

**ESTIMATING RISK PREFERENCES IN THE PRESENCE OF BIFURCATED WEALTH DYNAMICS:
DO WE MISATTRIBUTE DYNAMIC RISK RESPONSES TO STATIC RISK AVERSION?**

Travis J. Lybbert, UC Davis, tlybbert@ucdavis.edu

David R. Just, Cornell University

Christopher B. Barrett, Cornell University

March 2010

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 behavior to risk aversion and bias estimated preferences. Agents often modify risky decisions in response to dynamic wealth or asset thresholds, where they 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. Joint estimation of risk preferences and wealth dynamics may remedy the problem, but can be empirically challenging.

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

JEL Code: D81, O12, D90

**ESTIMATING RISK PREFERENCES IN THE PRESENCE OF BIFURCATED WEALTH DYNAMICS:
DO WE MISATTRIBUTE DYNAMIC RISK RESPONSES TO STATIC RISK AVERSION?**

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 work in Kenya (Robinson and Yeh, 2008), 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 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, nonlinear path dynamics.

We illustrate this fairly general problem using the example of non-linear wealth dynamics associated with poverty traps that based on a growing literature exist in at least some developing economies (Adato et al., 2006; Azariadis and Stachurski, 2005; Barrett et al., 2006; Bowles et al., 2006; Carter and Barrett, 2006; Dasgupta, 1997; Hoddinott, 2006; Lybbert et al., 2004; 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 recognize thresholds in underlying wealth dynamics, their valuation of risk responds in part to the underlying dynamics she perceives; risk response is not solely a function of the level of risk aversion that characterizes her static preferences (Lybbert and Barrett, 2007; Lybbert and Barrett, forthcoming). Ignoring this dynamic risk response thus introduces an attribution bias in static

estimates of risk aversion. We demonstrate this potential pitfall. Specifically, we show that dynamic wealth forces can make it very difficult empirically to ferret out the difference between reactions to dynamic wealth changes and risk preferences.

Our exploration of potential 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 recognized 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 maximize 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 cause risk preferences to (appear to) change over time in our formulation.

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 Brunneimeier and Nagel (2008).

We first construct a simple analytical model to illustrate how a dynamic risk response can bias estimates of risk preferences. We then turn to a Monte Carlo simulation to explore the implications of this problem for empirical estimation. 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 dodge this attribution bias, risk preference estimates are sensitive to the fit of the estimated recursion function.

2. ANALYTICAL MODEL

The potential to misattribute a dynamic risk response to an innate risk preference can be easily demonstrated with a stylized analytical model. In this model, risk averse agents who face a known wealth dynamic 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. Agents' valuation of the gamble is a function of both their innate aversion to risk and their initial wealth relative to these dynamics. If we naïvely ignore agents' dynamic risk response near bifurcating wealth thresholds, risk aversion estimates may be severely biased due to misattribution.

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

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

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

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

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

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

where $\delta \in (0,1)$ is the discount factor. Given the simple recursion function in (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 $(z, \rho; -z, 1-\rho)$, 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 (2) and to the size of z according to:

$$(4) \quad U_t(w_{it} + C_i, T(w_{it} + C_i), T(w_{i,t+1}), \dots) = EU_t(w_{it} + \tilde{z}, T(w_{it} + \tilde{z}), T(w_{i,t+1}), \dots | z, \rho)$$

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:

$$(5) \quad \begin{aligned} C_i^{--} &= C_i^{++} = \left[\rho(w_{it} + z)^{1-R} + (1-\rho)(w_{it} - z)^{1-R} \right]^{R-1} - w_{it} \\ C_i^{-} &= \left[\rho \frac{(w_{it} + z)^{1-R} (1 - \delta \varphi^{1-R})}{1 - \delta} + (1-\rho)(w_{it} - z)^{1-R} \right]^{R-1} - w_{it} \\ C_i^{+} &= \left[\rho(w_{it} + z)^{1-R} + (1-\rho) \frac{(w_{it} - z)^{1-R} (1 - \delta)}{1 - \delta \varphi^{1-R}} \right]^{R-1} - w_{it} \end{aligned}$$

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.

Panel (a) of figure 1 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 favorable 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 1(b)) demonstrate the misattribution bias. If we naïvely estimate R using dynamic certainty equivalents

but ignore agents' behavioral 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 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 parameterization of the model and the distribution of agents across the wealth spectrum (and especially near w^0), the possibility of 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 realized. 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 misattribution bias.

² For a related perspective on the potential problems of attributing this curvature exclusively to risk aversion see Just and Peterson (2003).

In contrast to the simple piecewise-linear recursion function in (2), the wealth dynamic in this simulation is characterized by a recursion function of the form

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

where w_t is wealth in period t , $\varepsilon_t \sim N(0, \sigma_d^2)$, and $\kappa, \alpha, \beta, \nu$ 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 nonlinearity 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 nonlinear, 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 realized return. Each individual decides how much to invest in this asset by solving:

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

where $\tilde{T}^t(\cdot)$ indicates 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 normalized 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 $i=1, \dots, N$ draws, the utility maximizing 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 behavioral parameters using naïve and informed approaches.

The first naïve estimation approach assumes that investment decisions are only based 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 minimizes the sum of squared first order conditions:³

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

The second naïve estimation approach also disallows any dynamic risk response, but allows for heterogenous coefficients of relative risk aversion such that R_i is based on w_{1i} and z_i as follows:

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

In contrast, the informed estimation approaches explicitly control for any dynamic risk response by extracting the wealth recursion function before estimating risk preferences. The first informed estimation approach assumes we know the true functional form of this function and uses the generated wealth panel data to estimate its parameters. With this estimated recursion function, this approach then uses a grid search over R and δ ⁴ to determine the parameters that solve

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

⁴ The discount factor was assumed away in the static approaches because they focus on instantaneous utility.

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

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 6th order polynomial (for other polynomial approximations see Antman and McKenzie, 2007; Barrett et al., 2006).

The simulations we report are based on a $L = 5$ decision problem with $N = 500$ observations and $M = 100$ draws. Initial wealth is drawn from the distribution $w_1 \sim U(0,10)$. The parameters of the wealth dynamic function are $\kappa = 5$, $\alpha = .75$, $\beta = 0.5$, $\nu = 2.5$, $\gamma = 3$, $\sigma_d^2 = 0.01$.⁵ 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 this bias is severe. Although the mean of the individual estimates of R is not statistically different than the true value of R , misattribution bias is again severe near the unstable threshold. As shown in figure 2, these individual R 's are significantly underestimated wherever initial wealth given the (ignored) recursion implies expected wealth erosion in the next period. These are also underestimated when expected wealth accumulation is positive and the risk of falling below the threshold is trivial. As in figure 1, they are overestimated when individuals are close enough to the threshold that the gamble could easily push them below it.⁶

The informed estimation results in column 3 have been stripped of the dynamic risk response and are consequently much closer to the true value of $R=3$. This approach assumes we

⁵ 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.

⁶ 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 parameterization and is therefore not predictable *ex ante*.

know the true form of the recursion function, which enables us to estimate the function very precisely. Indeed, this estimated recursion function is graphically indistinguishable from the true function in figure 2. As a result, we can effectively extract the dynamic risk response and produce point estimates of R and δ that are close to their true values.⁷ Unfortunately, implementing this remedy in practice may be wrought with complications. One complication emerges from our second informed estimation approach that approximates the recursion function. The estimated polynomial approximation to the recursion function is shown in figure 3. Although the approximation appears to be close to the true recursion, it is not perfect and the estimates in column 4 seem quite sensitive to even these minor imperfections. The way the recursion function enters the optimization problem in (10) ensures that small imperfections in approximation are quickly magnified and can limit our ability to accurately extract the dynamic risk response. In effect, small changes in the dynamic function create a festering undercurrent for wealth that over time looms large in wealth and consumption decisions through compounding. 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 (recall figure 2).

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 , but further complications lurk. The possibility of heterogeneous (e.g., ability-conditioned) wealth dynamics could severely complicate the estimation of the recursion function (Antman and McKenzie, 2007; Santos and Barrett, 2006). Lastly, 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. This could complicate the estimation of the recursion function near critical thresholds where dynamic risk responses may be strongest.⁸

⁷ We have omitted the calculation of standard errors. Given the non-linear nature of the estimation technique, there is no closed form solution for the standard error function, and the time required for using bootstrap methods is prohibitive. Over several estimation runs, however, the results were very stable, suggesting that the estimates are robust.

⁸ We have explored the effects of this complication on the estimation of the behavioral parameters via simulation by burning-in the recursion function several periods before generating the data used to estimate the parameters in the

4. 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.

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 behavioral models of decision making. As a policy matter in the poverty traps setting we use to illustrate the problem, since both the magnitude and 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 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. It seems reasonable that in many of these contexts individuals and households appreciated 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.

dynamic forces that shape the path long before economists understand them. In recognition of our late arrival on the scene, we should at least acknowledge the possibility of behavioral wrinkles induced by these dynamics.

References

- Adato M, Carter M, May J. Exploring poverty traps and social exclusion in South Africa using qualitative and quantitative data. *The Journal of Development Studies* 2006;42; 226-247.
- Antman F, McKenzie D. Poverty Traps and Nonlinear Income Dynamics with Measurement Error and Individual Heterogeneity. *Journal of Development Studies* 2007;43; 1057-1083.
- Azariadis C, Stachurski J. Poverty Traps. In: Aghion P, Durlauf SN (Eds), *Handbook of Economic Growth*. North Holland; 2005.
- Barrett CB, Marenya PP, McPeak JG, Minten B, Murithi FM, Kosura WO, Place F, Randrianarisoa JC, Rasambainarivo J, Wangila J. Welfare Dynamics in Rural Kenya and Madagascar. *Journal of Development Studies* 2006;42; 248-277.
- Bowles S, Durlauf SN, Hoff KR. *Poverty traps*. Princeton University Press: Princeton, N.J.; 2006.
- Brunnermeier M, Nagel S. Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals' asset allocation. *American Economic Review* 2008;98; 713-736.
- Campbell J, Cochrane J. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 1999;107; 205-251.
- Carter MR, Barrett CB. The Economics of Poverty Traps and Persistent Poverty: An Asset-Based Approach. *Journal of Development Studies* 2006;42; 178-199.
- Chetty R, Szeidl A. Consumption Commitments and Risk Preferences*. *The Quarterly Journal of Economics* 2007;122; 831-877.
- Chevalier J, Ellison G. Risk Taking by Mutual Funds as a Response to Incentives. *Journal of Political Economy* 1997;105; 1167-1200.
- Cohen WM, Klepper S. A reprise of size and R & D. *The Economic Journal* 1996;106; 925-951.
- Dasgupta P. Nutritional status, the capacity for work, and poverty traps. *Journal of Econometrics* 1997;77; 5-37.
- Foster D, Hart S. An operational measure of riskiness. *Journal of Political Economy* 2009;117.
- Friedman M, Savage LJ. The Utility Analysis of Choices Involving Risk. *Journal of Political Economy* 1948;56; 279-304.
- Golec J, Tamarkin M. Bettors Love Skewness, Not Risk, at the Horse Track. *Journal of Political Economy* 1998;106; 205-225.

- Hansen BE. Sample Splitting and Threshold Estimation. *Econometrica* 2000;68; 575-603.
- Hoddinott J. Shocks and Their Consequences across and within Households in Rural Zimbabwe. *Journal of Development Studies* 2006;42; 301-321.
- Indro DC, Jiang CX, Hu MY, Lee WY. Mutual fund performance: does fund size matter? *Financial Analysts Journal* 1999;55; 74-87.
- Just DR, Peterson HH. Diminishing Marginal Utility of Wealth and Calibration of Risk in Agriculture. *American Journal of Agricultural Economics* 2003;85; 1234-1241.
- Kristof ND, WuDunn S. *Half the Sky: Turning Oppression Into Opportunity for Women Worldwide*. Alfred A. Knopf Inc; 2009.
- Lybbert T-J, Barrett C-B. Risk Responses to Dynamic Asset Thresholds. *Review of Agricultural Economics* 2007;29; 412-418.
- Lybbert TJ, Barrett CB. Risk taking behavior in the presence of nonconvex asset dynamics. *Economic Inquiry* forthcoming.
- Lybbert TJ, Barrett CB, Desta S, Coppock DL. Stochastic Wealth Dynamics and Risk Management Among a Poor Population. *Economic Journal* 2004;114; 750-777.
- Merton RC. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics* 1969;51; 247-257.
- Munk C, Sorensen C. Optimal Real Consumption and Investment Strategies in Dynamic Stochastic Economies. In: Jensen BS, Palokangas T (Eds), *Stochastic economic dynamics*. CBS Press: Copenhagen, Denmark; 2007. p. 271-316.
- Robinson J, Yeh E. *Transactional Sex as a Response to Risk in Western Kenya*. Munich Personal RePEc Archive No. 7405. 2008.
- Santos P, Barrett CB. *Heterogeneous wealth dynamics: On the roles of risk and ability*. Cornell University; 2006.
- Vargas-Hill R. Using stated preferences and beliefs to identify the impact of risk on poor households. *Journal of Development Studies* 2009;45; 151-171.
- Yew Kwang N. Why Do People Buy Lottery Tickets? Choices Involving Risk and the Indivisibility of Expenditure. *Journal of Political Economy* 1965;73; 530-535.

Table 1 Results for Static and Informed Estimation Approaches

	True Parameters	Naïve Estimation		Informed Estimation	
		(1) Common R	(2) Individual R	(3) True Wealth Recursion Functional Form ^c	(4) Approximation of Wealth Recursion Function ^d
Relative Risk Aversion (R)	3.00	13.18	1.91 ^a	3.18	2.14
Discount Factor (δ)	0.95	-	-	0.85	0.65

^a This is the mean of individual estimates with standard deviation is 1.56.

^b P-values in parentheses of test that estimates equal true parameters (based on standard errors generated by 1000 bootstrapped samples).

^c 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.

^d This estimated approximation is given by 6th order polynomial: \dots . These estimated parameters are all significant at the 1% level. See Figure 2 for a comparison between this approximation and the true recursion function.

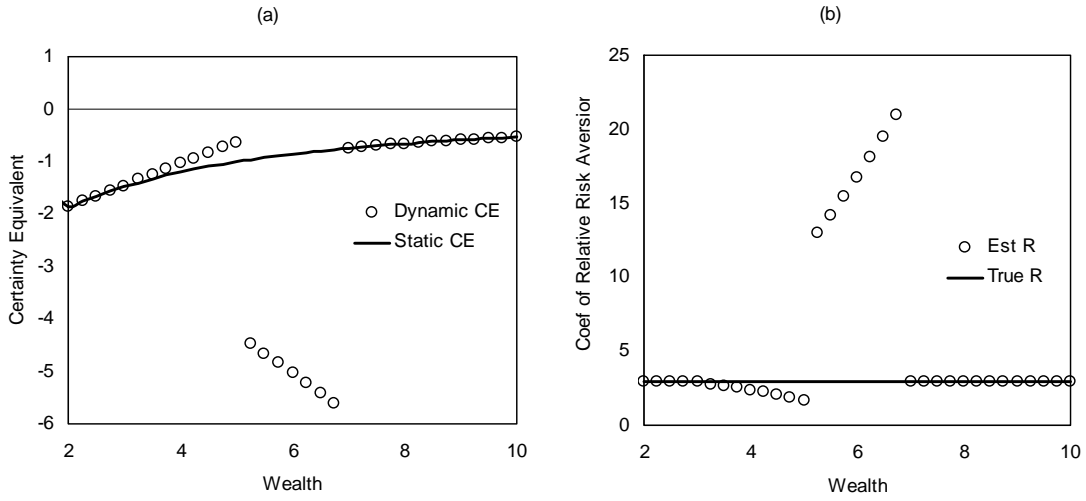


Figure 1 Static versus dynamic certainty equivalent (a) and naïvely estimated versus true coefficient of relative risk aversion (b) ($\rho=0.5$, $z=1.9$, $\varphi=0.95$, $R=3$, $\delta=0.9$, and $w^0=5$)




Figure 2 The Impact of Wealth Dynamics on Naïve Individual Estimates of Risk Aversion

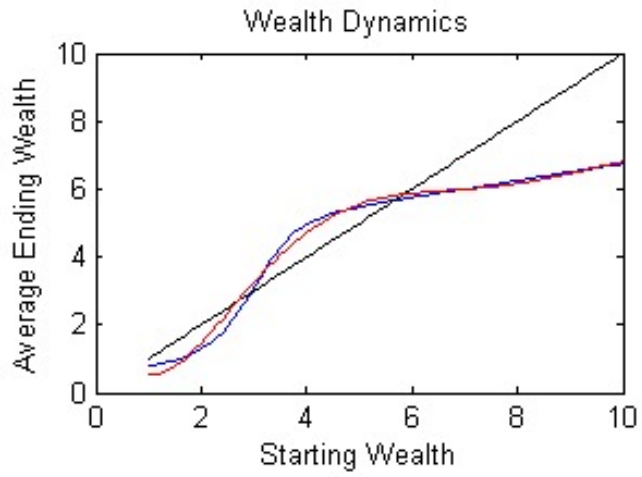


Figure 3 True Wealth Recursion Function (blue) with the Estimated 6th Order Polynomial Approximation (red)