

To study consumption dynamics in the absence of poverty traps, we set the fixed time commitment of being an entrepreneur to zero ($A = 0$). Under this assumption, all skill types will participate in the entrepreneurial livelihood. While the optimal steady state holding of capital is increasing with entrepreneurial skill, α , all households are converging toward an entrepreneurial equilibrium and there is no casual wage labor poverty trap. In this case, the equivalent of Figure A1 reveals high probabilities of the entrepreneurial occupation for all initial positions in the wealth-entrepreneurial skill space. We use the data from this no poverty traps data generation process in our primary analysis in Section 4.

We also solve the model with the fixed cost of being an entrepreneur set to be strictly positive ($A > 0$). Under this parameter value, Type B individuals are subject to multiple equilibrium poverty traps, as shown in Figure A1, while Types A and C are not. This heterogeneity in the applicability of multiple equilibrium poverty traps has implications for resilience measurement, as discussed in Section 5.

For both parameter specifications, we extracted samples of 10,000 households. Half of each sample was selected so that in season 4, the household received a substantial shock, destroying 40-60% of assets. In the other half of the sample, no such large shock was received in year 4. Histories were chosen such that no other large shocks occurred in any other season of the history. The sub-samples were also selected to be experimentally well balanced in terms of the distribution of skills and initial assets. A modest amount of classical measurement error was added to each variable.

In what follows, we will refer to the households that received the shock as the treated sample and households that did not receive the shock as the control sample. In other words, the control sample provides a balanced counterfactual for determining the present and future economic well-being of the treated sample had they not received a shock.

Appendix C Parameters for the Dynamic Model

TABLE A1: Functional Forms and Parameters used in Numerical Simulations

Production Technology and Parameters	
$F_{jt}^w = w_0 + k_{jt}^{\gamma_L}$ $F_{jt}^e = (w_0 - A) + \alpha_j k_{jt}^{\gamma_H}$ $\gamma_L = 0$ $\gamma_H = 0.56$ $A = 3.95 \text{ (0, for the no poverty trap case)}$ $w_0 = 3.95$	
Utility Function and Parameters	
Adaptive preferences utility function: $u(c_{it}) = \begin{cases} u^\ell(c_{it}) & \text{if } c_{it} < \tilde{c}(c_{g(i)}) \\ u^h(c_{it}) & \text{otherwise} \end{cases}$ Conventional preferences utility function: $u^l(c_t) = \frac{c_t^{1-\rho_l} - 1}{1-\rho_l}$ $\beta = 0.95$ $\rho_l = 0.75$ $\rho_h = 2.5$	
Distribution of Shocks	
The probability of θ_{jt} is assumed to be: $\text{density of } \theta_{jt} = \begin{cases} 0.3 & \theta_{jt} = 0.11 \\ 0.18 & \theta_{jt} = 0.021 \\ 0.13 & \theta_{jt} = 0.031 \\ 0.11 & \theta_{jt} = 0.041 \\ 0.10 & \theta_{jt} = 0.051 \\ 0.02 & \theta_{jt} = 0.061 \\ 0.01 & \theta_{jt} = \{0.071, 0.081, \dots, 0.191\} \end{cases}$	

Appendix D Individual resilience measures with real data

We first introduce our measure of resilience using household-level thought experiments in Section 2. We outline the need to recover a counterfactual income path to estimate our proposed measure. In Section 4, we show how experimental and quasi-random variation can be used to recover counterfactuals using averages at the group level to estimate average resilience. This is intentional—this type of as-if random group assignment allows us to identify the counterfactual: the no shock counterfactual and the untreated counterfactual both of which allow us to estimate the average resilience and the improvement in it due to treatment.

In section 5.2, we apply our approach to measure individual-level resilience with simulated data to illustrate how, in the presence of poverty traps, different measures of well-being may lead to different resilience outcomes. In our simulated data, we have full information which includes counterfactual income paths based on averages of non-shocked agents. We estimate individual measures of resilience easily because of this. In practice, estimating counterfactual income paths for households is not an easy task. The strength of our approach comes for estimating resilience at the group level; however, a lot of the resilience literature aims to estimate household/individual-level measures of resilience. This requires stronger assumptions. Below we make suggestions on how to estimate the no-shock counterfactual that projects the income path a household or an individual would have taken without a shock. We outline three different approaches to estimate this counterfactual:

1. **Matching:** Matching methods (such as nearest neighbor) could be used to estimate a counterfactual $\hat{y}_{it}^C(y_{i0}, \alpha_i)$ —either averages or maximum incomes could potentially provide an estimate for the counterfactual. Both this and the projection (next suggestion) approaches are likely to suffer from noisy estimates. Averaging over multiple nearest neighbors reduces some of this noise but this approach gets us a household- or an individual-level counterfactual only under the additional assumption of rank similarity across matched units, otherwise, matching only gives us a counterfactual distribution.
2. **Dynamic Panel Projection:** If panel data is available, either pre-shock or no shock data can be used to estimate a dynamic model which then can estimate the projected income without a shock for each individual in the data. These estimates can be used for $\hat{y}_{it}^C(y_{i0}, \alpha_i)$. Additional control variables can be used in the estimation equations


to reduce noise.

3. **Pre-shock Well-being:** The final, least noisy, approach is to use the income level prior to the shock as an estimate of the counterfactual. This should be adjusted for inflation or even general growth levels in the economy. This is similar to approaches used by [Alfani et al. \(2015\)](#) and [Knippenberg, Jensen and Conostas \(2019\)](#) who gauge the impact of shocks by comparing post-shock income to pre-shock income. Because it lacks a counterfactual, this approach assumes that nothing else would have changed following a shock. This approach is our least preferred approach.


When we apply methods 2 and 3 using our simulated data, these approaches give us individual measures of resilience that are relatively highly correlated with our original measure with the true counterfactual. Using dynamic panel projection for estimating our counterfactual (proposed method 2) gives individual resilience measures that have a correlation of 0.95 with the main true counterfactual measure. Using pre-shock well-being (proposed method 3) as a proxy for the counterfactual gives a resilience measures that is less correlated with our main counterfactual-based measure at 0.8.

Appendix E Resilience Heterogeneity in the LSMS-ISA Data

To estimate the resilience of different sub-populations using the LSMS-ISA data, we simply expand the model used for the core estimate in section 6.1 and include a binary variable for the characteristic and the characteristic interacted with the shock indicator. As can be seen, both households observed with formal education and those who employ intercropping techniques that are hypothesized to stabilize farm income show visibly higher levels of resilience reported in the main text.



(A) Education



(B) Intercrop

FIGURE A2: Measuring resilience using consumption versus income may lead to different conclusions in the presence of poverty traps.