On the Definition and Estimation of Economic Resilience using Counterfactuals

Mo Alloush and Michael R. Carter*

September 2025

Abstract

A household is economically resilient if it manages a shock with minimal reduction in current and future well-being. This paper develops the first measure of resilience consistent with this simple, but compelling definition. Using simulated data, we implement the measure and show it can be used to judge the efficacy of resilience-promoting interventions. Simulated data allows us to explore the measure under a variety of scenarios, interventions and data generation processes. We also show how the measure performs in the with and without poverty traps and show that when data are limited, income is a more reliable measure of economic well-being for measuring resilience than is consumption. We also implement the measure using real data from a recent RCT on an intervention intended to promote resilience and show that the financial returns to promoting resilience are high and can save public resources compared to a reactive social protection policy.

JEL Codes: G51, H53, H84, O12

*Alloush: Colorado State University. 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 acknowledged. We thank seminar audiences at Cornell University, Universidad Politécnica de Madrid, the University of Göttingen and the World Bank for many helpful critiques and comments.

1 Introduction

Economic resilience can be defined as the ability of a household or other economic unit to manage a shock with minimal reduction in current and future well-being.[1] This paper develops the first measure of economic resilience that is consistent with this definition. Our measure is based on the cumulative losses that a shock-exposed household experiences relative to a counterfactual measure of what their 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 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 to evaluate whether or not the pursuit of resilience is wise public policy. It is as if policy wanted to reduce poverty without cogent measures of poverty or its depth. The resilience measure we develop here is meant to fill this lacuna. The measure also 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 loss that a household (or other unit[2]) experiences over time in the wake of a shock. For ease of our initial discussion, we use income as the indicator of economic well-being. Other measures of economic well-being, such as consumption or assets, can be used.[3] We define this loss $\mathcal{L}$ 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):

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

[1]The definition in the text is derived from the definition put forward by the US Agency for International Development (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" (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 which stirs together resilience with economic mobility.

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

[3]As we discuss later, these other indicators of economic loss may be preferable in studies with a short- to medium period of post-shock observation.

where β is the discount factor and we assume that the shock occurs in time period *0* and that we observe the household for T periods thereafter.[4] Normalizing $\mathcal{L}$ by the counterfactual stream of income the household would have had absent the shock provides a resilience measure that ranges between 0 and 1. Comparison of this resilience measure between otherwise comparable households, some of whom enjoyed access to a market or a program that enhances resilience, lays the foundation for a benefit-cost evaluation of the market or program in terms of its impact on cumulative avoided losses. The measure can also be used to track resilience over time to gauge if a population is becoming more or less resilient.

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 by households that manage their consumption and asset accumulation decisions to optimize their long-term, expected level of economic well-being. Simulated data allows us cleanly explore how the measure operates under different scenarios with different observational frequencies. Using the model, we randomly assign some of these households a severe economic shock. Other households do not suffer the shock and their income or consumption trajectories serve as counterfactuals for what would have happened to the shocked households had they not suffered a shock.

Using these generated data, we implement an econometric approach that allows us to cleanly 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 insurance program for productive assets by assigning a subset of shocked households access to the program and tracking their recovery trajectories.

In addition, we propose ways in which an individual-specific resilience measure can be estimated, opening the door to the empirical investigation of the socio-economic and geographic characteristics that are associated with greater resilience. We build on these key economic resilience concepts by extending 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

[4] In 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 $\mathcal{L}^T = \sum_{t=0}^{T} max\,[0, (\hat{y}_t^c - y_t^s)]$. Section 6 below discusses how to augment this measure to separately account for the building back better or resilience dividend phenomenon.

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 an asset insurance policy increases and the individual resilience metric becomes close to zero 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. Thus, households falling into a poverty trap may 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.

Finally, we use data from a randomized controlled trial of resilience-promoting technologies (stress tolerant seeds bundled with insurance) among smallholder farmers in Mozambique and Tanzania to show that our proposed measure of resilience can be implemented with real data. Interestingly, the technologies in the real data have a more pronounced positive impact on resilience than the simulated asset insurance program, exhibiting a benefit-cost ratio of 11:1 as opposed to 2:1 in the simulated data. We also use the estimates to show how they can be used to value the intervention using an approach that only considers humanitarian assistance as a benefit. Under this narrower definition of benefit, the program still exhibits a benefit-cost ratio of 1.8:1.

The remainder of this paper is structured as follows. In Section 2 we graphically present the intuition behind our proposed measure of economic resilience. In Section 3, we review 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 review the resilience literature that primarily focuses on creating indices using observable household and community that have been posited to promote resilience. In Section 4, we use simulation data from a known dynamic stochastic optimization model to estimate the resilience metric. This section also considers the impact of a simulated publicly-provided asset insurance policy on resilience, and also shows how the resilience metric can be modified to provide a benefit-cost

Figure 1: Cumulative Loss-based Resilience Measurement

(a) Elements for Resilience Measure (b) Enhanced Resilience Two-ways

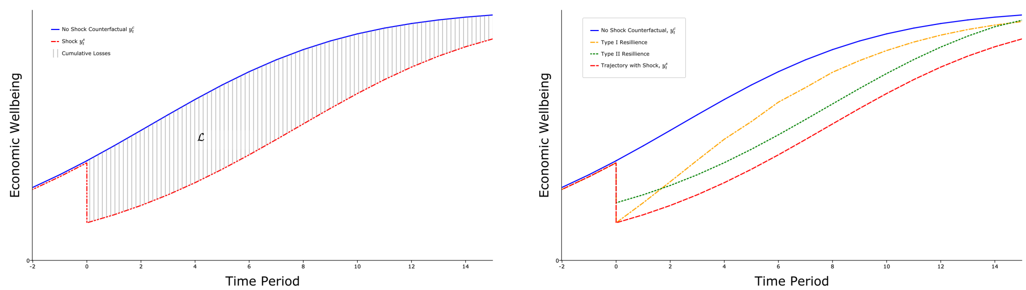

measure of this policy intervention. In Section 5, we show how the same methods can be used to estimate individual-level resilience measures. We also show how the measure can give insights about multiple equilibrium poverty traps. In Section 6, we illustrate our approach using real world data from eastern and southern Africa. 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 dashed 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 solid 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:

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

where $\mathcal{L}^T$ is the cumulative loss measure defined in expression 1 above and β is the discount factor. The term in square brackets captures the share of loss in current and future well-being induced by the shock. When those losses are zero, the resilience measure $\mathcal{R}^T = 1$, indicating full resilience. Lower values of $\mathcal{R}^T$ would signal larger cumulative avoided 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 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 measure 2 as the *CAL* Resilience Measure.

The *CAL* Resilience measure, $\mathcal{R}^T$, 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 of the economic or welfare measure of interest. For example, Alfani et al. (2015) define a quantitative measure of resilience based 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. While 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 phenomenon. 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, our we define resilience with respect to a counterfactual well-being path over time. Only in the special case in which the counterfactual trajectory is flat will these before-after comparison accurately measure the cumulative losses experienced by the shocked household.[5]

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 households (or groups of households), i and j, that suffer the same shock.

[5]The 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) who 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.

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:

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

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

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. Insurance is an example of a market or intervention that might operate through this mechanism and we label this type of resilience as Type A resilience.

A household can also become more resilient if 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 B resilience. These two types of resilience promotion are not mutually exclusive.[6]

The challenge is ultimately to estimate measure $\mathcal{R}$ 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 $\mathcal{R}$.

[6]Empirically, 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 et al. (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.

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). The Resilience Indicators for Measurement and Analysis (RIMA) put out by the UN Food and Agricultural Organization was the first such measure (Alinovi et al. (2008)) and it has been further modified over time (see d'Errico et al. (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 et al. 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 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. In an effort to build on the resilience capacities literature, Cissé and Barrett (2018) direct the resilience literature towards a focus on economic outcomes. These authors define resilience as the ability of households to avoid poverty over time as they traverse whatever good and bad events come their way. Specifically, their measure is based on the probability of avoiding poverty and can be used to identify which households are resilient in the sense that they avoid poverty in the future. The approach of these authors still does not specifically consider the impact of shocks and instead focuses on the estimated distribution of the error term from a well-being regression.[7]

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 dy-

[7] Upton et al. (2022) criticizes 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 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.

namics. 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 focussing the measure of resilience only on those households rather than shifting the meaning of resilience.

In contrast to this existing resilience literature, our measure of economic resilience defined above is based on observing shocks and the economic damage they cause, immediately and over time. The natural starting point for our approach is the very rich literature, spawned by the permanent income hypothesis (PIH) that explores the impact of shocks on household consumption and savings decisions. Unfortunately, the burgeoning resilience literature has heretofore 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.[8] Instead, Friedman argued that the household would spread out the impact of the shock over time (i.e., smooth consumption) by dis-saving 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

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

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.[9] Again the cumulative losses induced by shocks would be small.

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

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 et al. (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 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 even Deaton's analysis would suggest.

A possible explanation for these dissonant empirical findings is that for agricultural and other self-employed households, assets are not simply accumulated to buffer shocks but 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 under 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

[9] Assuming shocks are not auto-corollated.

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 reanalyze the same data used by Fafchamps et al. (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.

Stepping back, the empirical analysis of the economic impact of shocks has uncovered 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.[10] 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. Individuals in the model are credit-constrained (a la Deaton, 1991) and are subject to shocks that destroy capital. 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. 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 achieve the higher-income

[10]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}

entrepreneurial occupation (see Appendix B). In Section 5, we consider the case in which this is no longer true and it may no longer be optimal for some households to pursue the higher income entrepreneurial occupation, depending on their endowments and history of shocks.

While obviously 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, 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 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.

After using these first two populations to define resilience measures, Section 4.2 introduces the two additional randomly selected sub-populations shown in the second row of Table 1. Both groups receive an externally funded asset 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 how the insurance improves resilience and then use the resilience measure to show the cost effectiveness of the insurance program.

4.1 The Cumulative Loss-based Measure of Population Resilience

We consider 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 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 spirt 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 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:

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

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 $\delta's$ will be indistinguishable from zero if the shocked and non-shocked households are well balanced. As is obvious from this simple structure, our sub-populations estimates needed for cumulative loss measure Equation 1 are:

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

Figure 2: Cumulative Income Loss from Shock

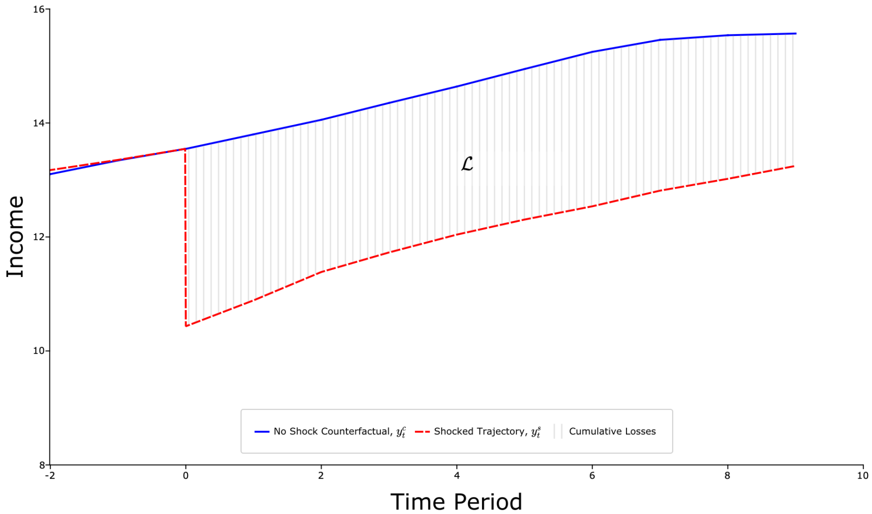

and

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

To illustrate the use of this machinery, we now turn to data simulated using the 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 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 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 plans), then the non-shocked households are in fact a good counterfactual for the shocked households. While 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.[11] 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 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 ($\hat{y}_t^c$) 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.

4.2 Using the *CAL*-Resilience Measure to Evaluate Asset Insurance

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 interest in evaluation 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 economic well-being that is visible in Figure 2. In this section, we "experimentally" introduce an asset 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 this analysis, we assume that the government provides every household an asset insurance policy against catastrophes that has the following characteristics:

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

- 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.

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 contemporaneous fall from the shock itself.[12]

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

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

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

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.

[12] 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. We consider such impacts in Section 6 when turn to analyze real world data in which such impacts can occur.

Figure 3: The Impact of Asset Insurance on Resilience

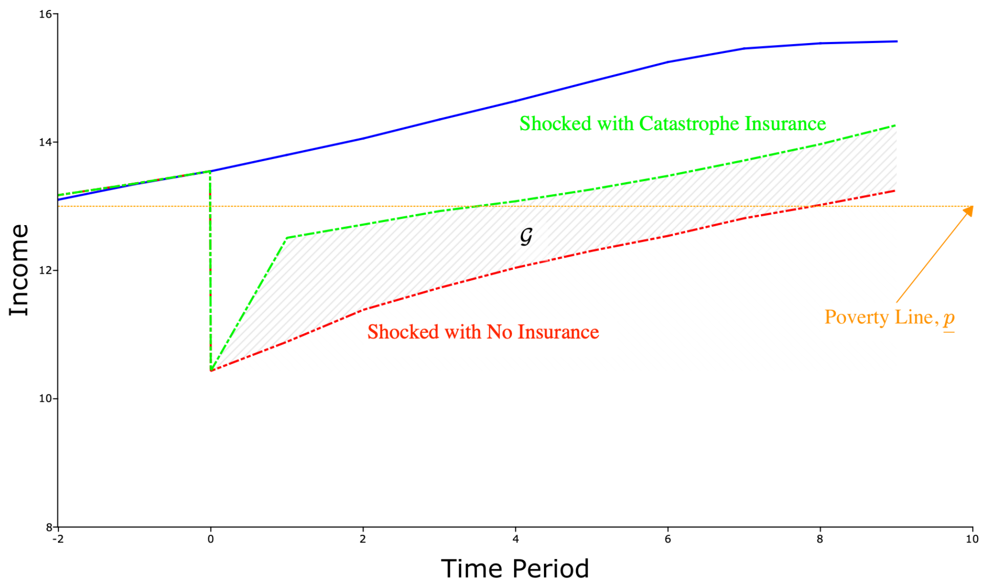

Table 2: Is Asset Insurance Good Policy Policy?

	Standard Data Generation Process (No Poverty Trap)		**Data Generation Process with Poverty Trap**	
	No Insurance	Asset Insurance	No Insurance	Asset Insurance
Mean $\mathcal{CAL}$ Resilience, $\mathcal{R}$	0.77	0.88	0.69	0.76
Benefit Cost Ratio	1.8		2.2	

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 asset 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 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.[13]

5 Poverty Traps and Individual Resilience Measurement

This section explores how the $\mathcal{CAL}$ 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. After exploring the implications of multiple equilibrium poverty traps on a population-level measure of resilience and the returns to the asset insurance policy introduced earlier, we develop an individual-specific resilience measure that provides additional insight into the economics of resilience and how to measure it.

5.1 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 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

[13] We define resilience without regard to the absolute income level of the shocked household unit. Barrett and Constas (2014) and Cissé and Barrett (2018) have argued that conceptually resilience is primarily of interest when focused on low-income or poor households. 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 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 are 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 later sections, we illustrate using the CAL resilience measure to measure benefits and costs when only humanitarian assistance avoided is valued as a benefit.

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 can measure resilience as a concept distinct from poverty dynamics.[14] In general, because the 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 shoes 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 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 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.

[14] As discussed in Section 3, some discussions of economic resilience that build on Barrett and Constas (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).

Figure 4: The Distribution of Individual Resilience

(a) No Poverty Traps (b) Multiple Equilibrium Poverty Traps

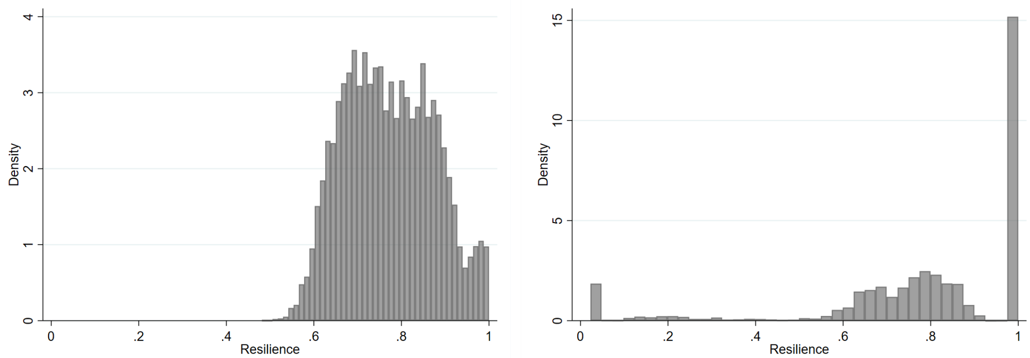

5.2 Individual Resilience Measurement

To measure resilience at the individual level, we need to identify a well-matched counterfactual for each individual. Define $\hat{y}_{it}^{C}(y_{i-1}, \alpha_i)$ is the matched counterfactual estimated 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 initial well-being and entrepreneurial ability. In real data, kernel and other methods of locating near neighbors for each treated observation could be used, or, with panel data, one could estimate income dynamics and use predicted income conditional on observable variables. We discuss these options in Appendix D.

We 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 their matched counterfactual.

Converting this to the proposed resilience measure again requires an appropriate normalization. Similar to our 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 77%, the individual measures range from 45% to 100%. Even the least resilient household has recovered from 45% of potential income loss.

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 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 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 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.

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 becomes negative when gauging their economic well-being. But what if we were instead to measure resilience using consumption?

Because it is no longer always economically optimal for middle ability individuals to aspire to accumulate and return to the time path to a high income equilibrium after a shock that puts them below their bifurcation threshold, they may find themselves with asset holdings that are above the steady state level for the lower income, non-entrepreneurial equilibrium they are now approaching. As such, it is optimal for these individuals to slowly de-accumulate 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 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. Graph 5a on the left shows that those who have fallen into a poverty trap

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

(a) Income (b) Consumption

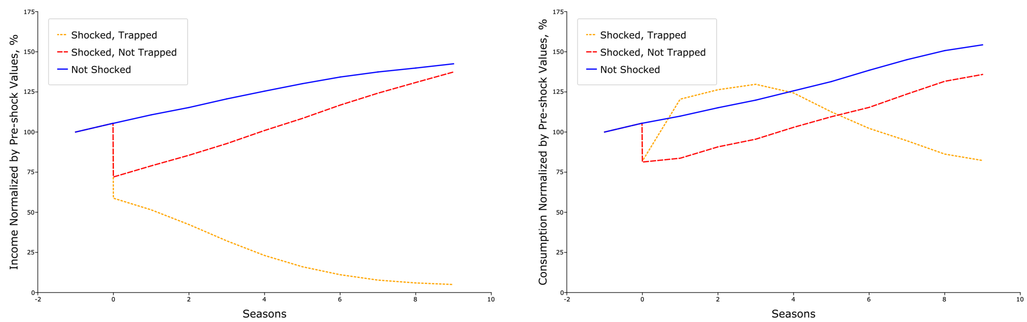

(whose trajectory is graphed as the dotted, orange curve) have an immediate income decline as they shift towards their new, lower income level. In contrast, Figure 5b shows that if we were to examine the consumption of these same households, they would appear more resilient for at least four seasons 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 of resilience of middle skills households, trapped households becomes apparent.

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 or assets will be the more reliable measure of resilience than consumption, at least in the presence of multiple equilibrium poverty traps.

6 Analysis of a Resilience-building Intervention in East and Southern Africa

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 ratio of a simulated policy intended to bolster household resilience. In this section, we introduce these tools to the real world. In particular, we use messier, shorter duration data from a recent randomized controlled trial (RCT) in Mozambique and Tanzania that covered the three year period stretching from 2015/16 through 2017/18.[15] The RCT offered some farm-

[15]Measurement of individual resilience, as discussed in Section 5.2 is more challenging using real

ers stress tolerant maize bundled with an index insurance product designed to cover losses not covered by the seeds that only protected farmers against droughts during the mid-season. Matching this controlled experiment, nature provided a natural experiment with drought and other shocks that affected farmers in both the control and treatment groups of the RCT, particularly during the 2016/17 agricultural year. The RCT in combination with the natural experiment assures that the 2×2 research design in Table 1 is fully populated. Full details on the technologies and the study are reported in Boucher et al. (2024).

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.[16] Figure 6 illustrates the effect of the shock on the RCT control group, where the counterfactual is simply the income level for control farmers who did not suffer the shock.[17]

Figure 6 graphs 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.[18] 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 $\mathcal{R}^3 = 64\%$, meaning that shocked farmers had a discounted stream of income over 3 years that is only 64% of the income enjoyed by non-shocked households. The total discounted cumulative losses of these households total $993.

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. (2024) shows a large and significant 25% increase in food insecurity

world data as it requires identification of a counterfactual for each observation rather than simply a valid counterfactual for a sub-population. Efforts to measure individual resilience using data from the Mozambique/Tanzania RCT resulted in quite noisy measures based on the difficulty of obtaining adequate counterfactual matches at the individual level. We show some of these in Appendix D.

[16]Boucher et al. (2024) estimate yields and area cultivated separately. Here we use their data and estimation model but simply estimate total maize income as a single variable.

[17]Because 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.

[18]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: *CAL*-resilience Amongst Maize Farming Households in Mozambique & Tanzania

Out/DT:II Analysis/con_only_august_25.png in the year following the shock. In short, the diagnosis for these farmers is clear—their resilience is extremely low. How then did the bundled genetic and financial intervention affect resilience and was it cost-effective?

A striking finding of the Boucher et al. (2024) is that farmers with access to the risk management technologies changed their behavior and began to invest more in their farms. This finding is consistent with prior studies of other risk management technologies, including index insurance and stress tolerant seed (see the review in Boucher et al., 2024 and also Lane, forthcoming). To account for this resilience dividend (as Boucher et al., 2024 label it), we expand expression 1 into a total economic impact measure by comparing the economic well-being of shocked households offered the resilience-building intervention ($\hat{y}_t^{sI}$) to the well-being of non-shocked households not offered the intervention ($\hat{y}_t^{cN}$):[19]

$$\mathcal{T}^T = \sum_{t=0}^{T} max\left[0, \left(\hat{y}_t^{cN} - \hat{y}_t^{sI}\right)\right] + \sum_{t=0}^{T} max\left[0, \left(\hat{y}_t^{sI} - \hat{y}_t^{cN}\right)\right] = \mathcal{L}^T + \mathcal{D}^T,$$

where the new term ($\mathcal{D}^T$) measures the resilience dividend for households offered the intervention.

Figure 7 augments Figure 6 by adding the income path of those shocked but were exposed to the 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 $\hat{\mathcal{L}}^s$ 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 impact of the DT genetic technology) and because recovery the year following the shock is immediate. Indeed, as discussed in detail by Boucher et al. (2024), 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 the resilience dividend, $\mathcal{D}$. 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 umps from 64% to 91%, meaning that treated households suffered only a modest loss of income compared to the counterfactual households who did not suffer a shock. This *CAL* measure is based on the intention to treat estimates from the Boucher et al. (2024) study. Because only 50%

[19] The measure of economic well-being from this study is total maize income for the household. Maize income accounts for approximately 65% of total family income. While a more complete measure would be preferable, maize income allows us to illustrate the resilience measure.

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

Out/DT:II Analysis/con_treat_august_25.png

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

	Experimental Status	
	Not Offered DT/II Bundle	Offered DT/II Bundle
Discounted Losses, $\mathcal{L}$*	\$993	\$226
Resilience Measure, $\mathcal{R}$*	64%	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.

of the farmers offered the DT/II bundle adopted it, these numbers understates the resilience gains experienced by those farmers who adopted the bundle. The time paths graphed in Figure 7 are the estimates for those offered the treatments, not those who adopted it.

Table 3 also shows the benefit cost ratio for this intervention.[20] 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.[21] 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 (space $\mathcal{D}$ in Figure 7), the benefit cost ratio rises to 11.2 to 1.

[20]Here are numbers used to calculateLosses only: =(993-226)/(291); HA only: =(561-42)/291; losses plus dividend: =((993-226)+2490)/291. Under HAA, are counting in extra cost of getting resilience plus, but do not value it.

[21]Note that this cost accounting does not account the fixed costs that underlay the development of the genetic and financial technologies. Because the products were sold commercially, the costs do account for product delivery and administration expenditures.

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 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 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 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.

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 intended to bolster resilience. Unfortunately, the existing literature on economic resilience is conceptually unclear and not up to these important tasks (Upton et al. (2022)). This paper reboots this literature and derives the first resilience metric that captures what is in fact meant by economic resilience understood as the cumulative losses caused by a shock. 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 income 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 resilience-promoting policies. We find that the benefit-cost ratio of an asset insurance policy against catastrophic shocks 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 a world in which poverty traps are operative. Finally, we use data from a just completed study of program promoted insured drought tolerant seeds to show that our resilience measure is informative 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. The benefit-cost ratio for this program is as high as 11-to-1, suggesting that resilience can be a very good public investment and substantially cheaper than responding to shocks with the cash transfers needed to restore people to their pre-shock positions.

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," *Quarterly Journal of Economics*, 2017, *132* (2), 811–870.

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.

_ , 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," *Economic Journal*, 2024, p. ueae012.

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," *Econometrica*, 1991, *59* (5), 1221–1248.

d'Errico, M., A. Garbero, and M Constas, "Quantitative Analyses for Resilience Measurement. Guidance for constructing variables and exploring relationships among variables," Technical Series 7, Food Security Information Network 2016.

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

Friedman, Milton, *A Theory of the Consumption Function*, Princeton University Press, 1957.

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, "Are the poor less well insured? Evidence on vulnerability to income risk in rural China," *Journal of Development Economics*, 1999, *58* (1), 61–81.

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.

Lane, Gregory, "dapting to Floods with Guaranteed Credit: Evidence from Bangladesh," *Econometrica*, forthcoming.

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.

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.

US Agency for International Development, "Resilience Evidence Forum Report," Technical Report, USAID April 2018.

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, "Psychosocial Constraints, Impact Heterogeneity and Spillovers in a Multifaceted Graduation Program in Kenya," *NBER Working Paper*, March 2023, (w31611).

Zimmerman, Frederick J. and Michael R. Carter, "Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints," *Journal of Development Economics*, aug 2003, *71* (2), 233–260.

Appendices (For Online Publication)

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 Dynamic 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 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.[22] Combining these two income functions yields a a non-concave set with locally increasing returns to scale: $F(\alpha, k) = max\left[F^w, F^e\right]$.

Following Ikegami et al. (2019), we assume that capital is subject to shocks and

[22]Give an example based on Bandiera et al. (2017).

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 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.[23] Appendix Appendix C below gives numerical values for parameters and shock distribution that underlie Figure 1, including the assumption that $A > 0$, meaning a discrete amount of time must be withdrawn from the labor market to operate the entrepreneurial technology.

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

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

shock sequences. Specifically, for each of 1500 initial positions evenly distributed across the initial endowment space shown in Figure 1. 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.

This entire process was then repeated 1000 times, generating 1000 histories for each of the 1500 initial endowment positions. The heat map in Figure 1 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 1 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

Figure 1: Long-term Occupational Choice as a Function of Initial Endowments

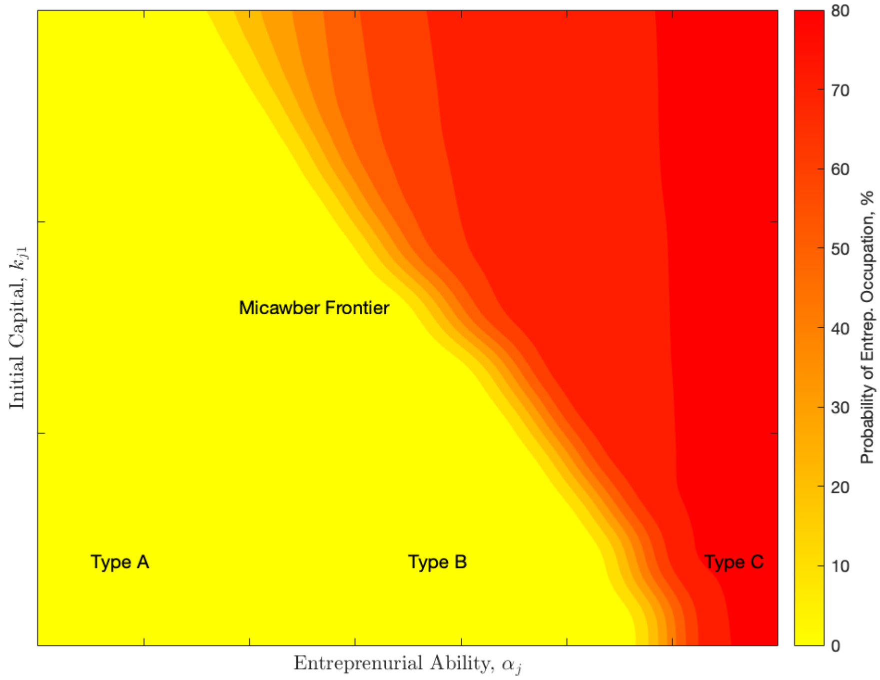

Source: Adapted from Zheng et al. (2023)

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 poverty trap. In this case, the equivalent of Figure 1 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 1, 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 1: 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: density of $\theta_{jt} = \begin{cases} 0.3 & \theta_{jt} = 0.11 \\ 0.18 & \theta_{jt} = 0.021 \\ 0.13 & \theta_{jt} = 0.031 \\ 0.11 & \theta_{jt} = 0.041 \\ 0.10 & \theta_{jt} = 0.051 \\ 0.02 & \theta_{jt} = 0.061 \\ 0.01 & \theta_{jt} = \{0.071, 0.081, ..., 0.191\} \end{cases}$

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 the treatment. In section 5.2, we also illustrate how our approach could be extended and applied to estimate individual level measures with simulated data. We note here that the counterfactual needed to estimate resilience at the individual level is not always readily available. In this Appendix, we propose easy-to-implement solutions that lead to individual measures of resilience that are highly correlated with our main counterfactual-based measure.

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. Below we make suggestions on how to estimate the no shock counterfactual that projects the income path an individual would have taken without a shock. We outline three different approaches to estimate this counterfactual:

1. **Matching Methods:** 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 approaches are likely to suffer from noisy estimates.

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 can be adjusted for inflation or even general growth levels in the economy. This is similar to approaches using by Alfani et al. (2015) and Knippenberg et al.

(2019) on return to pre-shock well-being levels.

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 less correlated with our main counterfactual-based measure at 0.8.