On the Definition and Estimation of Economic Resilience using Counterfactuals

Mo Alloush and Michael R. Carter*

November 2023

Abstract

This paper derives a quantitative metric of economic resilience based on the cumulative current and future losses a shock-exposed household experiences relative to a counterfactural measure of what their economic well-being would have been absent the shock. Drawing on the rich economics literature on the sensitivity of household consumption and income to shocks, we derive a resilience metric that can be estimated with panel data using standard impact evaluation and matching econometric methods. To illustrate these methods, we rely on a dynamic optimization model to generate data from a known data generation process. By manipulating the parameters of the model, we are able to explore the robustness of our resilience metric to presence or absence of multiple equilibrium poverty traps. We also show how this metric can be used to not only evaluate the impact of a policy (catastrophic insurance) on resilience but also to judge the public finance efficacy of that same policy by showing how the cumulative–loss based resilience measure can be used for cost-benefit analysis. We also show that reliance on income as a measure of economic well-being may be wiser in the absence of long-term data. Finally, we use data from a recent experiment in Mozambique and Tanzania to show that these methods can be informative even with relatively short duration data.

*Alloush: Hamilton College & UCSD. Carter: University of California, Davis, NBER & University of Cape Town; Financial support from the BASIS Markets, Risk Resilience Innovation Lab and the Gates Foundation is gratefully acknowedged.

1 Introduction

Economic resilience can be defined as the ability of a household or other economic unit to manage a climate shock or other adversity with minimal compromise of current and future economic well-being. While this and related definitions of resilience[1] have a qualitative, or at least a quantitatively imprecise, element (what does "minimal" compromise mean?), the quantitative measurement of economic resilience has become increasingly of interest as the frequency and severity of climate and other shocks increases. Governments and development agencies have launched a variety of policies intended stabilize livelihoods in the face of shocks and promote economic resilience. But, absent a reliable measure of resilience, it is hard to gauge the efficacy of these policies and whether or not the pursuit of resilience is in fact wise public policy.

The goal of this paper is derive a quantitative metric of economic resilience consistent with the above definition of resilience. Our analysis builds on an estimable measure of the cumulative economic losses that a household (or other unit[2]) experiences over time in the wake of shock, defined as the difference between the household's post-shock levels of economic well-being (y_t^s) compared to a counterfactual estimate of what the household's well-being would have been without the shock ($\hat{y}_t^c$):

$$(1) \qquad \mathcal{L}_T = \sum_{t=0}^{T} \left(\hat{y}_t^c - y_t^s\right),$$

where some shock occurs in time period *0* and we observe the household for T periods thereafter.

To illustrate the use of this loss-based metric, we first utilize data created by a known data generation process, namely a dynamic stochastic programming model that we use to create data by households that manage their consumption and asset

[1]For example resilience has been 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" (DFID, 2011; Walker et al. 2004; World Bank, 2013). Barrett and Constas (2014) define resilience in terms of the capacity to avoid poverty in the face of shocks and stresses, an approach with combines economic mobility with resilience as defined here.

[2]The measure and approach discussed here could be applied to larger units of analysis, including communities, value chains, etc.

accumulation decisions to optimize their long-term, expected level of economic well-being. We arbitrarily assign some of these households a severe economic shock. Other households do not suffer the shock and their income and consumption trajectories serve as a counterfactual for what would have happened to the shocked households had they not suffered a shock.

Using these generated data, we first define an easy-to-interpret cumulative loss based measure of economic resilience that can be used to gauge the resilience of particular populations, including measurement of the full economic costs that accrue when a population is not resilient. Second, we then show how the resilience measure can be used to evaluate the economic efficacy of a catastrophic insurance program by assigning a subset of shocked households access to the program and tracking their recovery trajectories. Third, we show how individual specific resilience measures can be estimated, opening the door to the empirical investigation of the socio-economic and geographic characteristics that are associated with greater resilience.

Building on these key economic resilience concepts, we extend our analysis in several directions. First, we alter the data generation process to open it up to multiple equilibrium poverty traps (in the sense of Barrett and Carter (2013a) and as analyzed by Ikegami et al. (2019) and Janzen et al. (2021)). When a subset of households fall into a poverty trap following a shock, population economic resilience declines, the benefit-cost ratio attached to a catastrophic insurance policy jumps substantially and the individual resilience metric becomes negative for those households who settle into a low level equilibrium trap and never return to to their prior economic level of well-being.

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 short. In the short-term, households that fall into a poverty trap will have higher consumption but lower income than households that do not–that is the same households will look more resilient using a consumption-based measure 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 the new low level equilibrium position. This finding implies that discussions about the choice of outcome variable and the frequency and length of data collection are inter-connected issues.

In its final contribution, this paper uses results from a recent randomized controlled trial of resilience promoting technologies (stress tolerant seeds bundled with catastrophic insurance) to show that the proposed resilience measures can be implemented with real world data. Interestingly, the technologies in the real world data have a much more pronounced impact on resilience than the simulated catastrophic insurance program and have a much higher benefit-cost ratio (7:1 as opposed to 2:1 in the simulated data).

The remainder of this paper is structured as follows. Section 2 reviews the rich economic literature on 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 briefly review the largely disconnected resilience literature that primarily focuses on indexing capacities that have been ex ante posited to promote resilience.

Section 3 then employs data from a known dynamic stochastic optimization model (without poverty traps) to derive a resilience metric based on the time path of consumption of households subjected to a shock versus their counterfactual trajectory without a shock. This section also considers the impact of a simulated publicly-provided catastrophic insurance policy on resilience (it raises resilience from 43% to 70%), and also shows how the resilience metric can be modified to provide a benefit-cost measure of this policy intervention (the policy has a benefit cost ratio of almost 2:1). Finally, this section shows how the same methods can be used to estimate individual resilience measures that can be used to study the determinants of resilience or resilience capacities. Section 4 relies on a data generation process that admits multiple equilibrium poverty traps and illustrates the additional insights that emerge in this case of more complex income and asset dynamics. Finally, Section 5 illustrates the use of resilience measure using shorter-term, less frequent real world data, while Section 6 concludes.

2 Sensitivity to Shocks and the Measurement of Resilience

This section reviews the rich economic literatures on sensitivity of consumption to shocks, as well as the literature on economic resilience. Unfortunately, these two literatures have largely remained separate. In this paper, we argue that by unifying them, we can arrive at a richer resilience metric that is descriptively useful and a powerful tool for analyzing the impact of policy intended to promote resilience of households or other economic units.

Economics has a rich theoretical and empirical literatures that explore the impact of shocks on households' consumption and asset holdings over time. Theoretically, the permanent income hypothesis posited that the consumption of credit-unconstrained households would respond very little to transitory shocks that temporarily lowered household income (cite). 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. While not explicitly related to resilience, in retrospect this early literature in suggests a possible approach to measuring resilience by comparing a post-shock time path of consumption to a benchmark standard of what would be expected in a world of full and complete credit markets (a perfect markets counterfactual).

These theoretical ideas in turn gave birth to a stream of empirical literature (*e.g.*, Paxson (1992)) that tested whether household consumption and asset choices in the wake of shocks conformed to these theoretical expectations. As we discuss below, this empirical literature uncovered behavior that strayed far from the expectations of the permanent income hypothesis, or even Deaton's modification of that theory to account for credit constraints. A more recent empirical literature, spawned by the proliferation of impact evaluations of public policies, explores the impact of different policies on the sensitivity of household consumption to shocks (*e.g.*, Premand and Stoeffler (2022)). While largely divorced from the strict counterfactual suggested by the permanent income hypothesis, this literature's reliance on what are essentially a two-way experimental/quasi-experimental designs (households with and without the policy treatment, and with and without the natural shock) points the way toward the creation of a resilience measure based on comparing a shocked household's consumption trajectory with a well-defined and relevant counterfactual for what that trajectory would have been without the shock.

2.1 Sensitivity of Consumption and Income to Shocks

Both the classic, full and complete markets version of the canonical consumption model and Deaton's credit constrained variant are "one and done" models. In the canonical model, consumption falls only by the ratio $\frac{r}{1+r}$ (where r is the interest and discount rate) as the household borrows against future permanent income to deal with the deleterious consequences of shocks. Effectively, credit markets are used to smooth the impact of the shock out over the full household lifecycle. In the Deaton (1991) analysis, "impatient" households build up low returning buffer assets and draw down those assets to neutralize the negative of a shock as much as possible. Even in those cases where buffer assets are inadequate, the impact is one and done because the assume income generation process is a wage process and the household returns to business as usual after one period.[3] While this assumption faithfully represents some economies, it clearly is not an adequate representation of many parts of the developing world (rural and urban) where assets are needed not only to buffer shocks but also to generate future income. For example, Section 5 below (which builds on Boucher et al., 2022) show that transitory climate shocks have long lasting effects on future incomes as household balance the desire to smooth consumption with the need to preserve capital for future production periods.

Efforts to test whether or not consumption is smoothed (as predicted by the canonical model) or if savings covers the impact of shocks (as in the Deaton model) led to a very rich literature. Key papers here include Paxson (1992), Jalan and Ravallion (2002), Fafchamps et al. (1996) and Kazianga and Udry (2006). While these papers find some behavior consistent with the standard theory, they also find substantial

[3] Assuming shocks are not auto-corollated.

dissonant evidence that many households (even those with positive amounts of assets) suffer large consumption and income losses that are much larger than theory predicts. 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) and that the non-convex production sets found in poverty trap theory (e.g., Ikegami et al. (2019)) will lead to behavior that departs even more from that predicted by consumption smoothing theory. In essence they argue that consumption smoothing is not a goal, but is the optimal strategy to maximize inter-temporal expected utility under only rather specific circumstances.

2.2 RCT literature

Perhaps discouraged by the seeming lack of usefulness of theory to generate a standard for optimal response to shocks, a more recent empirical literature, based on RCT's has begun to more agnostically as if certain interventions (e.g., cash transfers) lessen household sensitivity to shocks. Important examples of this literature include Macours et al. (2022), Premand and Stoeffler (2022), etc. While these approaches 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 counterfacutal against which resilience can be measured in experimental and quasi-experimental situations.

2.3 The Resilience Literature

As a concept, resilience is not new: it has been applied in ecology, engineering, and some social science fields for decades. Each has defined and measured resilience to fit its goals, but overall, it is seen as the ability of systems to absorb shocks or changes and persist. Given the increased focus in economics and specifically in development economics on resilience, Barrett and Constas (2014) argue that defining the concept within this field is important. In their paper, Barrett and Constas (2014) define resilience as "the capacity over time of a person, household, or other aggregate unit to avoid poverty in the face of various stressors and in the wake of myriad shocks. If and only if that capacity is and remains high over time, the unit is resilient." Our approach differentiates resilience from poverty and poverty dynamics. We define resilience as the ability to maintain well-being as close as possible to the no-shock counterfactual most closely following the definition of resilience proposed by Pimm (1984) that is based on the speed of return to an equilibrium in the context of ecosystem stability. With a few exceptions, existing measures of development resilience do not tie measurement to a counterfactual. Our proposed measure differs significantly from these, we highlight three main ways resilience is currently measured in the development economics context including methods used extensively among development oriented policy institutions.

Barrett et al. (2021) provide a clear review of the literature on development resilience. They divide up quantitative methods of measuring resilience into three broad approaches: (1) resilience as a capacity; (2) resilience as a normative condition; and (3) resilience as a return to equilibrium. Currently, the most common measurement tool is the first: resilience as a capacity. This approach treats resilience as a latent variable and in practice uses a set of indicators thought to capture resilience creating an index thought to capture ex ante resilience capacity. Proponents of this approach consider resilience multi-faceted and cannot be observed in one dimension but is related to a number of context-specific dimensions that we can observe. These observable variables can be reduced to a single (or multiple) dimension(s) using factor analysis. For example, Alinovi et al. (2008) use 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. A weakness of this approach is that it never actually measures resilience and assumes that we already know what generates it. In fact, in practice what these measures capture are explanatory variables rather than outcomes—this makes this approach particularly unsuitable to evaluate the impact of policies or programs specifically intended to improve resilience.

Another thoughtful latent variable approach that has gained ground among development agencies explicitly abstracts away from measuring resilience itself and instead defines resilience capacity as a set of conditions that enable households to achieve resilience in the face of shocks. Smith and Frankenberger (2018) recognize three types of resilience: absorptive capacity, adaptive capacity, and transformative capacity. Latent variables on resilience capacity are then measured by creating three indices using factor-analysis of relevant observable variables at the individual, household, and community levels. We note that this approach does not measure resilience *per se* but factors correlated with what one might think enables people to be resilient in the face of shocks. Our goal is to provide a measure of resilience itself which one can then use to see what individual, household, or location characteristics lead to higher levels of resilience.

The second way resilience is conceptualized and measured is as a normative condition where the concept is tied directly to a pre-determined well-being standard.Cissé and Barrett (2018) propose a conditional moments approach to measure resilience defined as the ability to avoid low levels of well-being over time. Their measure, based on the probability of avoiding some pre-defined level of well-being (for example a poverty line), can be used to identify which households are resilient. The probability of avoiding falling below some pre-defined well-being threshold measure is conceptually appealing and their measure can be aggregated to create general population or group-level measures of resilience. Cissé and Barrett (2018) further argue that their proposed measure of resilience is both forward-looking and allows for nonlinear well-being dynamics. Moreover, it ensures that resilience is a pro-poor concept which Barrett et al. (2021) argue should be a priority of any resilience measure. However, we find that this approach confounds poverty measurement with resilience and does not necessarily speak to speed of recovery to no-shock counterfactuals and how it affects resilience. Wealthy households that never recover to their counterfactuals but can avoid very low levels of well-being are just as resilient as those who fully recover. Futhermore, poor households that recover to their pre-shock equilibrium faster than other similarly poor households are more resilient, and giving decreasing marginal utility of wealth, more increasing resilience (as measured by speed of recovery) among poor households can have important welfare implications. Our proposed measure clearly differentiates resilience from measurement of poverty in the steady states.

A notable downside to these two approaches is that neither one does well in predicting out of sample well-being outcomes relative to, for example, simply using lagged well-being (Upton et al., 2022). Moreover, the approaches are inconsistent with one another in who it identifies as resilient. In recent work, Lee et al. (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 in which households are identified as resilient.

Finally, the third approach relates resilience to the speed of recovery or return to equilibrium. Our study is most directly related to this conceptualization which is popular in ecology (Perrings, 2006). Within the development resilience literature, a few studies conduct pre- and post-shock comparisons with speed of recovery in mind and are thus similar to our proposed measure. 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 what they define as *realized* resilience as the ability to maintain or improve food security after a shock. Moreover, Knippenberg and Hoddinott (2019) propose an approach to measure how a program increases resilience that implicitly uses a no-program counterfactual well-being path while Knippenberg et al. (2019) use high frequency data and subjective shock persistence measures to compare across households where, conditional on the same shock, households who show lower levels of shock persistence over time are considered more resilient. Zaharia et al. (2021) propose an approach that measures resilience as an asymmetric mean reversion conceptualizing it on speed of recovery. These approaches most closely think of resilience in the way we define and measure it in this paper.

3 A Counterfactual-based Measure of Economic Resilience

In order to clearly develop our resilience measure, we create (noisy[4]) artificial data generated by a known dynamic stochastic optimization model. The appendix below details this dynamic optimization model of occupational choice in which individuals choose how much of their available wealth to consume versus how much to save and invest in order to improve their expected economic well-being over time by becoming entrepreneurs as opposed to casual wage workers. Individuals face risk that a shock will occur and destroy some of their wealth. The model assumes that individuals understand the probability that shocks occur and take that risk into consideration as they plot their optimal trajectory of consumption and savings. After making their optimal choices, individuals are exposed to shocks that destroy wealth. Depending on the shock received, individuals optimally adjust their consumption and savings trajectories moving forward.

In undertaking this simulation analysis, we consider an economy comprised of a large number of individuals, each born with a different levels of inherited wealth as well as different levels of entrepreneurial skill. For this section, we fix parameters of the underlying model such that it is optimal for all individuals (rich or poor, high skill or low skill) to try to accumulate assets and shift from wage labor to a higher returning entrepreneurial occupation. Section 4 will modify the parameters of the model in order to generate data where at least some households are subject to multiple equilibrium poverty traps in the sense of Barrett and Carter (2013b) and Ikegami et al. (2019).

While obviously avoiding the messiness of real world data, this approach allows us to create well defined, ethical experiments by exposing randomly selected households to shocks and, or to policy interventions intended to promote resilience. Table 1 shows the research design we employ in this section to develop our key counterfactual-based

[4] The data are noisy in the sense that we introduce classical measurement error to the results from the simulated data.

Table 1: Experimental Design

		Shock Exposure	
		Not Exposed	Exposed
Resilience Intervention	Not Treated	P_{00}	P_{01}
	Treated	P_{10}	P_{11}

measures of economic resilience. In section 3.1, 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. In section 3.2, we will show how these same data can be utilized to recover individual-specific measures of resilience, which can in turn be used to identify the factors that make some individuals more resilient than others.

After using this first two populations to define resilience measures, section 3.3 introduces the two additional randomly selected sub-populations shown in the second row of Table 1. Both groups receive a publicly funded catastrophic insurance (or contingent social protection) program, whereas only sub-population P_{11} suffers a large economic shock. Using this additional data, we will show how to use the resilience measure to measure the impact and the cost effectiveness of the catastrophic insurance.

3.1 The $\mathcal{L}r$ Measure of Resilience

We first consider a simple regression model that can be applied to panel data that includes measures of household well-being (consumption, assets or income) that span a shock event that affects a subset of the households. In our artificial data, we are able to apply the shock to a well-balanced random subset of households, so that the shock treatment is expected to be orthogonal to all variables, latent or otherwise. In real data, fixed effects or other control variables might be required.

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 analyze the data from that period and the 9 subsequent time periods or seasons. Letting y_{it} represent an economic well-being measure for household i in time period t, we write the basic resilience regression model as:

$$(2) \quad y_{it} = \sum_{t=0}^{T} \left(\beta_t^c d_t + \delta_t^S \left(S_i \times d_t\right)\right) + \varepsilon_{it},$$

where there are $T+1$ time periods in the panel data set, d_t is a vector of time period binary variables and β_t^C and $\delta_t^S$are vectors of coefficients for control and treated (shocked) households, respectively. As is obvious from this simple structure, our sub-populations estimates are

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

and

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

Using equation 1 above, the estimated cumulative loss measure that captures current and future losses from the shock is $\hat{\mathcal{L}^{10}} = \sum_{t=0}^{9} \hat{\delta}_t$. Interpretation of $\hat{\mathcal{L}}^{10}$ 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 flood plans), then the non-shocked households are in fact a good counterfactual for the shocked households. While structure of our data generation process guarantees that these conditions are met in this simple, real world data of course requires greater caution, as we discuss in Section 5 below.

Figure 1: Cumulative Income Loss from Shock

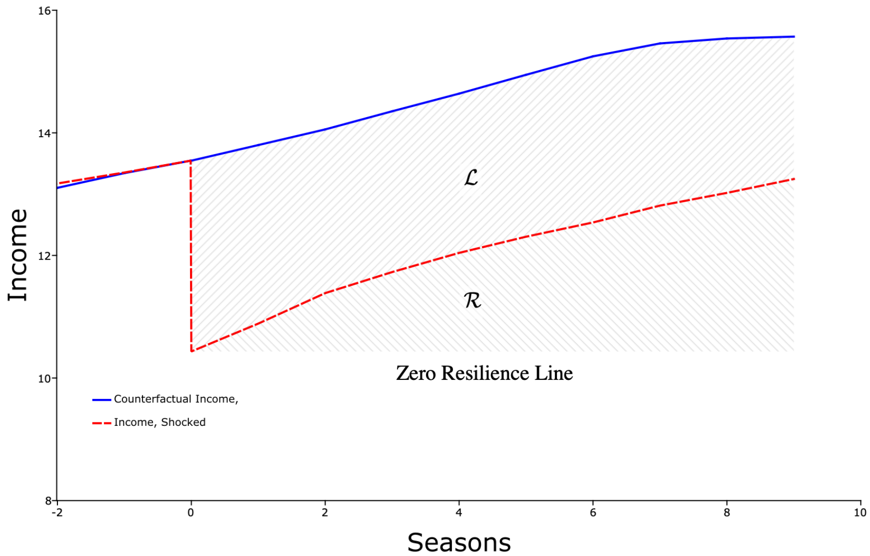

Figure 1 plots the the estimates of $\hat{y}_t^c$ and $\hat{y}_t^s$ from regression equation 2. The solid or blue line traces out the counterfactual trajectory for non-shocked households, whereas the dashed or red line shows the same for households that suffered a severe drought in year 0, but did not experience any further losses. The total cumulative loss in economic well-being is represented by the cross-hatched area $\mathcal{L}^{10}$. Total cumulative losses are $\mathcal{L}^{10} = \$2675$ per-household, and their discounted present value[5] is $\mathcal{L}'^{10} = \$2084$, where the superscript 10 indicates that in this case resilience is measured over a ten year time period (the shock year plus the nine proceeding periods).[6] Note that if the household had been completely protected (say, by elaborate social protection schemes), then the area L would shrink to nothing. On the other hand, the less resilient the household, the larger the area L would become.

In order to create an interpretable welfare metric out of loss measure $\mathcal{L}$, we normalize the cumulative economic loss caused by the shock by the income trajectory

[5]Define the discounted present value of current and future losses from a shock as:

$$\mathcal{L}' = \sum_{t=1}^{T} \beta^t \left(y_t^c - y_t^s\right)$$

[6]These figures assume that the income numbers given in thee figure are measured in hundreds of dollars and that the discount rate is 5%.

that the household would have experienced had they not recovered from the shock at all and remained at their immediate post-shock income level. In Figure 1, the bottom of the lower shaded area traces out this hypothetical zero resilience income trajectory. Defining the total area between the zero resilience line and the counterfactual income trajectory $E\left(y_{it}^S | S_i = 1\right)$ (the red, dashed line in the figure) as R (the recovery area), we can define the normalized resilience metric for our study population as:

$$\mathcal{L}r^{10} = 1 - \left(\frac{\mathcal{L}^{10}}{\mathcal{L}^{10} + R}\right).$$

Note that this measure has the property that as cumulative losses approach zero, $\mathcal{L}r^{10}$ approaches 1, whereas if the population fails to recover at all, then $\mathcal{L}r^{10} = 0$. Except in the case of poverty traps (discussed below), we would expect $0 \leq \mathcal{L}r^{10} \leq 1$, with greater values of $\mathcal{L}r^{10}$ signaling a more resilient population that managed the shock with less compromise of current and future economic well-being. In this particular case, the estimated average resilience for our study population is $\mathcal{L}r^{10} = 43\%$, meaning that on average, households' partial recovery closes just over 40% of the total losses a household would have experienced if it had not recovered at all.

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 ($\hat{y}_t^c$) 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.

3.2 Identifying Resilience Capacities with Individual $\mathcal{L}r$ Measures

While the method detailed in the prior section recovers an estimate of average resilience for the study population, there are several ways to derive resilience metrics for either population sub-groups (e.g., women, members of savings groups, etc.), or even individuals. The former is easily doable by adopting the regression approach to account for heterogeneous treatment effects in the usual way. Here we consider the use of matching methods that allow the estimation of an individual-specific resilience metric.

Define $\hat{y}^C_{it}(y_{i-1}, \alpha_i)$ is the matched counterfactual estimate 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 initial well-being and entrepreneurial ability, but in real data, kernel and other methods of locating near neighbors for each treated observation could be used. We define individual resilience as:

$$(3) \quad \mathcal{L}r_i = 1 - \left(\frac{\sum_{t=0}^{9} \left(\hat{y}^C_{it}(y_{i0}, \alpha_i) - y^S_{it}(y_{i0}, \alpha_i) \right)}{\sum_{t=0}^{9} \left(\hat{y}^C_{it}(y_{i0}, \alpha_i) - y^T_{i0}(y_{i0}, \alpha_i) \right)} \right)$$

where the numerator in the fraction is simply the cumulative losses for shocked household i relative to their matched counterfactual. The denominator is simply the cumulative loss between the counterfactual and the household i's income in the immediate post-shock period projected forward in time (the zero recovery normalization).

Figure 2 displays a histogram of the individual resilience metrics using our artificial data. While the average resilience is 43%, the individual measures range from 25% to 70%. Consistent with the underlying data generation process, which does not admit multiple equilibrium poverty traps, the resilience measures for all households are positive as even the least resilient household has recovered from 25% of original income loss.

3.3 Using the $\mathcal{L}r$ Measure to Evaluate Policy

The prior two sub-sections have shown how the $\mathcal{L}r$ measure can be used to characterize the resilience of a population and of the individuals who comprise that population. Much of the interest in resilience measurement has stemmed from the introduction of policies designed to make households (and other units of analysis) better positioned to withstand shocks without the sort of long-lasting negative impact on household

Figure 2: Distribution of Individual Resilience under DGP-1

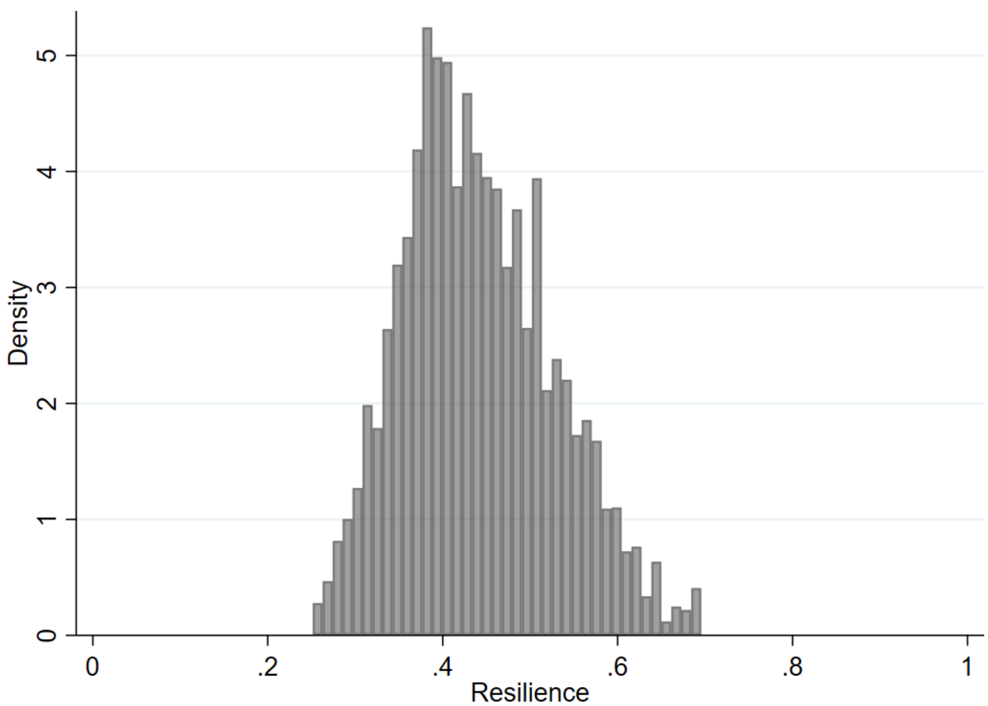

economic well-being that is visible in Figure 1. In this section, we "experimentally" introduce a catastrophic insurance policy that rebuilds assets for households following a severe shock. That is, we introduce the additional sub-populations defined in the second row of Table 1. We first look at the impact of the policy on average resilience. In order to allow for more reliable policy evaluation, we then use the discounted cumulative loss measure ($\mathcal{L}'$, defined in footnote 5 above) to create benefit cost measure of the effectiveness of the catastrophic insurance policy.

In this analysis, we assume that the government provides every household a catastrophic insurance policy that has the following characteristics:

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

Using the probabilities in our underlying model, we can calculate the actuarially fair price of this insurance policy. We further assume that the policy is sold to the government at a 25% mark-up over the actuarially fair pice. Importantly we ignore the behavioral consequences of insurance discussed by Janzen et al. (2021), which as they show can add substantially to the resilience-promoting impacts of this kind of insurance through what they call a behavioral, investment incentive effect.

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

$$(4) \quad y_{it} = \sum_{t=1}^{T} \left(\beta_t^c d_t + \delta_t^S \left(S_i \times d_t\right)\right) + I_{it} \left[\beta_I + \sum_{t=1}^{T} \left(\delta_t^I d_t + \delta_t^{IS} \left(S_i \times d_t\right)\right)\right] + \varepsilon_{it},$$

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

To allow for a fair evaluation of the benefits and costs of this policy, we assume that the government has been buying the contract for the entire population for a decade. Given that the severe loss events happen about 5% of the time, 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.

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 restock household assets and assist the recovery of income. The shaded area marked G measures the resilience gain from the policy, that 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 "No Poverty Trap" columns of Table 2, the resilience for the sub-population covered with the catastrophic insurance policy rises from 43% to 70%.

In order to evaluate the economic efficacy of this catastrophic insurance policy, we calculate the present value of the resilience gain (the area G in Figure 3) and

Figure 3: Measuring the Impact of Catastrophic Insurance on Resilience

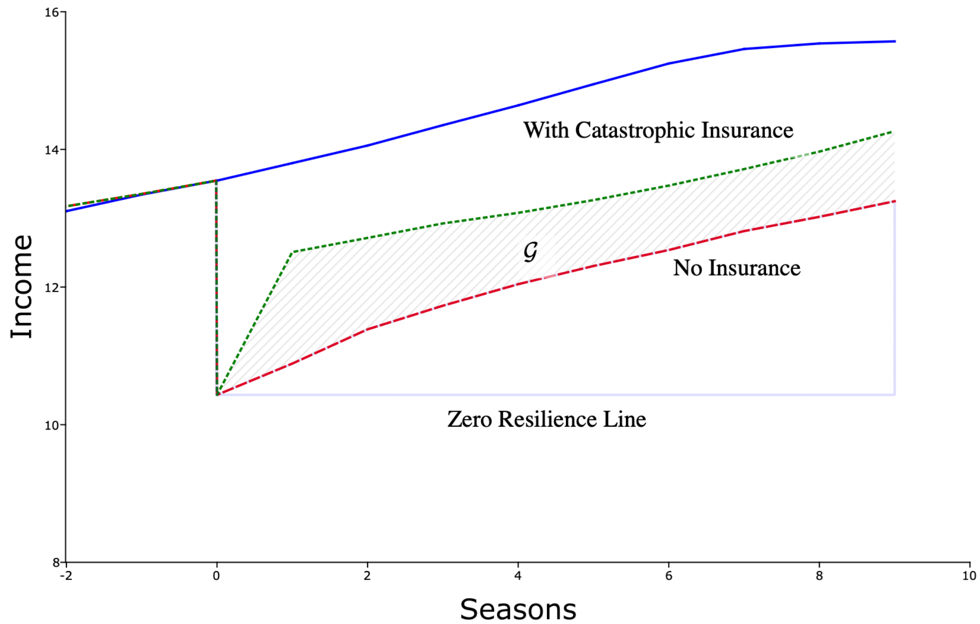

Table 2: Does Catastrophic Insurance Pay?

	No Poverty Trap		**Poverty Trap**	
	No Insurance	*Catastrophic Insurance*	*No Insurance*	*Catastrophic Insurance*
Mean Resilience	43%	70%	41%	72%
Benefit Cost Ratio	1.8		2.2	

then compare it to the present value of the full public expenditure on insurance for the entire population (shocked or not). As shown in Table 2, this benefit cost ratio is 1.8, meaning that every public dollar spent on insurance reduces cumulative losses by 1.8 dollars. Another way to think about this exercise is to imagine that the government has a reactive 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 catastrophic insurance policy would save $1.8 in post-shock social protection expenditures.

4 Poverty Traps and the Choice of Economic Outcome to Measure Resilience

The data generating process underlying the analysis in Section 3 set the fixed time cost associated with becoming an entrepreneur to zero. Under this specification, the income set becomes concave and all households stochastically approach a higher income entrepreneurial equilibrium. In this section, we modify the parameters of the model generating the data such that the fixed time cost parameter of being an entrepreneur (A) is strictly positive. As discussed in the appendix, this modest change in specification exposes a subset of middle ability individuals (labelled B-types in Figure A1 in the appendix) to multiple equilibrium poverty, meaning that if these individuals are born to poor, or 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 (A-types in the appendix figure) 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 can measure resilience as a concept distinct from poverty dynamics.[7] In general, because

[7]As discussed in Section 2, some discussions of economic resilience that build on Barrett and Constas (2014) appear to conflate 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

the Type A subset of the population will operate with a much lower capital stock under this modified data generation process (ie., 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 3. At the same time, Type B 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 (Type C in the appendix) are largely unaffected by the change in specification between the no poverty trap and poverty trap cases.

It is thus unclear whether average resilience will be higher or lower in the presence of poverty traps. However, we unambiguously expect the variability of the individual resilience measures to increase. The remainder of this section explores these issues and revisits the impact of the same catastrophic insurance policy considered in section 3.3.

4.1 Resilience Measurement in the Presence of Poverty traps

In the interest of space, we do repeat Figures 1 and 3 for the poverty trap data generation process. Table 2 does report average resilience and the impact of insurance. As can be seen, the resilience of the overall population does decline modestly under the poverty trap scenario. As discussed above, this change in the population average is a mix of households that are more resilient (but poorer) and others who become less resilient because they fall into a poverty trap.

As also shown in Table 2, the catastrophic insurance policy discussed in section 3.3 continues to have a major impact on population resilience (increasing from 41 to 72%). More interestingly, the catastrophic insurance is much more effective from a public finance perspective as its benefit cost ratio rises from 1.8 in the no poverty

may improve resilience but not resolve poverty (which may require asset transfers). To say that such a policy does not increase resilience because it does not eliminate poverty would seem to confuse the conversation.

Figure 4: Individual Resilience in the Presence of Poverty Traps

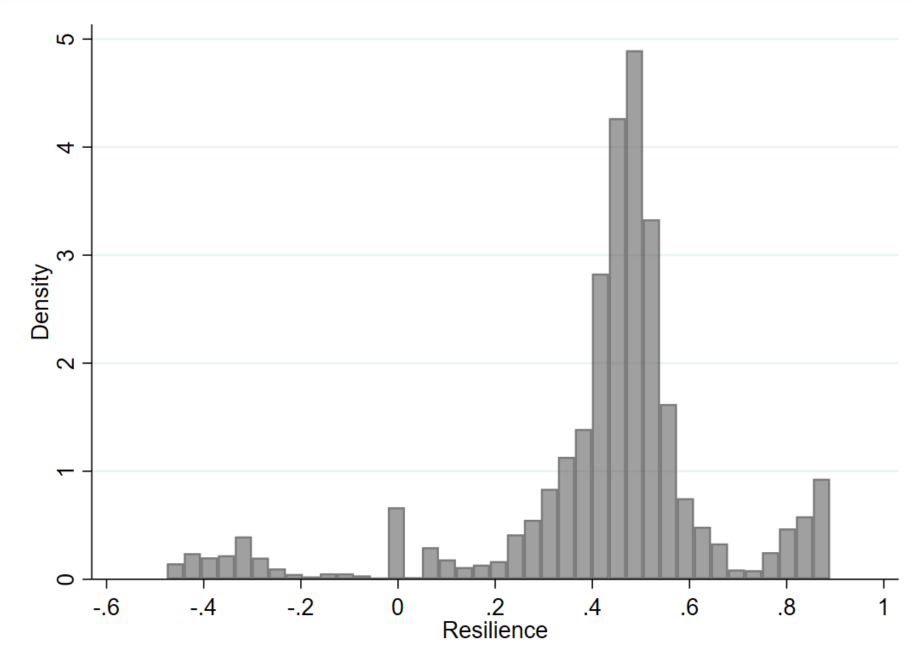

trap case to a hefty 2.2 in the presence of poverty trap. This increase is a clear signal that the policy has a major impact on Type B individuals who absent the insurance policy fall into a low income stochastic steady state. For at least a sub-set of these individuals, the catastrophic insurance transfer restores them to a position of economic viability from which they can move back to level of well-being of their non-shocked comparison group.

Analyzing individual resilience reveals important additional clues about the efficacy of catastrophic insurance in the presence of poverty traps. Using equation 3, we can again measure the distribution of individual resilience in the presence of poverty traps. Figure 4 shows the resulting distribution of individual resilience.

Comparing Figures 2 and 4, we see the clear presence of poverty traps in the latter figure. A not inconsequential number of households exhibit negative resilience as they not only fail to recover to their matched counterfactual position, but are also approaching a lower level equilibrium. At the upper end of the distribution, we also see some households with resilience measures in excess of 80% under the poverty trap data generation process.. These new highly resilient types are low skill Type A individuals whose counterfactual comparison group relies little on capital and hence they 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.

4.2 Why Consumption versus Income May Give Different Resilience Measures in the Short Term

In the presence of poverty traps, Type B individuals are liable to fall into a position of economic non-viability following a severe shock. As just shown, the resilience of these individuals becomes negative when gauging their economic well-being. But what if we were instead to measure resilience using consumption?

Because is no longer economically optimal for Type B individuals to aspire to accumulate and return to the time path to a high income equilibrium, they may find themselves with asset holdings that are above their steady state level for the lower income, non-entrepreneurial equilibrium they are now approaching. As such, it is optimal for these individuals to slowly deaccumulate their excess assets, allowing them to temporarily boost their consumption as they descend toward their new long run equilibrium position. This logic suggests that in the short run these Type B 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 indeed the case in our simulated data in the poverty traps case. Panel 5a on the left side of the figure 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, panel 5b shows that if we were to examine the consumption of these same households, they would appear more resilient for at least 4 seasons than their counterparts who did not fall into a poverty trap and are actually moving toward a higher long-term economic equilibrium. However, past that point, the true long-term lack fo resilience of the type B, trapped households becomes apparent.

Figure 5: Measuring Resilience with Consumption versus Income in the Presence of Poverty Traps

(a) Income (b) Consumption

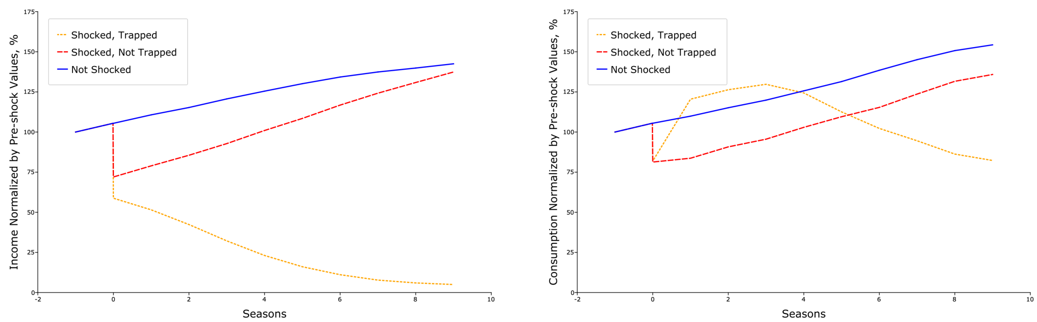

While our controlled data generation process allows us to clearly see what is going on, it raises a cautionary note about the economic outcome used to measure resilience and the time period over which we can measure that outcome. For shorter duration post-shock time series, this analysis suggests that income is the more reliable measure of resilience than consumption.

5 Measuring Resilience and the Impact of Resilience-promoting Policy with Short Duration Real World Data

This paper has so far used simulated data from known data generation processes to develop cumulative loss-based measures of resilience that can be used to diagnose the resilience of a population and analyze the benefit-cost ration of a simulated policy intended to bolster household resilience. In this section we use messier, shorter duration real world data to show how these ideas can in fact be implemented in practice and diagnose the level of resilience and gauge the impact of a particular intervention.

5.1 $\mathcal{L}$−resilience as a Diagnostic tool

To illustrate thee use of the $\mathcal{L}$−resilience metric as a diagnostic tool with real data, we utilize the control group data from an RCT studying the impact of a bundle of drought-tolerant maize and insurance on farmers using a sample of 3000 households Mozambique and Tanzania (see Boucher et al. (2022) for details). Over the course of the study, which was spatially diversified within and between countries, drought and other shocks hit a random selection of farmers in the sample. Using a variant of equation 2 above, Boucher et al. (2022) recover estimates of the contemporaneous and lingering future effects of severe shocks on control group farmers. Figure 6 illustrates the effect of the shock on control group, where the counterfactual is simply the income level for control farmers who did not suffer the shock. Figure 6 illustrates the estimated time paths for these two groups, where again the red line shows the average income of those who were shocked while the blue line shows the expected income of counterfactual group of households who did not experience the drought. The shock was large, resulting in a nearly 50% reduction in income and recovery was slow over the two-year time period we observe.[8] As analyzed in detail by Boucher et al. (2022), this slow recovery results from coping strategies that badly decapitalized farms in the wake of the shock. For this population, the $\mathcal{L}$−resilience measure $\mathcal{L}r^3 = 0.18$. The total, 3-year cumulative loss for these farmers averages $\mathcal{L}^3 = \$1443$, with the present value of the loss equal to $\mathcal{L}'^3 = \$1326$.

While further observations would be informative, even this short panel reveals that the resilience of the studied small-scale farming communities is quite weak. While we focus here only on losses of maize income, ancillary analysis of food insecurity by Boucher et al. (2022) shows a large and significant 25% increase in food insecurity in the year following the shock. In short, the diagnosis for these farmers is clear–their resilience is extremely low. We turn now to show how the same resilience measurement tool can be used to evaluate the intervention offered to treatment group farmers in

[8]The 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. As we discuss below, frequency with which longitudinal data are and can be collected post-shock is an important part of the proposed research.

Figure 6: $\mathcal{L}-$Resilience Amongst Maize Farming Households in Tanzania and Mozambique

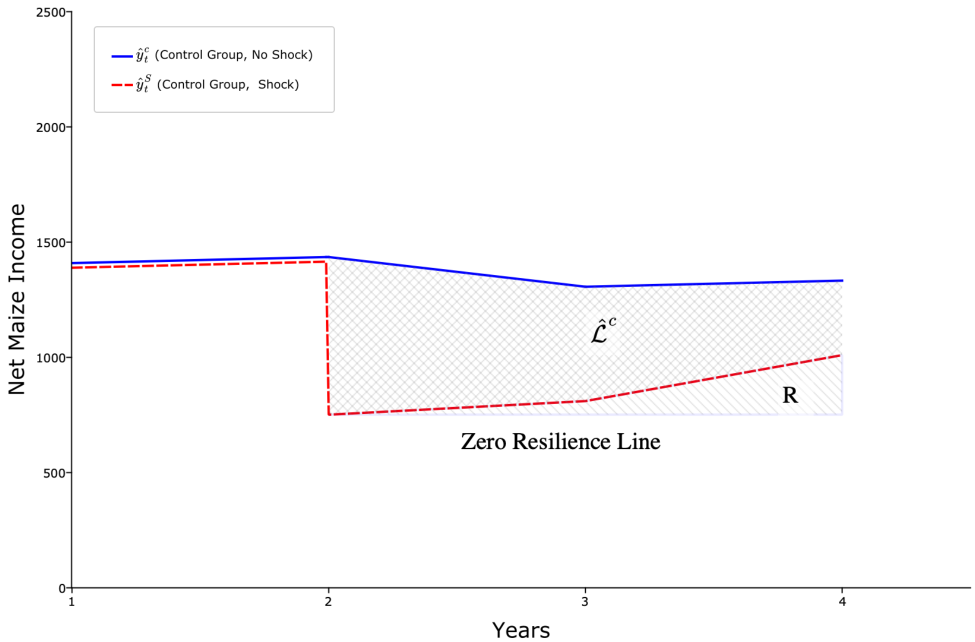

this same study.

5.2 Policy Analysis using $\mathcal{L}-$resilience

The Boucher et al. (2022) study that examined the impact of a bundled treatment of drought tolerant maize and fail-safe index insurance on the (decidedly non-resilient) population of maize farmers in Mozambique and Tanzania. While their experimental structure requires a slightly more complicated specification than that shown in equation 4, their analogue estimate allows recovery of predicted income time path for treatment group farmers who suffered the same shock that led to large losses of current and future income in the control group.

Figure 7 augments Figure 6 by adding the income path of those shocked but were exposed to the (putatively) resilience building bundle of drought tolerant maize and fail-safe index insurance. As can be seen in the figure, the cumulative loss area shrinks dramatically (from $\hat{\mathcal{L}}^c$ to $\hat{\mathcal{L}}^I$). This reduction in loss is the result of not having fallen as far because of the shock (due to the dramatic impact of the DT

Figure 7: Evaluating the Impact of a Resilience Intervention for Maize Farmers

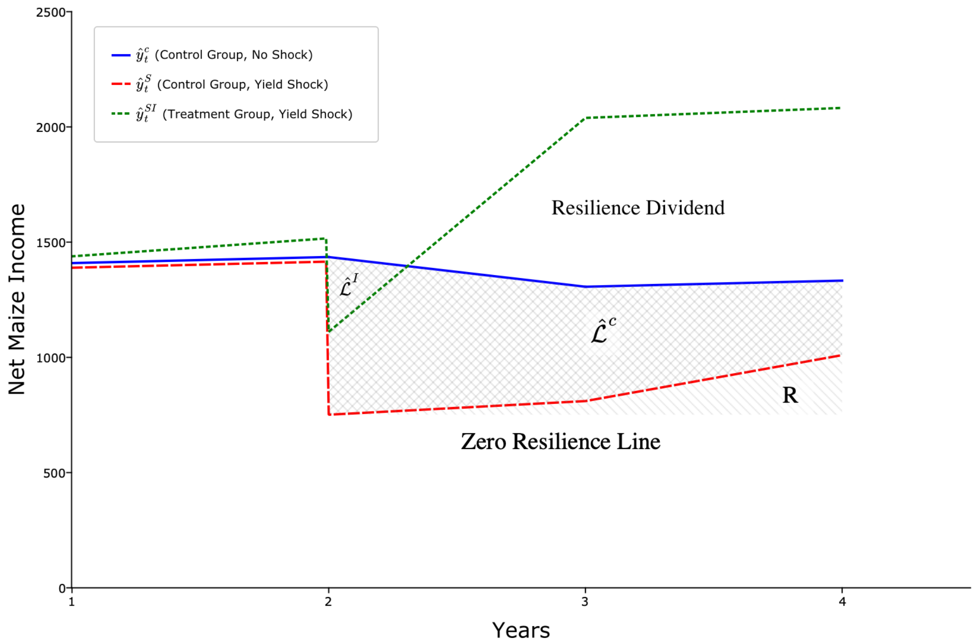

genetic technology) and because recovery the year following the shock is immediate. Indeed, as discussed in detail by Boucher et al. (2022), treated farmers more than recover form the shock as the experience of the resilience building technologies leads to a subsequent intensification of the adoption of these technologies at both extensive and intensive margins. This "seeing is believing" behavioral response generates what the figure labels as the resilience dividend.

Ignoring the resilience dividend (for discursive purposes), the $\mathcal{L}-$resilience for these farmers jumps 450% ($\mathcal{L}r$ rises from 18% to 82%) when they are treated with the genetic/financial bundle. More impressively, the benefit cost ratio for this resilience promoting intervention is 6.8:1, meaning every dollar spent on resilience generates a benefit in the form of reduced losses by almost \$7.[9] In this particular case, the intervention not only promotes resilience, but it is a good investment from a public finance perspective.[10]

[9]Formally, the benefit is defined as $\hat{\mathcal{L}}'^c - \hat{\mathcal{L}}'^I$ and the cost assumes the purchase of the technology package for the 5 years preceding the shock and the 5 years following the shock.

[10]An alternative way to measure the public finance impact of this resilience promotion effort would be to follow Janzen et al. (2021) and calculate the total change in government social protection expenditures under a hypothetical disaster response regime in which the government commits to

Finally, the econometric estimates graphed in Figure 7 reveal an additional benefit of resilience building technologies. As discussed in detail in the Boucher et al. (2022) paper, once farmers experience the usefulness of the resilience building technologies, they subsequently deepen their investments, expanding the area cultivated and the intensity of production (more improved seed). This behavior generates what is displayed in Figure 7 as the resilience dividend as farmers who experienced the technology ultimately move beyond the level of their non-treated, non-shocked counterfactual comparison. Adding in these additional benefits would further boost the benefit cost ratio for the resilience-promoting policy to over *10*.

6 Conclusion

With the onward advance of climate change and the development of policies meant to combat, it has become increasingly important to have measures of resilience that can be used to gauge the impact of those policies and their cost effectiveness. Unfortunately, what has come to be known as the resilience measurement literature is conceptually unclear and not up to these important tasks (Upton et al. (2022)). This paper has attempted to reboot this discussion and derive a resilience metric that captures what is in fact meant by economic resilience. Drawing on the rich economics literature on how households are theoretically expected to respond to shocks, and how they actually do, we derive an estimable resilience metric that is based based on a comparison between a shocked household's actual shock and post-shock income (or consumption) trajectory with a counterfactual measure of what that trajectory would have been absent the shock. Using data derived from known data generation processes (dynamic stochastic optimization models), we show how these metrics can be estimated on average for a population and at the level of the individual. Building on a recent literature that shows how RCT data can be used to evaluate the impact of policies on shock sensitivity, we also how this metric allows a through evaluation of

using income supports to restore households to their counterfactual status. This approach could be easily adopted to a policy that simply closes the poverty gap for all shocked households.

resilience-promoting policies. We find that the benefit-cost ratio of a catastrophic insurance policy is much higher in a world in which at least a fraction of the population is subject to poverty traps.

In addition to these core findings, we also show that when data are available for only a short time following a shock, using income to measure resilience is likely to be more reliable than consumption, at least in world in which poverty traps are operative. Finally, we use data from a just completed study of program promoted drought tolerant seed varieties in combination with insurance to show that these same methods can be used, even when the post-shock time series is short. Importantly, the real world data also shows programs that make households resilient can induce increased investment that generates "resilience-plus" or a resilience dividend in which the newly resilient household ultimately supersede the economic position to which they would have returned had they been merely resilient.

References

Alfani, Federica, Andrew Dabalen, Peter Fisker, and Vasco Molini, "Can we measure Resilience? A proposed method and Evidence from Countries in the Sahel," *World Bank: Policy Research Working Paper*, 2015, *7170* (January).

Alinovi, Luca, Ergdin Mane, and Donato Romano, "Towards the Measurement of Household Resilience to Food Insecurity : An Application to Palestinian Households," in R Sibrian, ed., *Deriving Food Security Information From National Household Budget Surveys. Experiences, Achievement, Challenges*, Rome: FAO, 2008, pp. 137–152.

Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman, "Labor markets and poverty in village economies," *Working paper*, 2016, (July), Advance online publication.

Barrett, Christopher B. and Mark Constas, "Toward a Theory of Resilience for International Development Applications," *Proceedings of the National Academy of Sciences*, 2014, *111* (40).

Barrett, Christopher B and Michael R Carter, "The economics of poverty traps and persistent poverty: Empirical and policy implications," *Journal of Development Studies*, 2013, *49* (7), 976–990.

Barrett, Christopher B. and Michael R. Carter, "The Economics of Poverty Traps and Persistent Poverty: Empirical and Policy Implications," *Journal of Development Studies*, jul 2013, *49* (7), 976–990.

Barrett, Christopher B, Kate Ghezzi-Kopel, John Hoddinott, Nima Homami, Elizabeth Tennant, Joanna Upton, and Tong Wu, "A scoping review of the development resilience literature: Theory, methods and evidence," *World Development*, 2021, *146*, 105612.

Boucher, Stephen, Michael R. Carter, Jon Einar Flatnes, Travis J Lybbert, Jonathan Malacarne, Paswel Marenya, and Laura Paul, "Bundling

Genetic and Financial Technologies for More Resilient and Productive Small-scale Agriculture," Technical Report w29234, NBER 2022.

Carter, Michael R. and Travis J. Lybbert, "Consumption versus asset smoothing: testing the implications of poverty trap theory in Burkina Faso," *Journal of Development Economics*, nov 2012, *99* (2), 255–264.

Cissé, Jennifer Denno and Christopher B. Barrett, "Estimating development resilience: A conditional moments-based approach," *Journal of Development Economics*, 2018.

Deaton, Angus, "Saving and Liquidity Constraints Published by : The Econometric Society Linked references are available on JSTOR for this article : IS CONCERNED with the optimal intertemporal consumption behav- ior of consumers who are restricted in their ability to borr," *1Econometrica*, 1991, *59* (5), 1221–1248.

Fafchamps, Marcel, Christopher Udry, and Katherine Czukas, "Drought and Saving in West Africa : Are Livestock a Buffer Stock ? 1," 1996, (December 1994).

Ikegami, Munenobu, Michael R Carter, Christopher B Barrett, and Sarah Janzen, *Poverty Traps and the Social Protection Paradox*, University of Chicago Press, 2019.

Jalan, Jyotsna and Martin Ravallion, "Geographic poverty traps? A micro model of consumption growth in rural China," *Journal of Applied Econometrics*, jul 2002, *17* (4), 329–346.

Janzen, Sarah A, Michael R Carter, and Munenobu Ikegami, "Can insurance alter poverty dynamics and reduce the cost of social protection in developing countries?," *Journal of Risk and Insurance*, 2021, *88* (2), 293–324.

Kazianga, Harounan and Christopher Udry, "Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso," *Journal of Development Economics*, apr 2006, *79* (2), 413–446.

Knippenberg, Erwin and John Hoddinott, "Shocks, social protection, and resilience: Evidence from Ethiopia," *Gates Open Res*, 2019, *3* (702), 702.

_, Nathaniel Jensen, and Mark Constas, "Quantifying household resilience with high frequency data: Temporal dynamics and methodological options," *World Development*, 2019, *121*, 1–15.

Lee, Seungmin, Kibrom Abay, and Christopher Barrett, "Estimating Multidimensional Development Resilience," *Working Paper*, 2023.

Macours, Karen, Patrick Premand, and Renos Vakis, "Transfers, Diversification and Household Risk Strategies: Can productive safety nets help households manage climatic variability?," *The Economic Journal*, 2022, *132* (647), 2438–2470.

Paxson, Christina H, "Using weather variability to estimate the response of savings to transitory income in Thailand," *The American Economic Review*, 1992, pp. 15–33.

Perrings, Charles, "Resilience and sustainable development," *Environment and Development economics*, 2006, *11* (4), 417–427.

Pimm, Stuart L, "The complexity and stability of ecosystems," *Nature*, 1984, *307* (5949), 321–326.

Premand, Patrick and Quentin Stoeffler, "Cash transfers, climatic shocks and resilience in the Sahel," *Journal of Environmental Economics and Management*, 2022, *116*, 102744.

Smith, Lisa C and Timothy R Frankenberger, "Does resilience capacity reduce the negative impact of shocks on household food security? Evidence from the 2014 floods in Northern Bangladesh," *World Development*, 2018, *102*, 358–376.

_ and _ , "Recovering from severe drought in the drylands of Ethiopia: Impact of Comprehensive Resilience Programming," *World Development*, 2022, *156*, 105829.

Upton, Joanna, Susana Constenla-Villoslada, and Christopher B Barrett, "Caveat utilitor: A comparative assessment of resilience measurement approaches," *Journal of Development Economics*, 2022, *157*, 102873.

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., "Recovery without resilience? A novel way to measure nutritional resilience in Nepal, Bangladesh, and Uganda," *Global Food Security*, 2021, *31*, 100573.

Zheng, Geyi, Michael R. Carter, Nathaniel Jenson, and Laurel Krovetz, "Unpacking the Psychological Channels that Underlie Spillovers and Impact Heterogeneity in Multi-faceted Graduation Programs," March 2023.

Appendices

These appendices lays 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.

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

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 (2014). In this model, individuals can devote their resources to one of two different occupations:

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

We assume both "livelihood 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 can be thought of as time that must be withdrawn from the casual labor market in order become an entrepreneur.[11] Combining these two livelihood functions yields a 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

[11] Give an example based on Bandiera et al. (2016).

depreciation.

To study the dynamics of occupational choice and consumption dynamics, we assume that individuals solve the following inter-temporal maximization problem:

$$\max_{c_{jt}} \qquad E_\theta \sum_{t=0}^{\infty} \beta^t u(c_{jt})$$

subject to:

$$c_{jt} \leq k_{jt} + F(\alpha_j, k_{jt})$$
$$F(\alpha, k) = max\,[F^w, F^e]$$
$$k_{jt+1} = (k_{jt} + f(k_{jt}) - c_{jt})\,(\theta_{jt+1} - \delta)$$
$$k_{jt} \geq 0$$

where E_θ is the expectation taken over the distribution of the negative shocks and β is the time discount factor. $u(c_{it})$ 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.[12] Appendix A below gives numerical values for parameters and shock distribution that underlie Figure A1, including the assumption that $A > 0$.

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, an indicator variable was formed indicating whether or not the individual was pursuing the wage labor or the entrepreneurial livelihood in period 60.

[12]Note also that this model assumes that capital is used for production and is not a strictly buffer asset.

Figure A1: Long-term Occupational Choice as a Function of Initial Endowments

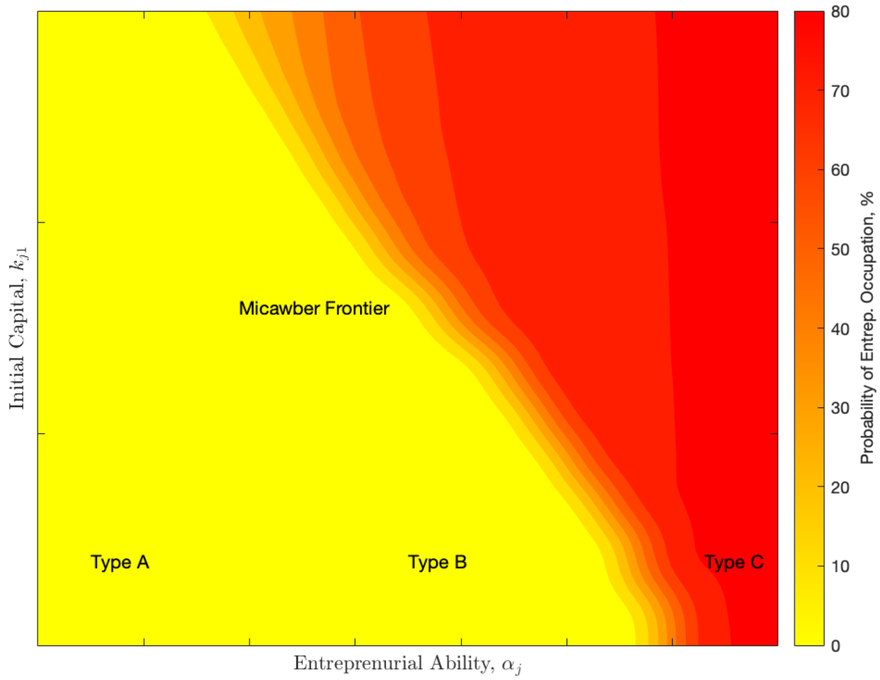

Source: Adapted from Zheng et al. (2023)

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 (2013b) 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 poverty trap. We denote this no poverty trap data generation process as DGP-1.

We also solve the model with the fixed cost of being an entrepreneur set to be strictly positive ($A > 0$). Under this parameter value, which we call DGP-2, type 2 individuals are subject to multiple equilibrium poverty traps.

For both data generation processes, 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.[13]

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
$F^w_{jt} = w_0 + k_{jt}^{\gamma_L}$ $F^e_{jt} = (w_0 - A) + \alpha_j k_{jt}^{\gamma_H}$ $\gamma_L = 0$ $\gamma_H = 0.56$ $A = 3.95$ (0, for the no poverty trap case) $w_0 = 3.95$
Utility Function and Parameters
Adaptive preferences utility function: $u(c_{it}) = \begin{cases} u^\ell(c_{it})\ if\ c_{it} < \tilde{c}(c_{g(i)}) \\ u^h(c_{it})\ otherwise \end{cases}$ 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: $\text{density of } \theta_{jt} = \begin{cases} 0.3 & \theta_{jt} = 0.11 \\ 0.18 & \theta_{jt} = 0.021 \\ 0.13 & \theta_{jt} = 0.031 \\ 0.11 & \theta_{jt} = 0.041 \\ 0.10 & \theta_{jt} = 0.051 \\ 0.02 & \theta_{jt} = 0.061 \\ 0.01 & \theta_{jt} = \{0.071, 0.081, ..., 0.191\} \end{cases}$

[13] Explain addition of classical measurement error