On the Definition and Estimation of Economic Resilience using Well-identified Counterfactuals

Mo Alloush and Michael R. Carter*

March 2022, Preliminary Draft

Abstract

This paper derives a quantitative metric of economic resilience 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. 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. These exercises reveal insights into the kinds of situations when promoting resilience is wise public policy. Finally, we are able to explore the sensitivity of our proposed resilience measures to the frequency with which data are collected post-shock. This exercise provides some guideposts for deployment of survey resources in the real world where data collection is expensive.

*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 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.[2]

The goal of this paper is derive a quantitative metric of economic resilience that measures what is meant by economic resilience as commonly defined above. Using data generated by a known dynamic stochastic model, we show how resilience can be measured and how it can be used to gauge the efficacy of policy interventions intended to promote resilience and to in turn measure the cost-effectiveness of such policies. We also show that our resilience measure is robust to multiple equilibrium poverty traps and brings out the additional gains from policies that promote resilience in this environment.

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)

[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, conflating economic mobility with resilience.

[2] In a theoretical exercise, Janzen et al. (2021) show that subsidy of an insurance, which keeps households in a position of economic viability, is less expensive than conventional, reactive social protection policies.

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.

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 Section 2 discusses in more detail, much of this 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. Specifically, the measure we put forward allows us to quantitatively measure the loss of current and future economic well-being relative to well-being levels without the shocks.

The remainder of this paper is structured as follows. Section 2 reviews the literature on the sensitivity of household well-being indicators to climate and other shocks. We also briefly review the largely disconnected resilience literature that primarily focusses on indexing characteristics that have been ex ante posited to promote resilience. Section 3 introduces a dynamic stochastic model of occupational choice that we use to generate the data on optimal consumption, income and assets that we use to develop and illustrate a resilience metric. Employing a known model to generate this data allows us to explore the robustness of the resilience metric to both data frequency and to the presence or absence of poverty traps. Section 4 then employs high frequency data from the 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 publicly-provided catastrophic insurance policy on resilience, and also shows how the resilience metric can be adapted to provide a benefit-cost measure of this policy intervention. Section 5 relies on a data generation process that admits multiple equilibrium poverty traps and illustrates the additional insights that emerge from the resilience metric and from the analysis of the catastrophic insurance policy. Finally, Section 6 shows the impact of reduced data frequency on the resilience metric, while Section 7 concludes.

2 Sensitivity to Shocks and the Measurement of Resilience

This section reviews the rich economic literatures on sensitivity of consumption to shocks, and the more modest literature on economic resilience per se. Unfortunately, these two literatures are 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.

2.1 Shock Sensitivity

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, 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 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 just a optimal, dynamic expected utility maximization 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

[3] Assuming shocks are not auto-corollated.

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" Measurement Literature

While the economics literature on shock sensitivity is theoretically well-grounded and empirically rich, the economics literature on resilience per se is less deep. 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 to 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 existing measures of development resilience do not tie measurement to a no-shock counterfactual. Our proposed measure differs significantly from these, we highlight these measures because they have been used extensively among development oriented policy institutions.

Currently, the most common measurement tool for resilience is one that uses latent variables. 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. This makes this approach particularly unsuitable to evaluate the impact of policies or programs specifically intended to improve resilience.

Another popular latent variable approach among development agencies abstracts away from 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. Resilience capacity is then measured by creating three indices using factor-analysis of relevant latent variables. 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.

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 and can be aggregated for general population or group-level measures. Cissé and Barrett (2018) argue that their proposed measure of resilience is both forward-looking and allows for nonlinear well-being dynamics. However, their measure 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. Our proposed measure differentiates resilience from measurement of poverty in the steady states.

Finally, a small number of studies conduct pre- and post-shock comparisons that are similar to our proposed measure. 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. 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. These two approaches most closely think of resilience in the way we define and measure it in this paper.

3 The Data Generation Process

In order to clearly develop our resilience measure, we create (noisy) artificial data generated by a known dynamic stochastic optimization model. While obviously avoiding the messiness of real world data, this approach affords three key advantages. First, we can create well defined experiments in the model, by exposing some households to shocks and, or to policy interventions intended to promote resilience.

Second, we can manipulate the model in order to study distinct data generation processes (DGPs) that are central to the study resilience. For DGP-1, we structure the model so that it has a single equilibrium, with all individuals trying to accumulate and converge toward a higher income equilibrium. For DGP-2, we extend the model so that it admits the case in which some individuals are subject to multiple equilibrium poverty traps (Barrett and Carter (2013) and Ikegami et al. (2019)). While distinguishing between these cases in real world data is challenging (see ... and Alloush and Carter, 2022), by controlling the data generation process in our artificial data we are able to explore the robustness of our measure to the presence of poverty traps.

Third, by using a model to generate data, we can manipulate the frequency with which we observe the consumption, asset and income data on our sample of household units. While we begin with high frequency data (biannual observations), in Section 6, we consider the performance of our resilience measures when data are obtained less frequently.

3.1 The Dynamic Stochastic 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.[4] Combining these two livelihood 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 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

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

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:

$$c_{jt} \leq k_{jt} + F(\alpha_j, k_{jt})$$

$$F(\alpha, k) = max\left[F^{w}, F^{e}\right]$$

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

$$k_{jt} \geq 0$$

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.[5] Appendix A below gives numerical values for parameters and shock distribution that underlie Figure 1, 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 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 intial 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

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

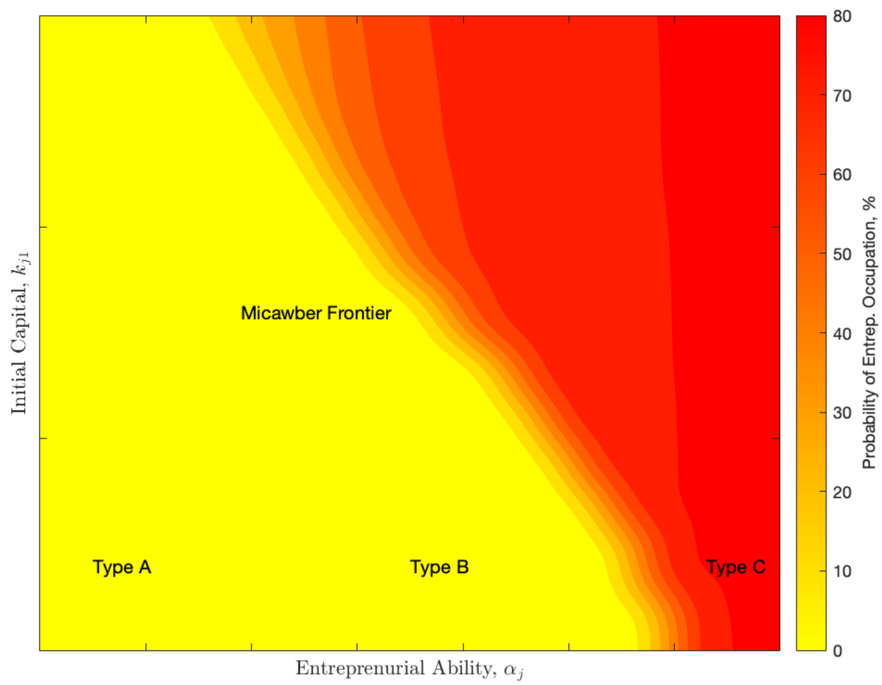

Source: Zheng et al. (2023)

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

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 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.

3.2 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 fo 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.

4 A Counterfactual-based Resilience Measure

This section uses the artificial data extracted from DGP-1 to develop a counterfactual-based resilience metric. We first use a standard treatment effects equation to show an average resilience measure based on the cumulative losses of households subjected to a shock relative to a valid counterfactual measure of what their well-being would have been without the shock, both at the time of the shock and in the future. A simple normalization of that cumulative loss measure results in an intuitively appealing and easy to interpret metric of average resilience. When the shock does zero damage to current and future well-being to the well-being of shocked households (compared to what their well-being would have counterfactually have been without the shock), the resilience metric is 1. When households become stuck and do not recover at all, then the metric drops to zero, indicating zero resilience. As we explore more deeply in Section 5, the metric becomes negative if the shock pushes a household into a poverty trap such that their economic well-being deteriorates further in the post-shock situation.

This section also shows how econometric matching methods can be used to derive the distribution of resilience across the sample population, pointing the way to an analysis of factors that promote resilience. Finally, we 'experimentally' introduce a catastrophic insurance policy to the households in our data set. Using this exercise, we are able to both explore the impact of the policy on resilience, and also calculate a benefit-cost measure of the policy's fiscal efficacy.

4.1 A Resilience Metric to Measure what We Mean

We first consider a simple regression model that can be applied to panel data that includes measures of household well-being (consumption, income and cash-on-hand, or wealth) 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 (as described above), and define $Post_t$ as a binary indicator for all time periods after the shock occurs. Letting w_{it} represent an economic well-being measure for household i in time period t, we write the regression model as:

$$w_{it} = \sum_{t=1}^{T} \left(\beta_t^c d_t + \beta_t^S \left(S_i \times d_t\right)\right) + Post_t \times \left(\sum_{t=1}^{T} \beta_t^S \left(S_i \times d_t\right)\right) + \varepsilon_{it},$$

where there are T time periods in the panel data set, d_t is a vector of T time period binary variables and β_t^C and $\beta_t^S$are vectors of coefficients for control and treated (shocked) households, respectively.

To illustrate ideas, and explore the anatomy of a shock (when households optimally manage their consumption and assets in the sense of the dynamic optimization problem above), we first estimate this regression equation for three measures of economic well-being: household consumption, household income and household cash-on-hand. The latter is the sum of the value of assets plus income and is the amount that constrains the consumption every time periods in the credit-constrained maximization problem in the prior section. Figure 2 displays the estimated outcome variables for the treated households, but only shows the estimated consumption level for control households. To make the graph more easily interpretable, all well-being measures are normalized by their immediate pre-shock value.

As can be seen, treated households received a substantial shock between seasons 3

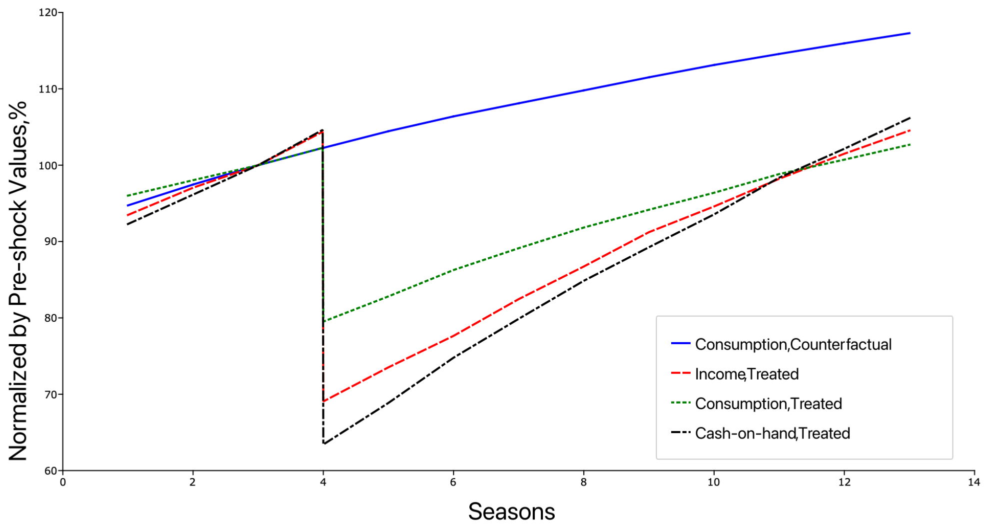

Figure 2: Anatomy of a Shock: Estimated Trajectories of Consumption, Income & Cash-on-hand

and 4. As can be seen, their cash-on-hand to fund immediate post-shock consumption and asset rebuilding plummets almost 40%. Income falls by nearly as much, while on average households defend consumption to some extent, by rebuilding their decimated asset stocks only slowly. The uppermost, solid blue line is the counterfactual estimate of what household's consumption trajectory would have been in the absence of a shock.

Figure 3 looks more closely at the predicted actual $\left(E\left(y_{it}^{C}|S_i = 0\right)\right)$ and counterfactual $\left(E\left(y_{it}^{S}|S_i = 1\right)\right)$ income trajectories.[6] The total loss in economic well-being is represented by the cross-hatched area "L". 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, we normalize the cumulative economic loss caused by the shock by the income trajectory 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 3, the bottom of the lower shaded area traces out this hypothetical zero resilience income trajectory. Defining the to-

[6] The same analysis could be undertaken with consumption and cash-on-hand.

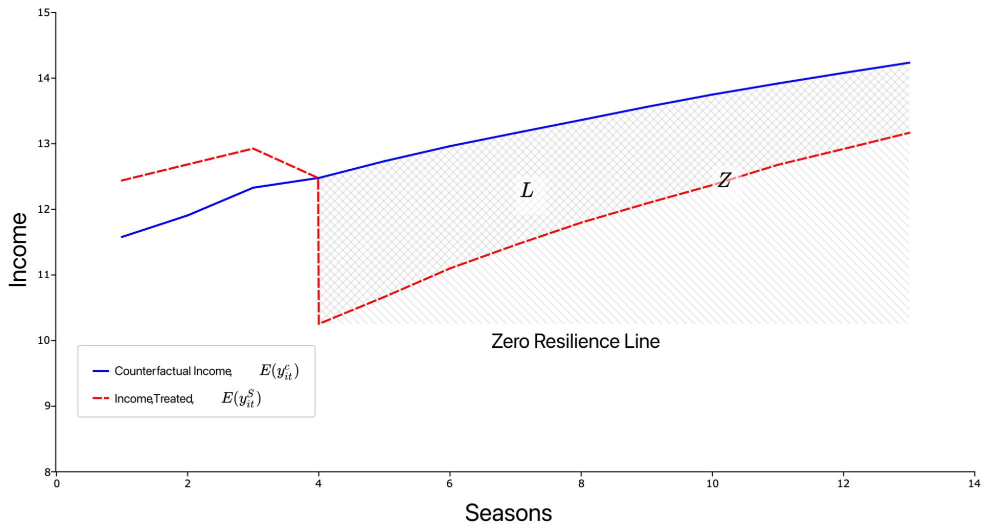

Figure 3: Cumulative Income Loss from Shock

tal area between the zero resilience line and the counterfactual income trajectory $E\left(y_{it}^{C}|S_i = 0\right)$ (the solid/blue line in the figure) as Z,we can define the normalized resilience metric for our study population as:

$$\bar{R} = 1 - \left(\frac{L}{Z}\right).$$

Note that this measure has the property that as cumulative losses approach zero, $\bar{R}$ approaches 1, whereas if the population fails to recover at all, then $\bar{R} = 0$. More generally, except in the case of poverty traps discussed below, we would expect $0 \leq \bar{R} \leq 1$, with greater values of $\bar{R}$ 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 43% (see Table 1), meaning that on average, households close just less than half of the losses over the 10, post-shock seasons covered by our data. Note that this measure does not discount future losses, and treats early and later losses the same way. We will return to the issue of discounting in section 4.3.

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 () 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.

4.2 Sub-group and Individual Resilience Metrics

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 $\hat{y}_{it}^C(y_{i0}, \alpha_i)$ is the matched counterfactual estimate for person i. In our simulated data, we are able to use exact matching based on initial wealth 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:

$$R_i = 1 - \left(\frac{\sum_{t=1}^{10} \left(\hat{y}_{it}^C(y_{i0}, \alpha_i) - y_{it}^S(y_{i0}, \alpha_i) \right)}{\sum_{t=1}^{10} \left(\hat{y}_{it}^C(y_{i0}, \alpha_i) - y_{i1}^T(y_{i0}, \alpha_i) \right)} \right) = 1 - \left(\frac{L_i}{Z_i} \right)$$

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.

Figure 4 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 aseven the least resilient household has recovered from 25% of original

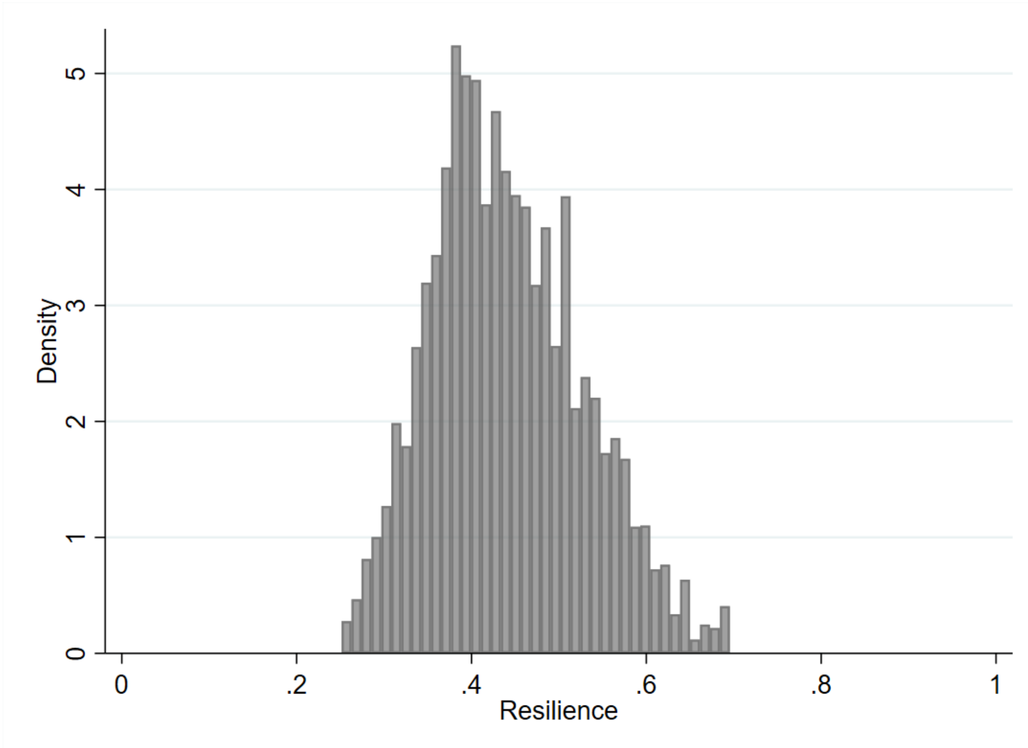

Figure 4: Distribution of Individual Resilience under DGP-1

income loss.

4.3 Using the Resilience Measure to Evaluate Policy

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 economic well-being that is visible in Figure 3. In this section, we "experimentally" introduce a catastrophic insurance policy that rebuilds assets for households following a severe shock.[7] We first look at the impact of the policy on average resilience. In order to allow for more reliable policy evaluation, we then introduce a version of the resilience metric which introduces discounting and then compare the discounted present value of the benefits of the policy compared to its costs.

In this analysis, we assume that the buys every household a catastrophic insurance

[7]Effectively, we allow the entire population to relive their identical stochastic history, except that those who receive a catastrophic shock receive a partial asset restocking grant the year after the shock.

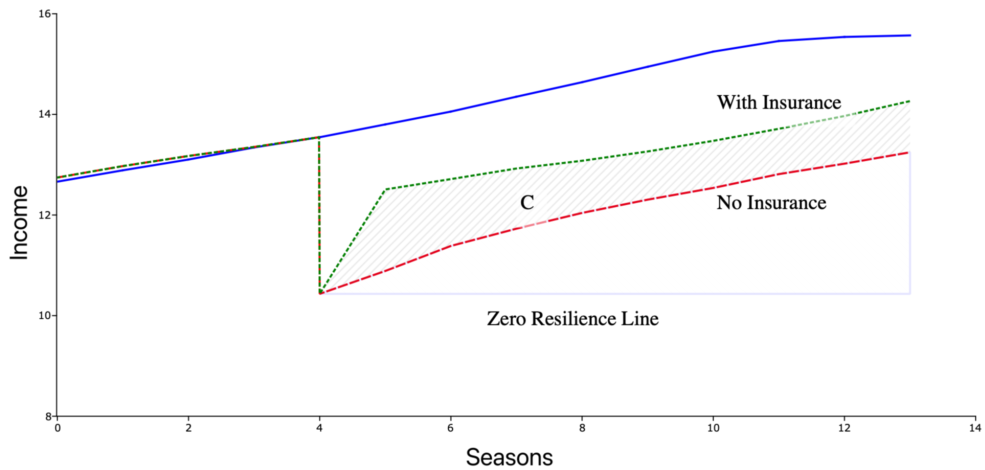

Figure 5: Measuring the Impact of Catastrophic Insurance on Resilience

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,
- The new 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 the show can add substantially to the resilience-promoting impacts of this kind of insurance through what they call a behavioral, investment incentive effect.

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.

	DGP1, No Poverty Trap		DGP2, Poverty Trap	
	Autarchy	*Catastrophic Insurance*	*Autarchy*	*Catastrophic Insurance*
Mean Resilience	43%	70%	41%	72%
Benefit Cost Ratio	1.8		2.2	

Table 1: Does Catastrophic Insurance Pay?

Figure 5 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 C measures the resilience gain from 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 1, mean resilience ($\bar{R}$) rises from 43% to 70% when the catastrophic insurance policy is implemented.

In order to evaluate the economic efficacy of this catastrophic insurance policy, we calculate the present value of the resilience gain (the area C in Figure 5) and 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 1, this benefit cost ratio is 1.8, meaning that every public dollar spent on insurance generates 1.8 dollars in additional income, compared to the no insurance case.

As shown in Table 1, the Catastrophic policy increases resilience from 0.45 for the uninsured to 0.XX for the insurance. The net present value of the resilience gain illustrated in Figure 5 is 71-44. Accounting for the cost of money, every dollar spent promoting resilience through insurance returns $1.8 in benefits under data generation process 1.

5 Resilience Measurement in the Presence of Multiple Equilibrium Poverty Traps

The data generating process underlying the analysis in Section 4set the fixed time cost associated with becoming an entrepreneur to zero. Under this specification, the income set becomes concave and all households stochastically approach the 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 Section 3, this modest change in specification exposes a subset of individuals with intermediate entrepreneurial skill (whom we labelled Type B) to a multiple equilibrium poverty trap. This change in specification also means that low skill, Type A 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.[8] In general, because the Type A subset of the population will operate with a much lower capital stock under DGP2 (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 DGP1. At the same time, Type 2 individuals no face the risk of falling into a poverty trap such that the impacts of the shock are long-lasting and irreversible. Finally note that individuals with high entreprenurial skill (Type C) are largely unaffected by the change in specification between DGP1 and DGP2.

It is thus unclear whether average resilience ($\bar{R}$) will be higher or lower under DGP2 compared to DGP1. 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 4.3.

[8]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 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.

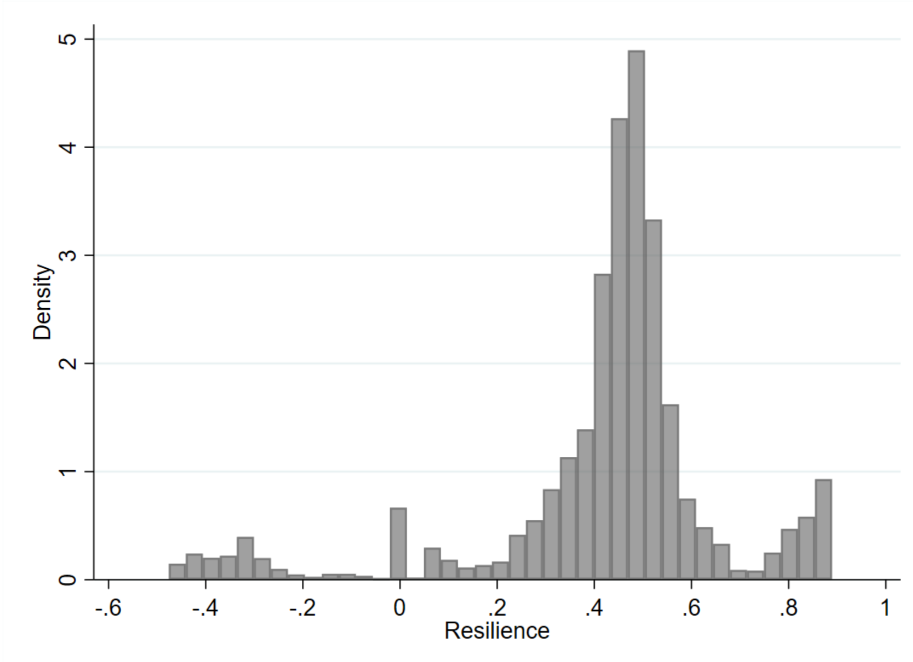

Figure 6: Individual Resilience in the Presence of Poverty Traps

5.1 Resilience Measurement in the Presence of Poverty traps

In the interest of space, we do repeat Figures 2 and 3 for the poverty trap data generation process. When we calculate average resilience $\bar{R}$ for this case we find that it is little changed from DGP1, falling from 43% to 41% (see Table 1).

More interesting is what happens the distribution of the individual resilience measures based on the matching methods presented earlier. Comparing Figures 4 and 6, 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 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 DGP2. While we still need to explore this further, this is almost surely Type A individuals whose counterfactual comparison group relies little on capital. The impact of shocks on these households dissipates more 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.

5.2 The impact of Catastrophic Insurance in the Presence of Poverty Traps

We also consider the impact of the catastrophic insurance policy. Using the same methodology described earlier, we see that in the presence of poverty traps, the catastrophic insurance is much more effective from a public finance perspective as the benefit cost ratio rises from 1.8 to a hefty 2.2. This 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 lease a sub-set of these individuals, the asset transfer pushes them back above the Micawber frontier illustrated in Figure 1 above.

6 Measuring Resilience with Less Frequent Data

While our ability to inexpensively generate high frequency using our model based data generation process is clearly artificial, we can choose to measure the impact on resilience measures of less frequent data. Figures 3 and 5 allow us to infer what will happen. Under optimal recovery and accumulation, we see the most rapid changes in income and consumption take place in the first few seasons following the shock. Failing to record economic well-being indicators during those crucial early periods would risk serious mis-measurement of resilience (or lack thereof). For example, looking at Figure 5, we can see that if we waited until season 8 (2 years post-shock), we would miss much of the welfare gain generated by the catastrophic insurance policy. Under DGP2, the degree of mis-measurement from delayed or infrequent survey would likely be less.

As we will explore further, there are two possible approaches to this problem. One would be to consider some kind of linear interpolation between say season 8 and season 4. Another, would be to consider the optimal deployment of survey resources in order to most closely approximate what could be learned from consistently high frequency surveys. For example a budget of 5 post-shock surveys might be optimally distributed in seasons 5, 6, 10 and 12.

6.1 Group-level Example with Infrequent Experimental Data

To illustrate how this method can be used with real data, we apply this method to group level data from a randomized controlled trial in Mozambique and Tanzania. In this RCT, rural farmers were offered a bundle of drought-tolerant maize and index insurance. This treatment is directly designed to help farmers recover faster after a drought—a large shock. A thorough investigation of consequences of this treatment can be found in Boucher et al. (2023). For our purposes, this RCT was spatially diversified within and across the two countries and a large drought affected some of the areas in the study.

Assuming that shock incidence is random after controlling for baseline characteristics, we can use those who were not exposed to the shock as a counterfactual for those who did. Figure 7 illustrates the two groups. The red line shows the average income of those who were shocked while the blue line shows the average income of counterfactual group of households who did not experience the drought. The shock was large—a nearly 50% reduction in income and recover was slow in the time period we observe. Using our method of calculating resilience, we get a number $\bar{R} = 0.18$.

Figure 8 adds the income path of those shocked but were exposed to the bundled treatment of drought tolerant maize and fail-safe index insurance. That group shows much higher resilience (0.82) and even shows a resilience dividend. On average, their initial loss from the shock was lower and their recovery quicker and the dividend showed that their treatment led them to invest more in the future as a reaction to experiencing the benefit of the treatment. This resilience-building treatment clearly had a large effect.

7 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

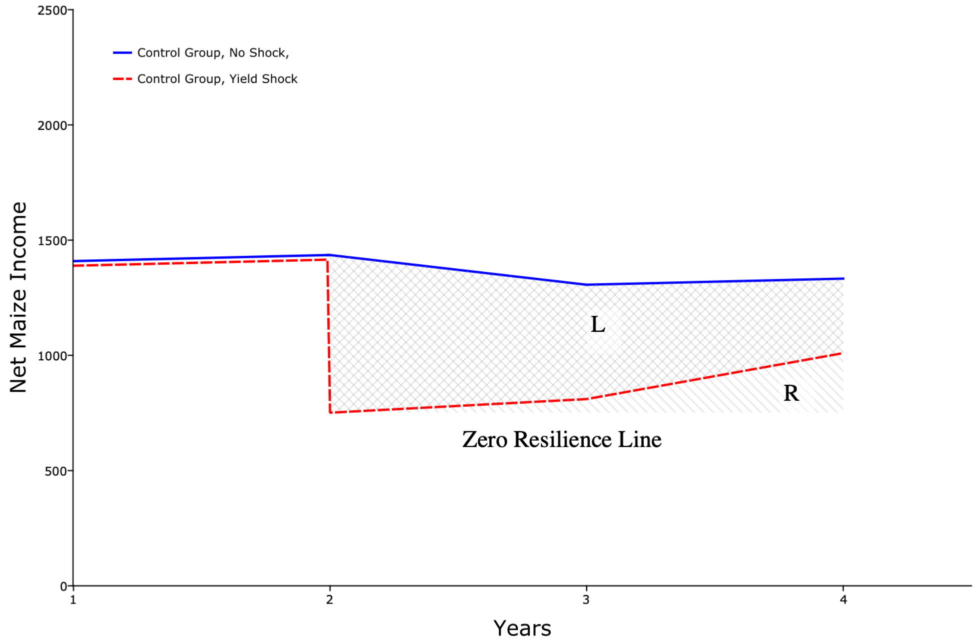

Figure 7: y

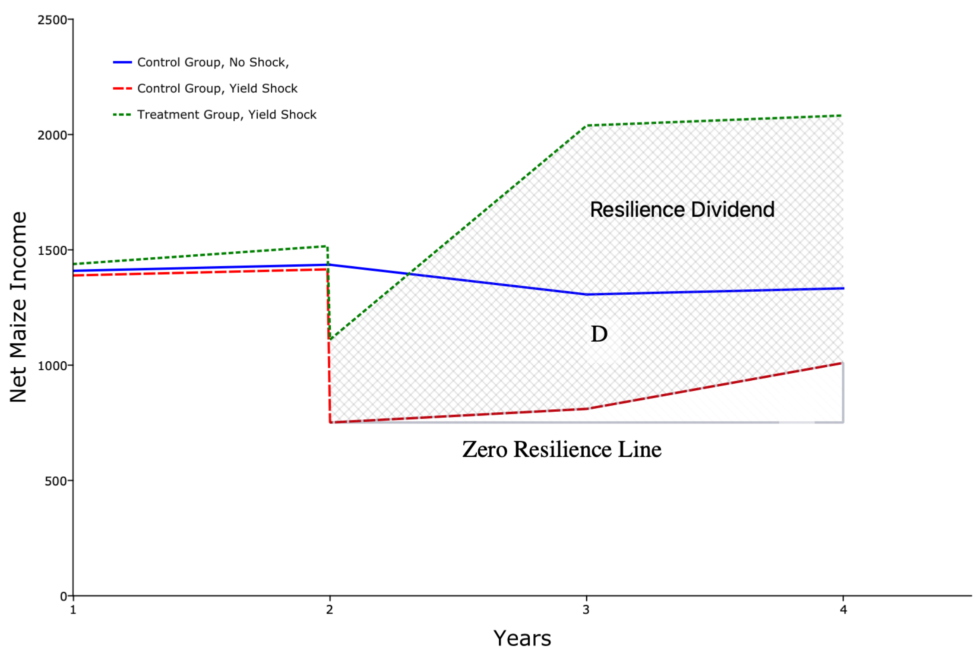

Figure 8: xx

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 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 a catastrophic insurance policy is much higher in a world in which at least a fraction of the population is subject to poverty traps. While our own analysis is based on generated data, we derive some simple rules for considering how frequently data must be collected in the real world to allow reliable measurement of economic resilience.

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).

_ **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.

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.

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.

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.

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.

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.

Appendix 1: Numerical Specification of Dynamic Optimization Model