European Review of Agricultural Economics Vol **40 (2)** (2013) pp. 361–377
doi:10.1093/erae/jbs027
Advance Access Publication 16 August 2012

Estimating risk preferences in the presence of bifurcated wealth dynamics: can we identify static risk aversion amidst dynamic risk responses?

Travis J. Lybbert†*, David R. Just‡ and Christopher B. Barrett‡
†University of California, Davis, USA; ‡Cornell University, Ithaca, NY, USA

Received December 2010; final version revised October 2011; final version accepted January 2012

Review coordinated by Thomas Heckelei

Abstract

Estimating risk preferences is tricky because controlling for confounding factors is difficult. Omitting or imperfectly controlling for these factors can attribute too much observable behaviour to risk aversion and bias estimated preferences. Agents often modify risky decisions in response to dynamic wealth or asset thresholds, where such thresholds exist. Ignoring this dynamic risk response introduces an attribution bias in static estimates of risk aversion. We demonstrate this pitfall using a simple model and a Monte Carlo simulation to explore the implications of this problem for empirical estimation. While an approach that jointly estimates risk preferences and wealth dynamics may remedy the problem by extracting dynamic risk responses from observed behaviour, it is likely to be challenging to implement in broader empirical settings for reasons we discuss.

Keywords: risk, uncertainty, wealth dynamics, risk aversion, risk preference estimation, poverty

JEL classification: D81, O12, D90

1. Introduction

People often take greater risks when facing real prospects of unemployment, hunger, home foreclosure or other imminent perils that they might avoid with a stroke of manufactured good fortune. To invoke a simple image, even a cautious person will jump from a burning building if they believe themselves doomed otherwise. Such an induced risk response is evident in risky sex

*Corresponding author: Giannini Foundation of Agricultural Economics/Department of Agricultural and Resource Economics, One Shields Avenue, Davis, CA 95616, USA. E-mail: tlybbert@ucdavis.edu

work in Kenya (Robinson and Yeh, 2011), illegal migration and consensual participation in human trafficking (Kristof and WuDunn, 2009), and skewness seeking in lottery participation (Yew Kwang, 1965), at the horse track (Golec and Tamarkin, 1998) and among mutual fund managers who gamble with riskier fourth-quarter portfolios in order to catch the market or make 'best fund' lists (Chevalier and Ellison, 1997). In making risky choices, real people seem to factor in how the outcome might change the path dynamic on which they find themselves. To date, however, economists have largely overlooked the effect of background path dynamics on choice under uncertainty. The long tradition of estimating risk preferences based on the moments of return distributions to represent choice-conditioned outcomes considers only the direct payoffs to a risky choice, not the longer term consequences that arise due to known, non-linear path dynamics.

We illustrate this fairly general problem using the example of non-linear wealth dynamics associated with poverty traps that a growing literature suggests exist in at least some developing economies (Dasgupta, 1997; Lybbert *et al*., 2004; Azariadis and Stachurski, 2005; Adato, Carter and May, 2006; Barrett *et al*., 2006; Bowles, Durlauf and Hoff, 2006; Carter and Barrett, 2006; Hoddinott, 2006; Vargas-Hill, 2009). As the empirical evidence grows that there often exist threshold effects in low-income economies where risk issues are especially salient, economists must consider the consequences of such contexts for risk preference estimation. When agents recognise thresholds in underlying wealth dynamics, their valuation of risk responds in part to the underlying dynamics they perceive; risk response is not solely a function of the level of risk aversion that characterises static preferences (Lybbert and Barrett, 2007; Lybbert and Barrett, 2011). Ignoring this dynamic risk response thus introduces an attribution bias in static estimates of risk aversion, and dynamic wealth forces can make it very difficult empirically to ferret out the difference between reactions to dynamic wealth changes and risk preferences.

The state dependence of asset dynamics is central to our argument. The microeconomics of development literature routinely finds that asset accumulation rates can vary non-monotonically with initial asset stocks and, correspondingly, that shocks can have permanent consequences. Consider human capital accumulation in children, for example. Shocks due to natural disasters can have permanent negative impacts on health (Jensen, 2000; Paxson and Schady, 2005; Alderman, Hoddinott and Kinsey, 2006), education (Ferreira and Schady, 2009) and productivity (Currie and Thomas, 2001; Case and Paxson, 2008), suggesting that the returns from investing in children are path-dependent. Investing too little in a child may lead to negative returns if it hampers her physical, cognitive or academic development, but investing enough to nurture the child's development can generate significant positive returns (e.g. future wages or remittances). Evidence that parents in poor countries often invest in children in ways that allow them to specialise (Horowitz and Wang, 2004; Emerson and Souza, 2008) suggests an awareness of seemingly non-linear returns to human capital accumulation. In such a setting, investment patterns are likely driven by risk and time preferences as well as parents' perceptions of these dynamic forces; ignoring the latter would lead to misleading estimates of the former.

While we draw our motivation and framing for this paper from the microeconomics of development, there have been somewhat related efforts to estimate risk preferences in the context of inherently dynamic or intertemporal structure. For example, an established literature has wrestled with theoretical and empirical approaches to modelling risk aversion when decisions have important intertemporal dimensions (e.g. Epstein and Zin, 1991; Bommier, 2007), including a recent effort to separate risk and intertemporal substitution preferences among Kenyan herders (Lybbert and McPeak, 2012). More directly relevant to our approach here, a few economists have used data from game shows that involve risk-taking to estimate risk preferences. Anderson *et al.* (2008) provide a survey of these efforts and offer their own approach to recovering risk preferences with data from the show *Deal Or No Deal* that explicitly accounts for the dynamic stochastic structure of this game. This effort provides useful econometric approaches to reflecting a dynamic structure in risk preference estimation in contrived contexts where this structure is known commonly and in detail. We aim to explore a related problem in which dynamics shape risk-taking, but in a way that sheds light on less contrived and more policy-relevant settings.

Our exploration of risk responses induced by non-linear wealth dynamics and associated implications for empirical preference estimation is distinct from seemingly related strands of the finance literature. Aside from clear differences stemming from our focus on poverty dynamics rather than asset pricing, several distinctions are worth highlighting. Prior work in finance has examined issues of dynamic risk in terms of serial correlation in investment returns (Merton, 1969) or more generally when the return on investment changes stochastically over time (for a recent review, see Munk and Sorensen, 2007). However, while this literature considers choices between investments with returns that change over time, the returns are generally independent of individual wealth or the size of investment. Some have recognised the impact of investment size in determining the return on investment in the management of mutual funds (Indro *et al.*, 1999) or the influence of firm size in determining the return on research and development expenditures (Cohen and Klepper, 1996). In these cases, investment decisions were either assumed to maximise expected profits – and hence risk preferences did not play a direct role in investment – or the investment decisions were treated as exogenous. In contrast to this treatment of dynamic risk in the finance literature, we explore the risk implications that arise when individuals face systematic background currents in personal wealth that influence their risky choices. While Foster and Hart (2009) share a similar concern about dynamic risk dimensions in their recent derivation of a theoretical measure of riskiness based on the probability of bankruptcy, they do not address the relationship between wealth dynamics and risk-taking and say nothing about the empirical implications of such a relationship, which is the void we aim to fill. Lastly, the notion of time-varying risk preferences is captured in several recent models in macroeconomics and finance.[1] The proposed causes of these time-varying preferences, including habit-formation preferences (Campbell and Cochrane, 1999) and consumption commitments (Chetty and Szeidl, 2007), are, however, completely unrelated to the dynamic risk response that may causes risk preferences to (appear to) change over time in our formulation.

In this paper, we build on the insight that observed behaviours often reflect both innate preferences and responses to known (or perceived) wealth dynamics by exploring more explicitly its implications for risk preference estimation. What, if anything, can be done econometrically to avoid misattributing dynamic risk responses to static risk preferences? To address this question, we first construct a simple model to illustrate how a dynamic risk response can bias estimates of risk preferences by attributing to preferences what actually belongs in the structure of wealth creation. Using this model as an analytical framework, we then develop and describe a Monte Carlo simulation that enables us to test different potential econometric remedies to this misattribution problem. Comparison of the resulting risk preference estimates indicates the substantial misattribution bias that can affect such estimates when dynamic risk responses are ignored. While joint estimation of risk preferences and the underlying wealth recursion function can help alleviate this bias, this approach is likely to be challenging to implement with observational data, suggesting that this misattribution problem may be difficult to remedy econometrically.

2. Analytical model

The potential to misattribute a dynamic risk response to an innate risk preference has been raised recently in the literature. As demonstrated in Lybbert and Barrett (2011), when agents face known non-convex wealth dynamics, observed behaviours can misleadingly suggest risk-loving preferences over some asset ranges even when unobservable preferences exhibit moderate risk aversion. That insight raises a crucial distinction between the static concept of risk aversion embedded in preferences and the forward-looking dynamic risk responses manifest in many observed behaviours. In this section, we first develop a stylised analytical model to illustrate several key points. While many of these points appeared in a different analytical form in Lybbert and Barrett (2011), they bear repeating briefly here in order to help motivate and provide the analytical foundation for the Monte Carlo simulation work in Section 3.

We assume that risk-averse, infinitely lived agents face a known wealth dynamic and are offered a gamble in the first period and then live forever with the consequences of the gamble, as shaped by the underlying wealth dynamics. This stark setup helps to convey how agents' valuation of the gamble

1 There is a growing microeconomic literature that seeks to test these models empirically, which generally finds weak or very limited support for these models of time-varying risk preferences. For an example and a review of this literature, see Brunnermeier and Nagel (2008).

is a function of both their innate aversion to risk and their initial wealth relative to the key thresholds that describe these dynamics. If we naïvely ignore agents' dynamic risk response near bifurcating wealth thresholds, risk aversion estimates may be severely biased due to misattribution. Similar to the approach in Lybbert and Barrett (2011), we identify three behavioural regimes. The first prevails for those well below and well above the key wealth threshold at which path dynamics bifurcate. The gamble has little impact on the behaviours of these two cohorts because it is unlikely to affect the underlying wealth dynamics they follow. In contrast, the second and third behavioural regimes prevail for those relatively near the wealth threshold: for those just below and above this point in asset space, the threshold induces a risk-taking response since the outcome of a gamble might shift them onto a different long-run wealth accumulation path.

Suppose that a representative agent's contemporaneous utility is given by the constant relative risk aversion function

$$u(w) = \frac{w^{1-R}}{1-R}, \tag{1}$$

where R is the coefficient of relative risk aversion and w is the agent's wealth. Suppose that the wealth recursion is characterised by a threshold w^0, above which expected wealth is static and below which wealth steadily falls in each subsequent period:

$$w_{t+1} = T(w_t) \equiv \begin{cases} \varphi w_t & \text{if } w_t < w^0 \\ w_t & \text{if } w_t \geq w^0 \end{cases}, \quad \varphi \in (0, 1). \tag{2}$$

Provided that agents accurately perceive this recursion function, their infinite horizon intertemporal utility is given by:

$$U_t(w_t, T(w_t), T(w_{t+1}), \ldots) = \begin{cases} \dfrac{u(w_t)}{1-\delta\varphi^{1-R}} & \text{if } w_t < w^0 \\ \dfrac{u(w_t)}{1-\delta} & \text{if } w_t \geq w^0 \end{cases}, \tag{3}$$

where $\delta \in (0, 1)$ is the discount factor. Given the simple recursion function in equation (2), the dynamic erosion of wealth below w^0 by the depreciation parameter φ (adjusted for the diminishing marginal utility of wealth implied by $1 - R$) simply raises the effective discount rate.

Given this setup, suppose we offer agents in this model a gamble (z, ρ; $-z$, $1-\rho$) which pays out z with probability ρ. Agents' certainty equivalent (C_i) of this gamble depends on their wealth relative to the recursion function in equation (2) and to the size of z according to:

$$\begin{aligned} &U_t(w_{it} + C_i, T(w_{it} + C_i), T(w_{i,t+1}), \ldots) \\ &\quad = EU_t(w_{it} + \tilde{z}, T(w_{it} + \tilde{z}), T(w_{i,t+1}), \ldots | z, \rho), \end{aligned} \tag{4}$$

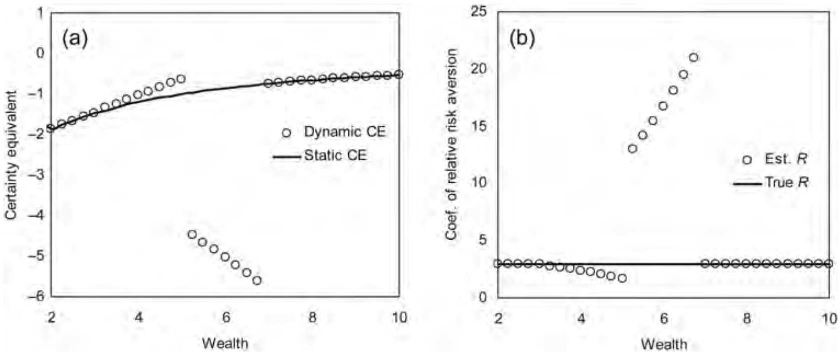

Fig. 1. Static versus dynamic certainty equivalent (a) and naïvely estimated versus true coefficient of relative risk aversion (b) ($\rho = 0.5, z = 1.9, \varphi = 0.95, R = 3, \delta = 0.9$ and $w^0 = 5$).

where $\tilde{z}$ is the random payout of the gamble. Four distinct wealth cohorts ($--$, $-$, $+$, $++$) emerge with the implied certainty equivalent for each wealth cohort given by:

$$C_i^{--} = C_i^{++} = [\rho(w_{it} + z)^{1-R} + (1-\rho)(w_{it} - z)^{1-R}]^{R-1} - w_{it},$$

$$C_i^{-} = \left[\rho\frac{(w_{it} + z)^{1-R}(1 - \delta\varphi^{1-R})}{1-\delta} + (1-\rho)(w_{it} - z)^{1-R}\right]^{R-1} - w_{it},$$

$$C_i^{+} = \left[\rho(w_{it} + z)^{1-R} + (1-\rho)\frac{(w_{it} - z)^{1-R}(1-\delta)}{1-\delta\varphi^{1-R}}\right]^{R-1} - w_{it}.$$

The poorest and richest cohorts value the gamble similarly because z is not big enough to change their position relative to w^0. The valuation of the $-$ and $+$ cohorts reflects their potential to cross this threshold by winning or losing z, respectively.

Figure 1a depicts the certainty equivalents for these wealth cohorts relative to a benchmark static certainty equivalent function with linear wealth dynamics (i.e. $\varphi = 1$). Those with wealth just below the threshold ($w^0 = 5$) value the gamble more because a lucky draw can push them to more favourable dynamic path, while those just above the threshold value it substantially less because the gamble threatens their otherwise safe wealth position. The coefficients of relative risk aversion R implied by the dynamic certainty equivalents (Figure 1b) demonstrate the misattribution bias. If we naïvely estimate R using dynamic certainty equivalents but ignore agents' behavioural response to the wealth recursion, we misattribute all the variation in these certainty equivalents to the curvature of the contemporaneous utility function.[2] The direction of this bias depends on agents' location relative to w^0. Naïve

2 For a related perspective on the potential problems of attributing the curvature implied by individuals' risky decisions exclusively to risk aversion, see Just and Peterson (2003).

estimation wildly exaggerates the degree of risk aversion of those playing it safe just above this threshold and understates the risk aversion of those going for broke just below it.

A common approach to estimating risk preferences involves pooling agents and estimating an average risk preference for a group or sub-group of individuals. Given a uniform wealth distribution, the average degree of risk aversion in this simple analytical model is $\bar{R} = 5.8$, nearly twice the actual degree of risk aversion of $R = 3$. While the direction of the bias in $\bar{R}$ depends on the parameterisation of the model and the distribution of agents across the wealth spectrum (and especially near w^0), the possibility of the misattribution bias in estimated risk preferences is obvious. The average coefficients of risk aversion by wealth cohort are $\bar{R}^{--} = \bar{R}^{++} = 3$, $\bar{R}^{-} = 2.3$ and $\bar{R}^{+} = 16.9$. With a sufficiently discerning sample splitting technique (e.g. Hansen, 2000), we might be able to define these cohorts and estimate these cohort averages, but these still entangle a dynamic risk response with static risk aversion. To dodge this misattribution pitfall, one must control properly for the structure of the underlying dynamics and the risk response that dynamic induces.

3. Monte Carlo simulation

Building on the simple analytical model above, we construct a simulation model to explore the joint estimation of risk preferences and wealth dynamics. As before, individuals in this model are offered the possibility of investing their wealth in a one-time risky investment. They can invest any portion of their wealth in this investment and all face a wealth recursion function after the stochastic investment outcome is realised. We use this model to generate decision data that we then use to demonstrate the pitfalls in risk preference estimation introduced by bifurcated wealth dynamics. We propose two naïve estimation approaches that ignore underlying wealth dynamics and two informed estimation approaches that explicitly account for these dynamics. We compare estimates from these approaches to illustrate the misattribution bias.

In contrast to the simple piecewise-linear recursion function used for analytical tractability in equation (2), the wealth dynamic in this numerical simulation is characterised by a more realistic recursion function of the form

$$w_{t+1} = \tilde{T}(w_t) \equiv \kappa\alpha \frac{1}{1 + \mathrm{e}^{-\nu(w_t - \gamma)}} + (1 - \alpha)w_t + \beta + \varepsilon_t, \tag{6}$$

where w_t is wealth in period t, $\varepsilon_t \sim N(0, \sigma_d^2)$, and κ, α, β, ν and γ are shape parameters, with $\alpha \in [0, 1]$, and all other parameters positive. The first term is a standard logistic function, the curvature of which depends on the ν and γ parameters, while the second term is obviously linear in w. Thus the parameter α measures the degree of non-linearity in the wealth recursion function, with $\alpha = 0$ corresponding to linear expected wealth growth of β, a random walk with drift process. When $\alpha = \beta = 0$, wealth follows a simple random walk process. The κ parameter scales the logistic component of the function so that there may or may not be multiple equilibria in the expected path dynamics. This is the simplest general functional form that allows for both linear and non-linear, unique and multiple equilibria in the underlying path dynamics.

Faced with this known wealth dynamic, individuals have the opportunity in the first period to purchase a risky asset with an instantaneously realised return. Each individual decides how much to invest in this asset by solving:

$$\max_{z \in [0, w_1]} \sum_{t=1}^{L} \delta^t Eu(\tilde{T}^t(w_1 + z(s-1))), \tag{7}$$

where $\tilde{T}^t(\cdot)$ indicates that the wealth recursion function $\tilde{T}(\cdot)$ has been applied to its argument t times in iteration, z is the amount of wealth invested in the risky asset, the purchase price of the asset is normalised to 1, and $s \sim N(\mu_s, \sigma_s^2)$ is the per unit gross return on the risky asset. Parameters are the same across all individuals. The only source of cross-sectional variation is initial wealth.

We draw a sample of N initial wealth observations from a uniform distribution. For each of these $i = 1, \ldots, N$ draws, the utility maximising z_i is determined by using a grid search over the possible values between 0 and w_{1i}. For each value in the grid, a sample of M draws of $\tilde{s}$is used to calculate the expected utility. The z_i yielding the highest expected utility is then recorded as the optimal choice for that observation. Next, the wealth dynamic is applied to the resulting wealth L times in order to generate a wealth time series of L periods following the risky decision. Thus, the simulation generates an initial wealth, an optimal investment decision and a resulting wealth time series for each of the N observations. We then use these generated data to estimate behavioural parameters using naïve and informed approaches.

The first naïve estimation approach assumes that investment decisions are based only on instantaneous utility (i.e. ignores any dynamic risk response) and that individuals share a common coefficient of relative risk aversion. This implies estimation of a unique R that minimises the sum of squared deviations from the first-order conditions[3]:

$$\min_R \sum_{i=1}^{N} \{E[w_{1i} + z_i(s-1)]^{-R}(s-1)\}^2. \tag{8}$$

The second naïve estimation approach also ignores any dynamic risk response, but allows for heterogeneous coefficients of relative risk aversion such that R_i

3 The instantaneous utility function is of the exponential form, and thus the decision problem can be written as $\max_{z \in [0, w_1]} \sum_{t=1}^{L} \delta^t E(\tilde{T}^t(w_1 + z(s-1)))^{1-R}/(1-R)$. For an interior solution, the first-order condition to this problem can be written as $\sum_{t=1}^{L} \delta^t E(\tilde{T}^t(w_1 + z(s-1)))^{-R}$ $\partial\{\tilde{T}^t(w_1 + z(s-1))\}(s-1)/\partial w$. Under the assumption of static wealth, the derivative of the wealth dynamic function is 1, thus the first-order condition under static estimation can be written as $\sum_{t=1}^{L} \delta^t E(w_1 + z(s-1))^{-R}(s-1) = 0$.

is based on w_{1i} and z_i as follows:

$$\max_{z_i \in [0, w_{1i}]} E \frac{[w_{1i} + z_i(s-1)]^{R_i}}{1 - R_i}. \tag{9}$$

In contrast, the informed estimation approaches explicitly account for dynamic risk responses by estimating the wealth recursion function and risk preferences jointly. In contrast to the naïve approaches, which focus simply on instantaneous utility, these informed approaches consider the intertemporal stream of utility and, consequently, estimate the discount rate in addition to risk aversion. The first informed estimation approach assumes we know the true functional form of the recursion and uses the generated panel data of wealth outcomes to estimate its parameters. With this estimated recursion function, this approach then uses a grid search over R and δ to determine the parameters that solve

$$\min_{R,\delta} \sum_{i=1}^{N} \left\{ z_i - \arg\max_{z_i \in [0, w_{1i}]} \sum_{t=1}^{L} \delta^t E u(\hat{\tilde{T}}^t (w_{1i} + z_i(s-1))) \right\}^2. \tag{10}$$

The second informed approach is nearly identical to the first, but instead assumes that we do not know the functional form of the wealth recursion and must approximate it using a sixth-order polynomial (for other polynomial approximations, see Barrett, *et al.*, 2006; Antman and McKenzie, 2007), $w_{t+1} = \phi_0 + \phi_1 w_t + \phi_2 w_t^2 + ... + \phi_6 w_t^6$.

The simulations we report are based on an $L = 5$ decision problem with $N = 500$ observations and $M = 100$ draws. Initial wealth is drawn from the distribution $w_1 \tilde{U}(0, 10)$. The parameters of the wealth dynamic function are $\kappa = 5$, $\alpha = 0.75$, $\beta = 0.5$, $\nu = 2.5$, $\gamma = 3$, $\sigma_d^2 = 0.01$.[4] These parameters imply an unstable equilibrium at $w \approx 3$ (analogous to w^0 in the analytical model above). The risky asset has return $\mu_s = 1.1$ and variance $\sigma_s^2 = 0.25$. The discount factor is set to $\delta = 0.95$, and relative risk aversion is $R = 3$.

The results of the four estimation approaches using this generated data are shown in Table 1. Columns 1 and 2 illustrate the misattribution bias inherent in the naïve estimation approaches. With a common R (column 1), this bias is severe. The mean of the individual estimates of R (column 2) is not statistically different than the true value of R as shown by the p-value for that estimate, but the misattribution bias in these individual estimates remains severe near the unstable threshold. Figure 2 illustrates the bias in these individual

4 The assumption of a uniform initial wealth distribution suggests a subtle inconsistency with the assumed values of the wealth recursion function parameters. The latter imply three stochastic dynamic equilibria – two stable, one unstable – which in principle could affect the cross-sectional wealth distribution at any point in time. Given the random component to the wealth recursion function, however, the underlying distribution is ergodic. We therefore opt for the computationally simpler approach of assuming a uniform initial wealth distribution rather than generating an initial wealth distribution and then sampling from it.

Table 1. Results for static and informed estimation approaches

	True parameters	Naïve estimation (1) Common *R*	(2) Individual *R*	Informed estimation (3) True wealth recursion functional form[a]	(4) Approximation of wealth recursion function[b]
Relative risk aversion (*R*)	3.00	13.18 (0.000)[c]	1.91[d] (0.486)[c]	3.18 (0.000)[c]	2.14 (0.008)[c]
Discount factor (δ)	0.95	–	–	0.85 (0.001)[c]	0.65 (0.012)[c]

[a]The estimated parameters in the true wealth dynamic functional form (equation (6)) are so precisely estimated as to be indistinguishable from the true parameter values at two decimal places.
[b]This estimated approximation is given by sixth-order polynomial: these estimated parameters are all significant at the 1 per cent level. See Figure 3 for a comparison between this approximation and the true recursion function.
[c]*p*-Values in parentheses test the null that the estimates equal the true parameters. *p*-Values in columns 3 and 4 are based on standard errors generated by 1,000 bootstrapped samples.
[d]This is the mean of individual estimates with a standard deviation of 1.56.

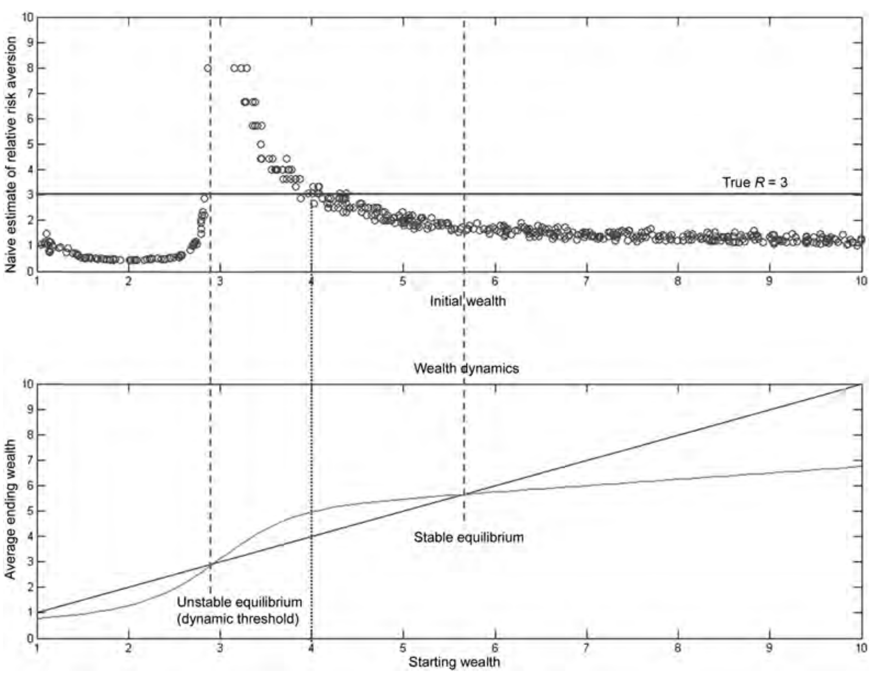

Fig. 2. The impact of wealth dynamics as captured by recursion function (bottom panel) on naïve individual estimates of risk aversion (top panel). Note: Dashed vertical lines correspond to equilibria as labelled.

estimates by graphing them in the initial wealth space relative to the true value of $R = 3$. Just above the dynamic threshold implied by the unstable equilibrium, we significantly overestimate risk aversion due to extremely cautious dynamic risk responses in this region stemming from the threat that the gamble could push them below the threshold. Elsewhere in the initial wealth space, we underestimate risk aversion due to more ambitious (i.e. risk loving) dynamic risk responses.[5] We correctly estimate risk aversion only around an initial wealth of about 4, which is the only point above the dynamic threshold where $\tilde{T}'(w_t) \approx 1$. More generally, any interpretation of the average of these individual risk aversion estimates (1.91 in this case) should appreciate the influence of outlier estimates for individuals near the dynamic threshold and acknowledge that relatively small changes in the structure of a context could dramatically change these outliers and thereby the average individual risk estimate.

Columns 3 and 4 of Table 1 show the informed estimates based on joint estimation of the underlying recursion function and risk preferences. Given the non-linear nature of this joint estimation technique, there is no closed form solution for the standard errors on these estimates, so we bootstrap to estimate standard errors using 1,000 simulated samples of size $N = 500$ as the basis for the p-values reported in Table 1. These standard errors are vanishingly small, suggesting that the estimation procedure is extremely accurate given the size of the sample and the chosen variance parameters. A comparison of results in columns 3 and 4 indicates how crucial it is to accurately estimate the recursion function in order to extract the dynamic risk response from our simulated gambling behaviour. When we estimate the true functional form for this recursion (column 3), our risk aversion estimate is much closer to $R = 3$, but still statistically significantly different from this true value. When we assume we do not know the true functional form of the recursion, we approximate it instead using a sixth-order polynomial, which appears to closely match the true function (Figure 3). Despite the apparent precision of our recursion function approximation, our estimates slip further from their true values (column 4). While our risk aversion estimate is closer to the true value of $R = 3$ than the naïve estimate in column 2, it is also estimated much more precisely and is therefore statistically significantly different from this true value. This suggests that our ability to alleviate the misattribution problem econometrically when estimating risk preferences may be limited by even seemingly small errors in our recursion function estimation.

4. Discussion

The estimation of the recursion function influences risk preference estimates in this joint estimation approach in several important ways, which highlight

5 We do not compare the direction of this misattribution bias in this simulation with that of the analytical model because the direction of bias due to an ignored dynamic risk response hinges on the parameterisation and is therefore not predictable *ex ante*.

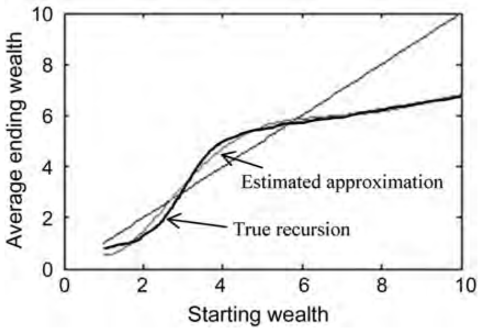

Fig. 3. True wealth recursion function (black) with the estimated sixth-order polynomial approximation (grey).

how complicated practical implementation of this strategy may be. As just mentioned, risk estimates are quite sensitive to even minor imperfections in our estimation of wealth dynamics. Specifically, the way the recursion function enters the optimisation problem in equation (10) ensures that small imperfections in approximation are quickly magnified and can limit our ability to accurately and reliably extract the dynamic risk response. Small changes in the dynamic function loom large over time in wealth and consumption decisions through compounding. Thus, small errors in the estimated recursion function can have large impacts on subsequent estimation of risk aversion parameters, which generally enter utility functions in an exponential, rather than multiplicative or additive, form. Errors near critical thresholds can especially hamper risk aversion estimates. To be clear, such errors do not introduce imprecision in statistical estimates. We can produce very precise estimates in the presence of these errors (recall column 4 of Table 1). But recursion function errors imply that these precise estimates may also be biased by an inaccurate remedy to the misattribution problem.

To explore these relationships further, we ran several additional simulation and estimation routines. First, we used the same simulation as above and approximated the recursion function as a third-order polynomial. This less flexible approximation is worse where it matters most: at the crossing of the 45° line, which determines the location of the dynamic threshold and strongly determines how much of the dynamic risk response can be extracted from the observed behaviour (recall Figure 2). Consequently, this simpler approximation yields risk estimates that are yet farther from the true values and virtually indistinguishable from the naïve estimate in column 2. In this case, we underestimate the risk parameter as the approximation introduces error in the recursion function, but such errors could also yield overestimates of this parameter depending on the direction of the error at the point of the actual threshold. As an alternative approximation approach, semi- or non-parametric estimation of the wealth recursion (e.g. Lybbert *et al.*, 2004) could improve the fit of the recursion function and thereby the estimates of R and δ.

Second, the location of the dynamic threshold relative to the density of the data has important implications. We ran a simulation with the recursion function calibrated, so it implied a lower threshold without changing the wealth distribution. This reduced the power of the data to recover the recursion function, increased the bias of the risk and discount factor estimates and switched the direction of the bias in the risk parameter. In general, the farther the threshold is from the centre of the wealth distribution, the harder it is to locate the threshold with accuracy and, hence, the greater the bias that remains in the risk estimate after trying to jointly estimate the wealth dynamic. The worst case scenario is one where a threshold lurks just on the edge of the wealth distribution for a given sample and therefore induces a dynamic risk response, but is impossible to identify empirically.

Third, the severity of this problem, which is due to fewer observations in the vicinity of the threshold with which to estimate the crossing point, also depends on the strength of the dynamic forces embedded in the recursion function. With a particularly strong dynamic, the function crosses the 45° line at an angle that is much higher than 45°. Fewer data points are required to identify such a crossing point. With a more subtle dynamic, the function is flatter at the crossing point and requires substantially more data points to locate the threshold accurately. Since we are unlikely to know the steepness of the recursion at the crossing point *ex ante* – particularly, if we do not know the location of the crossing point – we are also unlikely to know whether we need more or fewer observations to successfully estimate risk preferences while jointly estimating the recursion function. Moreover, the availability of data near the threshold may itself be a function of the strength of the dynamics. If the underlying dynamic forces in a system are particularly strong relative to stochastic shocks, a distinctly bi-modal wealth distribution can emerge with relatively few data points near any unstable equilibria, further complicating the estimation of the recursion function near critical thresholds where dynamic risk responses may be strongest.[6]

Fourth, implementing this joint estimation approach in a non-simulated empirical setting raises even more fundamental identification issues. Using alternative specifications and experimental data, Andersen *et al*. (2008) find that the parameters for time and risk preferences are substitutes for one another: estimating one requires some severe assumptions about the other, and the estimate of one is sensitive to assumptions about the form of the other. Just and Just (2009) provide a more general treatment of the challenges of identifying risk preferences empirically. In the analysis in this paper, we are able to make the necessary strict assumptions for identification because we take a simulation approach in which we know the exact functional form for these parameters. Thus, plots of the sum of squared error functions for the

6 We have explored the effects of this complication on the estimation of the behavioural parameters via simulation by burning in the recursion function several periods before generating the data used to estimate the parameters in the model. This quickly squeezes out information near bifurcation thresholds, making it difficult to get decent preference estimates even when explicitly controlling for the full dynamic structure.

above estimation routines reveal well-behaved convex functions in the parameter space. The same is less likely to be true in less controlled empirical settings even in the absence of non-convex wealth dynamics. Identification in non-simulation settings for the joint estimation of a recursion function and risk and time preferences we develop in this paper will be challenging.

Finally, two potential dimensions of how wealth dynamics work in practice further complicate the estimation of wealth recursion functions – whether in isolation or jointly with risk preferences. First, the dynamic path of wealth or asset accumulation is, in many settings, conditioned on ability or other unobservable individual attributes. The possibility of heterogeneous wealth dynamics can severely complicate the estimation of the recursion function (Santos and Barrett, 2006; Antman and McKenzie, 2007), which could limit any improvements in risk preference estimation. Second, dynamic risk responses are induced by perceptions of wealth dynamics rather than objective recursion functions *per se*. This is similar to individuals relying on subjective probability weighting functions rather than objective probabilities when making decisions under risk (Tversky and Kahneman, 1992), but at this point we know far too little about these subjective perceptions to expect any systematic patterns to emerge across individuals.

5. Conclusion

Estimating risk preferences is notoriously tricky because controlling for confounding factors is difficult. Our objective in this paper has been to showcase non-linear underlying path dynamics as an overlooked, but potentially important, confounding factor. When thresholds in these underlying dynamics are sufficiently distinct that they are apparent to agents in the system, thresholds may induce a dynamic risk response. Near such a threshold, part of an agent's valuation of risk is a response to the underlying dynamics she faces or perceives, rather than being simply a function of her level of risk aversion. Ignoring this dynamic risk response thus introduces an attribution bias in static estimates of risk aversion. We illustrate this general problem using the case of bifurcated wealth dynamics associated with poverty traps and demonstrate a potential econometric remedy. While joint estimation of wealth dynamics and risk preferences is feasible in a controlled simulation setting, implementation in broader empirical settings is likely to be very challenging. Consequently, we would advocate for additional research in controlled settings to better understand the nature of this problem. With carefully designed economic experiments, for example, one could continue to explore structural estimation techniques to account for dynamic risk responses. Advances in empirical methods to characterise asset dynamics could similarly improve our ability to tease apart dynamic risk responses and static risk aversion.

We conclude with a few final thoughts. First, does the misattribution of dynamic risk responses to static risk aversion really matter? As a theoretical matter, this seems to be an important distinction and speaks to the importance of balancing preferences and structure. Indeed, the observation that implicit wealth dynamics might directly shape preferences appears in Friedman and Savage's (1948) classic work. Moreover, dynamic misattribution bias seems relevant empirically. In particular, this potential source of bias in estimates of risk aversion highlights the importance of careful consideration and estimation of the structural features of behavioural models of decision-making. As a policy matter in the poverty traps setting we use to illustrate the problem, since both the magnitude and the direction of dynamic risk responses can change dramatically as wealth or assets change, misattribution could render the simulation of policies that create infra-marginal welfare impacts especially inaccurate. Even if a given policy changes wealth or assets marginally or not at all, our understanding of its effects might be hampered by the misattribution bias in average estimates of risk aversion.

Understanding non-linear dynamics is an important frontier of economic research, often with clear implications for the design and implementation of policy in areas as distinct as poverty reduction, environmental management, storable commodity markets and business cycles. We argue that while characterising these dynamics is important, more needs to be done to understand how these forces shape how individuals make decisions over time. In many of contexts, individuals and households surely appreciated the dynamics that shape the average wealth accumulation paths they see around them long before economists attempted to characterise them formally. In recognition of our late arrival on the scene, we should at least acknowledge the possibility of behavioural wrinkles induced by these dynamics.

References

Adato, M., Carter, M. and May, J. (2006). Exploring poverty traps and social exclusion in South Africa using qualitative and quantitative data. *The Journal of Development Studies* 42(2): 226–247.

Alderman, H., Hoddinott, J. and Kinsey, B. (2006). Long term consequences of early childhood malnutrition. *Oxford Economic Papers* 58(3): 450–474.

Andersen, S., Harrison, G. W. and Lau, M. I. (2008). Risk aversion in game shows. In: J. C. Cox and G. W. Harrison (eds), *Risk Aversion in Experiments*. New York: JAI Press.

Andersen, S., Harrison, G. W., Lau, M. I. and Rutström, E. E. (2008). Eliciting risk and time preferences. *Econometrica* 76(3): 583–618.

Antman, F. and McKenzie, D. (2007). Poverty traps and nonlinear income dynamics with measurement error and individual heterogeneity. *Journal of Development Studies* 43(6): 1057–1083.

Azariadis, C. and Stachurski, J. (2005). Poverty traps. In: P. Aghion and S. N. Durlauf (eds), *Handbook of Economic Growth*. Amsterdam: Elsevier North Holland.

Barrett, C. B., Marenya, P. P., McPeak, J. G., Minten, B., Murithi, F. M., Kosura, W. O., Place, F., Randrianarisoa, J. C., Rasambainarivo, J. and Wangila, J. (2006). Welfare dynamics in rural Kenya and Madagascar. *Journal of Development Studies* 42(2): 248–277.

Bommier, A. (2007). Risk aversion, intertemporal elasticity of substitution and correlation aversion. *Economics Bulletin* 4(29): 1–8.

Bowles, S., Durlauf, S. N. and Hoff, K. R. (2006). *Poverty Traps*. Princeton, NJ: Princeton University Press.

Brunnermeier, M. and Nagel, S. (2008). Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals' asset allocation. *American Economic Review* 98(3): 713–736.

Campbell, J. and Cochrane, J. (1999). By force of habit: a consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107(2): 205–251.

Carter, M. R. and Barrett, C. B. (2006). The economics of poverty traps and persistent poverty: an asset-based approach. *Journal of Development Studies* 42(2): 178–199.

Case, A. and Paxson, C. (2008). Stature and status: height, ability, and labor market outcomes. *The Journal of Political Economy* 116(3): 499.

Chetty, R. and Szeidl, A. (2007). Consumption commitments and risk preferences. *The Quarterly Journal of Economics* 122(2): 831–877.

Chevalier, J. and Ellison, G. (1997). Risk taking by mutual funds as a response to incentives. *Journal of Political Economy* 105(6): 1167–1200.

Cohen, W. M. and Klepper, S. (1996). A reprise of size and R&D. *The Economic Journal* 106(437): 925–951.

Currie, J. and Thomas, D. (2001). Early test scores, socioeconomic status, school quality and future outcomes. *Research in Labor Economics* 20: 103–132.

Dasgupta, P. (1997). Nutritional status, the capacity for work, and poverty traps. *Journal of Econometrics* 77(1): 5–37.

Emerson, P. M. and Souza, A. P. (2008). Birth order, child labor, and school attendance in Brazil. *World Development* 36(9): 1647–1664.

Epstein, L.-G. and Zin, S.-E. (1991). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis. *Journal of Political Economy* 99(2): 263–286.

Ferreira, F. and Schady, N. (2009). Aggregate economic shocks, child schooling, and child health. *World Bank Policy Research Working Paper*, No. 4701.

Foster, D. and Hart, S. (2009). An operational measure of riskiness. *Journal of Political Economy* 117(5): 785–814.

Friedman, M. and Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of Political Economy* 56(4): 279–304.

Golec, J. and Tamarkin, M. (1998). Bettors love skewness, not risk, at the horse track. *Journal of Political Economy* 106(1): 205–225.

Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica* 68(3): 575–603.

Hoddinott, J. (2006). Shocks and their consequences across and within households in rural Zimbabwe. *Journal of Development Studies* 42(2): 301–321.

Horowitz, A. and Wang, J. (2004). Favorite son? Specialized child laborers and students in poor LDC households. *Journal of Development Economics* 73(2): 631–642.

Indro, D. C., Jiang, C. X., Hu, M. Y. and Lee, W. Y. (1999). Mutual fund performance: does fund size matter? *Financial Analysts Journal* 55(3): 74–87.

Jensen, R. (2000). Agricultural volatility and investments in children. *American Economic Review* 90(2): 399–404.

Just, D. R. and Peterson, H. H. (2003). Diminishing marginal utility of wealth and calibration of risk in agriculture. *American Journal of Agricultural Economics* 85(5): 1234–1241.

Just, R. E. and Just, D. R. (2009). Global identification of risk preferences with revealed preference data. *Journal of Econometrics* 162(1): 6–17.

Kristof, N. D. and WuDunn, S. (2009). *Half the Sky: Turning Oppression into Opportunity for Women Worldwide*. New York: Knopf.

Lybbert, T. J. and Barrett, C. B. (2007). Risk responses to dynamic asset thresholds. *Review of Agricultural Economics* 29(3): 412–418.

Lybbert, T. J. and Barrett, C. B. (2011). Risk taking behavior in the presence of nonconvex asset dynamics. *Economic Inquiry* 49(4): 982–988.

Lybbert, T. J., Barrett, C. B., Desta, S. and Coppock, D. L. (2004). Stochastic wealth dynamics and risk management among a poor population. *Economic Journal* 114(498): 750–777.

Lybbert, T. J. and McPeak, J. (2012). Risk and intertemporal substitution: Livestock portfolios and off-take among Kenyan pastoralists. *Journal of Development Economics* 97(2): 415–426.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: the continuous-time case. *The Review of Economics and Statistics* 51(3): 247–257.

Munk, C. and Sorensen, C. (2007). Optimal real consumption and investment strategies in dynamic stochastic economies. In: B. S. Jensen and T. Palokangas (eds), *Stochastic Economic Dynamics*. Copenhagen, Denmark: CBS Press, 271–316.

Paxson, C. and Schady, N. (2005). Child health and economic crisis in Peru. *The World Bank Economic Review* 19(2): 203.

Robinson, J. and Yeh, E. (2011). Transactional sex as a response to risk in western Kenya. *American Economic Journal: Applied Economics* 3(1): 35–64.

Santos, P. and Barrett, C. B. (2006). Heterogeneous wealth dynamics: on the roles of risk and ability. Working Paper. Cornell University.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297–323.

Vargas-Hill, R. (2009). Using stated preferences and beliefs to identify the impact of risk on poor households. *Journal of Development Studies* 45(2): 151–171.

Yew Kwang, N. (1965). Why do people buy lottery tickets? Choices involving risk and the indivisibility of expenditure. *Journal of Political Economy* 73(Oct.): 530–535.