Saving and Risk in Developing Countries:

Theory and Evidence from the Thai Project

Executive Summary

Robert M. Townsend[1]

MIT and NBER

We begin by describing the underlying economic environment of an emerging market country and test for the efficiency of financial markets and institutions using benchmarks derived from economic theory. The primary round and most easily implementable set of tests assumes as a null that households and businesses can achieve the standards of complete markets for the optimal allocation of risk and of perfect, costless intermediation of savings to borrowers. A second round of tests allows for various obstacles to trade. We deliberately do not do just one thing. Rather, we conduct multiple tests using different variables, data stratifications, and data frequencies to detect common patterns. We seek to determine where the theory fits well and where there are anomalies that suggest possible policy intervention. Observed, actual policy variation and instrumented, econometric variation in access are also informative and consistent.

The data used to test the theory come from a large, long term project of the author in Thailand. The project covers various regions of the country, including urban and rural areas, with multiple survey instruments and spans over 12 years of both annual and monthly data. The Townsend Thai surveys began in l997 with the deliberate selection of four provinces, two relatively near Bangkok (Lop Buri and Chachoengsao), in the industrialized central region, and two in the relatively poor Northeast (Buriram and Sisaket). Each of the provinces had at least one amphoe (county) that had been selected randomly under the National Statistics Office's Socio-Economic Survey. After that, all selection was random: 12 tambons (subcounties) in each province, four villages in each tambon, and 15 households in each village. The only stratification involved an environmental weighting to pick up tambons with higher than average forest cover.

[1] Research on the Thai Project is supported by grants from the NICHD, the NSF, the John Templeton Foundation and the Bill and Melinda Gates Foundation through the Consortium on Financial Systems and Poverty at the University of Chicago, where Townsend is Research Associate. The surveys are implemented in Thailand though the Thai Family Research Program, headed by Khun Sombat Sakuntasathien and partially funded by the University of Chicago UTCC center, in Bangkok.

The Townsend Thai project also includes surveys with village headmen, the Bank for Agriculture and Agricultural Cooperatives (BAAC) joint liability groups, village financial institutions and a module on the environment, including soil samples. In l998, given the onset of the Asian financial crisis, which began in Thailand, a subset of four tambons in each province were chosen at random to constitute an annual panel. That panel continues in the field 12 years later, and it currently has over 10 years of public data available for analysis covering 1,230 households. The annual panel was expanded in 2003, to two provinces in the south and in 2004 to two in the north, with one in each of these two additional regions, Satun and Prae, respectively, continuously sampled every year to date. The annual panel was again expanded in 2005 to urban areas to include mostly towns and some city neighborhoods.

Beyond the annual panel, the Townsend Thai Project includes a rich set of monthly household data. One tambon from each province in the 1997 baseline survey was set aside for a high frequency interviews, chosen with the criterion that the environment be similar across the four villages but the institutional mix be dissimilar. This was done in anticipation of future evaluations of the impact of social organization. The monthly surveys contain detailed questions on village networks, crop operations, financial transactions, labor supply, health and more.

The detailed monthly data have been used to create financial statements for each of the households: an income statement, balance sheet, and statement of cash flow. This nontrivial endeavor conceptualized the households as corporate firms and is described in detail in the Econometric Society Monograph*,* with Samphantharak (2009). These household accounts can be aggregated to create village level income and product accounts, saving and investment accounts, and balance of payments accounts, as if each village were a small open economy (Paweenawat and Townsend 2010).

Evident from these data, and from other surveys and measurements, risk is abundant in the Thai economic environment. There are macro shocks as with the financial crisis, but even these can be obscured with regional agro-climatic shocks. Surveyed households also mention price changes of inputs and outputs, and idiosyncratic events such as acute illness, fire or large ceremonial expenses (even when planned). Rainfall data from Thai meteorological stations indicate a large amount of variability, even within the semi-arid

Northeast. Some tambons have floods while others nearby can suffer drought, all in the same year. Nationally, monsoon patterns differ with adjacent Gulf of Thailand and Pacific currents. International prices, such as for rubber exports, also fluctuate, though with some persistence, and this has big consequences for villages in which many households specialize in para rubber cultivation.

Typically, however, not all members of a household, nor all households in a village, do the same thing. Though rice is the predominate crop in the Northeast, there are households that run businesses. Others earn spot wages from day labor or from salaried employment. Lop Buri has dairy cattle and considerable variety in cash crops while Chachoengsao has shrimp ponds, which support a national export to Wal-Mart and Costco, among others.

Paxon (1992), Townsend and Vickery (2004) and others test for how well consumption is smoothed against rain and price shocks, using the standards of the permanent income model and Thai Socio-economic Survey (SES) data. Oddly, rainfall is well smoothed but rubber prices are not. Most of the benchmark tests reported below adopt an even more demanding standard, which comes from the theory of the optimal allocation of risk: household consumption should move with aggregate shocks of the risk sharing unit (kinship group, village, region) and not with idiosyncratic, household-specific income. This can be derived from maximizing a weighted sum of household utilities subject to resource constraints. An extension of the theory allows production and investment, as in a neoclassical growth model, generating Euler equations, and under parametric assumptions, this delivers a regression frequently used in finance: investment should depend on common fixed time effects of the risk-sharing group and not on household-specific cash flow, controlling for productivity.

Townsend Thai annual data show that the coefficient on idiosyncratic income, though not zero, is of the order of magnitude of the local interest rate, which is consistent with the permanent income model (Alem 2010). However, the relatively poor at the bottom 30% of wealth distribution have much higher coefficients, on both consumption and investment. The Townsend Thai monthly data show a remarkably similar pattern. In fact the benchmark is not rejected for households with kinship groups in the village (Chiappori et al 2010). For those without other family-related households around in the village, coefficients are positive and

significant, yet low, as reported in the annual data. Likewise, there is an exception for those in the lower quartile of the average income distribution, who display much higher coefficients. We conclude from these consumption and income data that the relatively poor and especially those without kin in the same village are vulnerable to idiosyncratic shocks. A buffer stock model with lumpy investment can generate such sensitivity to shocks, as in Kaboski and Townsend (2010).

Again, for those in family groups, the benchmark is not rejected, so one can go further and allow for heterogeneous risk preferences. Specifically, one can back out risk aversion by examining how much a household's consumption varies with village level shocks, i.e., their willingness to bear aggregate risk. These shocks vary across villages. Independently, one can do this from Euler equations, Mehra-Prescott (1985) style calculations using household specific consumption data and the return on household assets. See Chiappori et al (2010). Tests confirm heterogeneity and, further, a positive correlation in these two measures, so heterogeneity in risk preferences seems a key part of reality. We can then ask what would be the welfare implications of insuring each village against aggregate risk. The typical household has welfare losses from fluctuations at one half to one percent of consumption, and more risk adverse households up to four percent of consumption. However, the more risk tolerant households are hurt, ironically, by this policy intervention; they had been increasing their average consumption by the sale of insurance to their more risk averse village neighbors. This is a warning that not all seemingly needed interventions are universally welfare-improving.

Ongoing work with Bonhomme et al (2010) modifies the theory to allow for endogenous labor supply, which includes periods of zero participation in the labor market (adding wage labor and non-wage labor to household enterprise). In the data, of course, labor supply and participation both move with the wage rate, as they should. However, controlling for that, a one standard deviation increase in profits -- that is, in non-wage income -- yields small effects overall. Therefore, idiosyncratic income variation does not matter much at all for hours or participation. Again, however, a salient and now familiar exception emerges: the relatively poor are vulnerable, working more and working harder when household-specific incomes are low.

Returning to the Euler equations, one can also test for efficiency using the standards of the capital asset pricing model. Essentially an activity that has a high yield, measured as return on assets (ROA), when

the village aggregate yield is low, is an asset that is quite valuable. Conversely, a household running a portfolio of activities that commoves with the village aggregate has to yield a higher expected return to make the portfolio worthwhile to hold. Plotting these household specific "betas" against the time average of household returns reveals a remarkably good fit for households in kinship networks and households in the villages more generally. However, by various criteria, the fit deteriorates as one moves to the regional and national levels, pooling across spatially-separated units. The degree of risk aversion which rationalizes the equity premium is relatively low village by village, but it rises to higher numbers at the regional and especially national levels. Therefore, the null of perfect markets is more likely to be rejected as one moves to more inclusive geographic aggregates. An auxiliary conclusion is that production and return data are informative of the effective "financial regime."

If full risk sharing is rejected for some stratifications of the data, by wealth or geography, then what is the alternative model? Work with Alex Karaivanov (2010) uses both the consumption and investment data to see which financial regime is the closest approximation to the actual data generating process. Essentially we use histograms, that is, frequency distributions of joint consumption, income, and investment variables, both in the actual Thai data and as generated from models assuming various possible imperfections: moral hazard, unobserved investment, simple credit markets with a natural borrowing limit, and savings only (buffer stock). The moral hazard model does reasonably well in the Northeast, especially when using only consumption and income data, and overall in some of the cross sectional consumption runs when restricting attention to those in family networks, consistent with earlier results, reported above. Overall though, using all data jointly, the borrowing-lending model fits best, often tied with savings only. It seems these divergent results, from different variables, are driven by the struggle to fit the investment/cash flow data rather than consumption/income data.

Other work also rejects the null of perfect risk sharing but differs on the effective financial regime. Paulson, Townsend, and Karaivanov (2006) using wealth and occupation transitions finds a moral hazard model fits best in the Central region, with limited commitment (wealth constrained borrowing) an additional element in the Northeast. Ahlin and Townsend (2007a) come to a similar conclusion using data on loan

repayments; social penalties vary for default across Northeastern villages and has the impact theory would predict. Ahlin and Townsend (2007b) find evidence of adverse selection in the choice BAAC borrowers make about joint or individual liability loans.

These tests focus on outcomes, but there is abundant data on mechanisms. The creation of the financial statements for households in the monthly Thai data allows us to decompose budget deficits -- the difference between consumption plus investment less income -- into the various financing devices, to see what households actually do in bad times. We discover that increases and decreases in cash holdings are the predominate device for most households in all regions. Of course, this would be consistent with a simple buffer stock model. Borrowing and especially gifts also play a role, with the latter particularly large in the Northeast. These would allow more smoothing and insurance. Surprisingly, use of savings in formal sector banks is small in the overall monthly household sample. An ongoing project (Alvarez et al 2010) studies these data to see if the large level of cash holdings and infrequent use of formal sector savings accounts can be rationalized by a Miller Or model, working on some of Alvarez and Lippi (2010) more recent contributions.

Data summaries give us the levels of borrowing and saving by financial institution, and also show us changes in balances, adjusted for interest payments, as a documentation of active use. We find for loans that, on average, 40 percent of the population use the BAAC, and a higher 60 percent use village funds, though this is beginning to decline over time. The informal sector, including family, accounts for approximately 40 percent of loan volume, though that may be declining over time as well. Likewise, although up to 80 percent of the population has savings accounts, including mandatory village fund accounts, this percentage drops to 50 percent with a more stringent, active-use criterion. More generally, one can run probits to summarize who is a client and who is not, who is using which financial instrument, and so on. For example, the BAAC targets the middle wealth segment of the market. Commercial banks focus on the high end, and the informal sector is on the lower end, although in all cases there is some overlap in use by wealth distributions.

We take analysis beyond these data summaries by returning to the benchmark risk sharing regressions. Alem and Townsend (2010) look at individual service providers to see how well they are doing in allowing households and businesses to achieve the benchmark standards. As is already evident, there is a

great deal selection across providers in the data, so we need plausible exogenous variation that determines use or participation but does not influence error terms in outcomes or conditional use. Using distance to the branch of the financial institution in question, or past history of providers in the area, or even surprise variation that breaks the local geographic patterns, they find that the BAAC allows clients to come close to or, in some cases, completely achieve the benchmark consumption and investment standards.

This is not surprising given the BAAC operating system. As documented in Townsend and Yaron (2001), otherwise standard looking loans are really in fact mixed with options that allow deferred payment, or even forgiveness of some of the principal, in some cases. Here then, we begin to couple the analysis of impact with an understanding of business models, on the supply side.

Commercial banks are also helpful, presumably operating though savings accounts, as there are relatively few commercial bank loans in the data. Other institutions such as standalone agricultural cooperatives or village level rice banks do much less well by this scorecard rating system. Ironically, we judge an institution not by its capital asset ratio or firewalls that "protect" institutions and the economy from contagion, but rather by how well the institution allows households and businesses to achieve an optimal allocation of risk bearing, treating customers and clients as the risk sharing group and acknowledging that someone must bear macro shocks.

Returning to kinship and other networks, we can now judge whether indirect links to the BAAC or to commercial banks help otherwise disconnected households, those without direct access to financial institutions. As might have been anticipated from the earlier discussion, networks formed by observed transactions in gifts and loans do allow "indirect customers" to do as well in smoothing consumption. Having family in the village is also effective, but interestingly, the effect may run more through penalties for default and off-equilibrium behavior than through observed transactions, especially for commercial banks. That is, having family around helps in smoothing investment from cash flow fluctuations (more so for the rich than the poor (Samphantarak & Townsend 2009). But the mechanism is not what one might have anticipated (i.e., it is not coming from observed financial transactions).

It bears repeating that one cannot look at each financial institution in isolation. Kinship groups can intermediate locally, as was just noted. A household can also use one institution to pay off the loan of another, in effect providing a bridge loan for refinancing. (Sripakdeevong & Townsend 2010) Also, in some instances, institutions can be substitutes for each other. Savings mobilization may be less effective than it otherwise appears because the deposit and withdrawals to and from formal sector savings accounts compete with lending and borrowing in informal, village-level money markets.

Yet local systems can sometimes be harnessed to outside interventions, enabling something which would not happen otherwise. In 2002, then Prime Minister Thaskin seeded every village in Thailand with one million baht ($24,000) to form a local savings and loan association, run by a village committee. With variation in the number of households in a village, something documented to be not related to almost anything else, this created a natural policy instruments, that is, variation in per capita treatment.

An astructural reduced form paper with Kaboski (2009) establishes that the million baht fund program increased consumption, profits from businesses, labor income, agricultural investment, and total borrowing above and beyond village fund credit, while also raising default rates and lowering assets and savings. Kaboski and Townsend (2010) describe a mechanism consistent with these impacts. Savings falls as credit availability reduces the need for buffers. Difficulty in repayment happens when a household used the money for consumption or invested it in a longer term asset, which (only slowly) yields dividends, and suffers from an unusually bad year.

At estimated parameters, the typical, median households could have been made better off relative to the million baht program with lump sum transfers. Defaulting household are actually hurt by the program, as a weakened credit constraint means they continue to carry debt at interest. On the other hand, households that suffer from constraints, i.e., borrowing up to a credit limit or holding back on investment, prefer the enhanced intermediation over transfers. The structural model also allows counterfactual experiments. A restriction that a borrowing household must also be investing at the same time raises average welfare.

The intermediation that came somehow with the village fund is captured only crudely in the model by a weakened credit constraint. De la Huerta (2010) establishes that there is much variation in factors across

urban neighborhood and rural village funds correlated with success or failure in this enhanced intermediation. Conversely, ongoing work with Breza and Banerjee (2010) hints that, on average, village funds were used efficiently, in that there was bigger impact on those households with higher total factor productivity (TFP), oddly though, on consumption relatively quickly and on investment only in the longer term.

These findings are consistent with earlier work with the monthly data. Pawasutipaisit & Townsend (2010) establishes that the marginal product of capital varies across household and is particularly high for those with few productive assets. This is not just a matter of dividing by a small number. Those same households, with high TFP, are observed to put earned profits back into their own enterprise. Likewise, savings rates for them tend to be higher and consistent with theoretical models in the literature (Buera 2003, Buera & Moll 2010). Yet, at any moment in time, there are high cross-sectional disparities in productivity which are not overcome by the slightly higher debt to asset ratios for these more productive households.

To sum up, poor households without kinships groups in their village appear quite vulnerable in consumption and investment to fluctuations in idiosyncratic income. This risk sharing seems to fall off with geographic distance, however. Moreover, the Thai financial system is not allowing for sufficient intermediation of savings to borrowers with productive technologies. This and welfare improving interventions are under consideration in next steps.

Ahlin, Christian and Robert M. Townsend. (2007a) "Using Repayment Data to Test Across Models of Joint Liability Lending," Economic Journal 117(517), March 2007: F11-F51.

Ahlin, Christian and Robert M. Townsend. (2007b) "Selection Into and Across Credit Contracts: Theory and Field Research," Journal of Econometrics 136(2), February 2007: 665-98.

Alem, Mauro and Robert M. Townsend. (2010) "An Evaluation of Financial Institutions: Impact on Consumption and Investment Using Panel Data and the Theory of Risk-Bearing" Working Paper.

Alvarez, Fernando and Francesco Lippi . (2010) "Financial innovation and the transactions demand for cash." Working Paper.

Alvarez, Fernando, Anan Pawasutipaisit, and Robert M. Townsend. (2010) "Household as Firms: Cash Management in Thai Villages." Working paper, University of Chicago.

Breza, Banerjee, and Townsend (2010) manuscript, MIT.

Buera, Francisco J. (2003) "A Dynamic Model of Entrepreneurial Choice with Borrowing Constraints". University of Chicago.

Buera, Francisco J. and Benjamin Moll. (2010) "Dynastic Capital Misallocation." Working Paper.

Chiappori, Pierre Andre, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M. Townsend (2010) "Heterogeneity and Risk Sharing in Thai Villages". Working Paper.

Kaboski, Joseph and Robert M. Townsend (2010). "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative," Econometrica, forthcoming.

Kaboski, Joseph and Robert M. Townsend (2009) "The Impacts of Credit on Village Economies." Working Paper.

Karaivanov, Alexander and Robert Townsend. (2010) "Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes." Working Paper.

Anna Paulson, Robert Townsend and Alex Karaivanov. (2006) "Distinguishing Limited Liability from Moral Hazard in a Model of Entrepreneurship." Journal of Political Economy 114(1): 100-144.
Mehra, Rajnish, and Edward C. Prescott. "The Equity Premium: A Puzzle." Journal of Monetary Economics. 15(2), March 1985: 145-61.

Pawasutipaisit, Anan and Robert Townsend. (2010) "Wealth Accumulation and Factors Accounting for Success." Journal of Econometrics, Forthcoming.

Paweenawat, Archawa and Robert M. Townsend. (2010) "Villages as Small Open Economies." Working Paper.

Paxson, C. H. (1992) 'UsingWeather Variability to Estimate the Response of Savings to Transitory Income in Thailand', American Economic Review, 82: 15–33.

Samphantharak, Krislert and R. M. Townsend (2009) Households as Corporate Firms: An Analysis of Household Finance Using Integrated Household Surveys and Corporate Financial Accounting. Econometric Society Monographs. Cambridge: Cambridge University Press

Sripakdeevong, Parit and Robert M. Townsend. (2010) manuscript, MIT.

Townsend, Robert M. and Jacob Yaron. (2001) 'The Credit Risk Contingency System of an Asian Development Bank', Federal Reserve Bank of Chicago Economic Perspectives, 25: 31–48.