Published in final edited form as:

Econometrica. 2011 September ; 79(5): 1357–1406. doi:10.3982/ECTA7079.

A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative

Joseph P. Kaboski and Robert M. Townsend^{†,*}

Abstract

This paper uses a structural model to understand, predict, and evaluate the impact of an exogenous microcredit intervention program, the Thai Million Baht Village Fund program. We model household decisions in the face of borrowing constraints, income uncertainty, and high-yield indivisible investment opportunities. After estimation of parameters using pre-program data, we evaluate the model's ability to predict and interpret the impact of the village fund intervention. Simulations from the model mirror the data in yielding a greater increase in consumption than credit, which is interpreted as evidence of credit constraints. A cost-benefit analysis using the model indicates that some households value the program much more than its per household cost, but overall the program costs 20 percent more than the sum of these benefits.

1 Introduction

This paper uses a structural model to understand, predict, and evaluate the impact of an exogenous microcredit intervention program, the Thai Million Baht Village Fund program. Understanding and evaluating microfinance interventions, especially such a large scale government program, is a matter of great importance. Proponents argue that microfinance allows the provision of credit that is both effective in fighting poverty and more financially viable than other means; detractors point to high default rates, reliance on (implicit and explicit) subsidies, and the lack of hard evidence of their impacts on households. The few efforts to evaluate the impacts of microfinance institutions using reduced form methods and plausibly exogenous data have produced mixed and even contradictory results.¹ To our knowledge, this is the first structural attempt to model and evaluate the impact of microfinance. Three key advantages of the structural approach are the potential for quantitative interpretation of the data, counterfactual policy/out of sample prediction, and well-defined normative program evaluation.

The Thai Million Baht Village fund program is one of the largest scale government microfinance initiatives of its kind.² Started in 2001, the program involved the transfer of one million baht to each of the nearly 80,000 villages in Thailand to start village banks. The transfers themselves sum to about 1.5 percent of Thai GDP and substantially increased available credit. We study a panel of 960 households from sixty-four rural Thai villages in

[†]Kaboski <jkaboski@nd.edu>, Townsend <rtownsen@mit.edu>.

^{*}Research funded by NICHD grant R03 HD04776801, the Bill & Melinda Gates Foundation grant to the U. of Chicago Consortium on Financial Systems and Poverty, John Templeton Foundation, and NSF. We thank Sombat Sakuntasathien, Aleena Adam, Francisco Buera, Flavio Cunha, Xavier Gine, Donghoon Lee, Audrey Light, Ben Moll, Masao Ogaki, Anan Pawasutpaisit, Mark Rosenzweig, Shing-Yi Wang, Bruce Weinberg and participants at FRB-Chicago, FRB-Minneapolis, Harvard-MIT, Michigan, NIH, Ohio State, UW Milwaukee, NYU, Yale, NEUDC 2006, 2006 Econometric Society, BREAD 2008, and World Bank Microeconomics of Growth 2008, UC-UTCC, NYU Development Conference, and SED 2009 presentations. Bin Yu, Taehyun Ahn, and Jungick Lee provided excellent research assistance on this project.

¹Pitt and Khandker (1998), Pitt et al (2003), Morduch (1998), Coleman (1999), Gertler, Levine and Moretti (2003), Karlan and Zinman (2006), and Banerjee et al (2009) are examples. Kaboski and Townsend (2005) estimates positive impacts of microfinance in Thailand using non-experimental data.

the Townsend Thai Survey (Townsend et al, 1997). In these villages, funds were founded between the 2001 and 2002 survey years, and village fund loans amounted to eighty percent of new short-term loans and one third of total short-term credit in the 2002 data. If we count village funds as part of the formal sector, participation in the formal credit sector jumps from 60 to 80 percent.

Though not a randomized treatment, the program is viewed as a quasi-experiment that produced plausibly exogenous variation in credit over time and across villages. The program was unanticipated and rapidly introduced. More importantly, the total amount of funding given to each village was the same (one million baht) regardless of the number of households in the village. Although village size shows considerable variation within the rural regions we study, villages are administrative geopolitical units and are often subdivided or joined for administrative or political purposes. Indeed, using GIS maps, we have verified that village size patterns are not much related to underlying geographic features and vary from year to year in biannual data. Hence, there are *a priori* grounds for believing that this variation and the magnitude of the per capita intervention is exogenous with respect to the relevant variables. Finally, village size is not significantly related to pre-existing differences (in levels or trends) in credit market or relevant outcome variables.

Our companion paper, Kaboski and Townsend (2008), examines impacts of the program using a reduced form regression approach and many of the impacts are puzzling without an explicit theory of credit-constrained behavior.³ In particular, households increased their borrowing and their consumption roughly one for one with each dollar put into the funds. A perfect credit model, such as a permanent income model, would have trouble explaining the large increase in borrowing, since reported interest rates on borrowing did not fall as a result of the program. Similarly, even if households treated loans as a shock to income rather than a loan, they would only consume the interest of the shock (roughly seven percent) perpetually. Moreover, households were not initially more likely in default after the program was introduced, despite the increase in borrowing. Finally, household investment is an important aspect of household behavior. We observe an increase in the frequency of investment, but, oddly, impacts of the program on the *level* of investment were difficult to discern. This is *a priori* puzzling in a model with divisible investment, if credit constraints are deemed to play an important role.

The structural model we develop in this paper here sheds light on many of these findings. Given the prevalence of income shocks that are not fully insured in these villages (see Chiappori et al. (2008)), we start with a standard precautionary savings model (e.g., Aiyagari (1994), Carroll (1997), Deaton (1991)). We then add important features central to the evaluation of microfinance but also key characteristics of the pre-program data: borrowing, default, investment, and growth. Short-term borrowing exists but is limited, and so we naturally allow borrowing but only up to limits. Similarly, default exists in equilibrium, as does renegotiation of payment terms, and so our model incorporates default. Investment is relatively infrequent in the data but is sizable when it occurs. To capture this lumpiness, we allow households to make investments in indivisible, illiquid, high yield projects whose size follows an unobserved stochastic process.⁴ Finally, income growth is high but variable, averaging 7 percent but varying greatly over households, even after

²The Thai program involves approximately \$1.8 billion in initial funds. This injection of credit into the rural sector is much smaller than Brazilian experience in the 1970s, which saw a growth in credit from about \$2 billion in 1970 to \$20.5 billion in 1979. However, in terms of a government program implemented through village institutions and using micro-lending techniques, the only comparable government program in terms of scale would be Indonesia's KUPEDS village bank program, which was started in 1984 at a cost of \$20 million and supplemented by an additional \$107 million in 1987. (World Bank, 1996)

³This companion paper also provides additional evidence on the exogeneity of village size, examines impacts in greater detail, and looks for general equilibrium effects on wages and interest rates.

controlling for life cycle trends. Allowing for growth requires writing a model that is homogeneous in the permanent component of income, so that a suitably normalized version attains a steady state solution, giving us time-invariant value functions and (normalized) policy functions.

In an attempt to *quantitatively* match central features of the environment, we estimate the model using a Method of Simulated Moments (MSM) on only the *pre-program* data. The parsimonious model broadly reproduces many important aspects of the data, closely matching consumption and investment levels, and investment and default probabilities. Nonetheless, two features of the model are less successful, and the overidentifying restrictions of the model are rejected.⁵

For our purposes, however, a more relevant test of the estimated model's usefulness is its ability to predict out-of-sample responses to an increase in available credit, namely the village fund intervention. Methodologically, we model the microfinance intervention as an introduction of a borrowing/lending technology that relaxes household borrowing limits. These limits are relaxed differentially across villages in order to induce an additional one million baht of short-term credit in each village; hence, small villages get larger reductions of their borrowing constraint.

Given the relaxed borrowing limits, we then simulate the model with the stochastic income process to create 500 artificial datasets of the same size as the actual Thai panel. These simulated data do remarkably well in reproducing the above impact estimates. In particular, they predict an average response in consumption that is close to the dollar-to-dollar response in the data. Similarly, the model reproduces the fact that effects on average investment levels and investment probabilities are difficult to measure in the data.

In the simulated data, however, these aggregate effects mask considerable heterogeneity across households, much of which we treat as unobservable to us as econometricians. Increases in consumption come from roughly two groups. First, hand-to-mouth consumers are constrained in their consumption either because they have low current liquidity (income plus savings) or are using current (pre-program) liquidity to finance lumpy investments. These constrained households use additional availability of credit to finance current consumption. Second, households who are not constrained may increase their consumption even without borrowing, since the increase in available credit in the future lowers their desired bufferstock savings. Third, for some households, increased credit induces them to invest in their high yield projects. Some of these households may actually reduce their consumption, however, as they supplement credit with reduced consumption in order to finance sizable indivisible projects. (Again, the evidence we present for such behavior in the pre-intervention data is an important motivation for modeling investment indivisibility.) Finally, for households who would have defaulted without the program, available credit may simply be used to repay existing loans and so have little effect on consumption or investment. Perhaps most surprising is that these different types of households may all appear *ex ante* identical in terms of their observables.

The estimated model not only highlights this underlying heterogeneity, but also shows the quantitative importance of these behaviors. Namely, the large increase in consumption

⁴An important literature in development has examined the interaction between financial constraints and indivisible investments. See, for example, Banerjee and Newman (1993), Galor and Zeira (1993), Gine and Townsend (2004), Lloyd-Ellis and Bernhardt (2001), and Owen and Weil (1997).

⁵The income process of the model has trouble replicating the variance in the data, which is affected by the Thai financial crisis in the middle of our pre-intervention data, and the borrowing and lending rates differ in the data but are assumed equal in the model. Using the model to match year-to-year fluctuations is also difficult.

indicates the relative importance of the first two types of households, both of whom increase their consumption. Also, the estimated structural parameters capture the relatively low investment rates and large skew in investment sizes. Hence, overall investment relationships are driven by a relatively few, large investments, and so very large samples are needed to accurately measure effects on average investment. The model generates these effects but for data that are larger than the actual Thai sample. Second and related, given the lumpiness of projects, small amounts of credit are relatively unlikely to change investment decisions on the large projects that drive aggregate investment.

Finally, our normative evaluation compares the costs of the Million Baht program to the costs of a direct transfer program that is equivalent in the sense of providing the same utility benefit. The heterogeneity of households plays an important role, and indeed the welfare benefits of the program vary substantially across households and villages. Essentially, there are two major differences between the microfinance program and a well-directed transfer program. First, the microfinance program is potentially *less* beneficial because households face the interest costs of credit. In order to access liquidity, households borrow more, and while they can always carry forward more debt into the future, they are left with larger interest payments. Interest costs are particularly high for otherwise defaulting households, whose debts is augmented to the more liberal borrowing limit, and so they bear higher interest charges. On the other hand, the microfinance program is potentially *more* beneficial than a direct transfer program because it can also provide more liquidity to those who potentially have the highest marginal valuation of liquidity by lowering the borrowing constraint. Hence, the program is relatively more costeffective for non-defaulting households with urgent liquidity needs for consumption and investment. Quantitatively, given the high frequency of default in the data⁶ and the high interest rate, the benefits (i.e., the equivalent transfer) of the program are twenty percent less than the program costs, but this masks the interesting variation among losers and gainers.

Beyond the out-of-sample and normative analyses, we also perform several alternative exercises that build on the strengths of the structural model: long run out-of-sample predictions showing the time-varying impacts; a counterfactual “investment contingent credit” policy simulation that underperforms the actual policy; and re-estimation using the pooled sample, which confirmed the robustness of our exercise.

The paper contributes to several literatures. First, we add a structural modeling approach to a small literature that uses theory to test the importance of credit constraints in developing countries (e.g., Banerjee and Duflo (2002)). Second, we contribute to an active literature on consumption and liquidity constraints, and the bufferstock model, in particular. Studies with U.S. data have also found a high sensitivity of consumption to current available liquidity (e.g., Zeldes (1989), Souleles and Gross (2002), Aaronson, Agarwal, and French (2008)), but like Burgess and Pande (2005), we study this response with quasi-experimental data in a developing country.⁷ Their study used a relaxation of branching requirements in India that allowed for differential bank expansion across regions of India over twenty years in order to assess impacts on poverty headcount and wage data. Third, methodologically, our quasi-experimental analysis builds on an existing literature that has used out-of-sample prediction, and experiments in particular, to evaluate structural models (e.g., Lise et al., (2005a, 2005b), Todd and Wolpin (2006)). Finally, we contribute to the literature on measuring and interpreting treatment effects (e.g., Heckman, Urzua, and Vytlacil (2004)), which has

⁶Default rates on short-term credit overall were 19 percent of households, but less than 3 percent of village fund credit was in default, and one-fourth to one-third of households reported that they borrowed from other sources to repay the loans.

⁷Banerjee et al (2009) find large impacts on durable expenditures using a randomized microfinance experiment in Hyderabad, India.

emphasized unobserved heterogeneity, non-linearity and time-varying impacts. We develop an explicit behavioral model where all three play a role.

The remainder of the paper is organized as follows. The next section discusses the underlying economic environment, the Million Baht village fund intervention, and reviews the facts from reduced form impact regressions that motivate the model. The model, and resulting value and policy functions, are presented in Section 3. Section 4 discusses the data and presents the MSM estimation procedure and resulting estimates. Section 5 simulates the Million Baht intervention, performs policy counterfactuals, and presents the welfare analysis. Section 6 concludes.

2 Thai Million Baht Credit Intervention

The intervention that we consider is the founding of village-level microcredit institutions by the Thai government, the Million Baht Fund program. Former Thai Prime Minister Thaksin Shinawatra implemented the program in Thailand in 2001, shortly after winning election. One million baht (about \$24,000) was distributed to each of the 77,000 villages in Thailand to found self-sustaining village microfinance banks. Every village, whether poor or wealthy, urban⁸ or rural was eligible to receive the funds. The size of the transfers alone, about \$1.8 billion, amounts to about 1.5 percent of GDP in 2001. The program was overwhelmingly a credit intervention; no training or other social services were tied to the program, and although the program did increase the fraction of households with formal savings accounts, savings constituted a small fraction (averaging 14,000 baht or less than two percent) of available funds, and we measured no effect on the actual levels of formal savings during the years we study.

The design of the program was peculiar in that the money was a grant program to village funds (because no repayment was expected or made), yet the money reaches borrowers as microcredit loans with an obligation to repay to the fund. As noted earlier default rates to these funds themselves were low (less than 3 percent up through available 2005 data), and all village funds in the sample we use continue in operation, indicating that the borrowers obligation to repay was well understood in the rural villages we study. (In contrast, default rates to village funds in urban areas are substantially higher, roughly 15 percent.) Also, the quasi-experiment is quite different and less clean than typical randomizations, since the villagers themselves get to organize the funds, and in randomizations there is typically much greater control over what happens. Thus, one must be careful not to extrapolate our results across all environments and microfinance interventions. We are not evaluating a microfinance product via randomized trials.

The design and organization of the funds were intended to allow all existing villagers equal access to these loans through competitive application and loan evaluation handled at the village level. Villages elected committees who then drew up the rules for operation. These rules needed to satisfy government standards, however, and the village fund committees were relatively large (consisting of 9–15 members) and representative (e.g., half women, no more than one member per household) with short, two year terms. In order to obtain funds from the government, the committees wrote proposals to the government administrators outlining the proposed policies for the fund.⁹ For these rural villages, funds were disbursed

⁸The village (moo ban) is an official political unit in Thailand, the smallest such unit, and is under the sub-district (tambon), district (amphoe), and province (changwat) levels. Thus, “villages” can be thought of as just small communities of households that exist in both urban and rural areas.

⁹These policies varied somewhat, but were not related to village size. For example, some funds required membership fees but all were under 100 baht (\$2.50), interest rates averaged 7 percent, but the standard deviation was 2 percent, the number of required guarantors varied with an average of 2.6 and a standard deviation of one.

to and held at the Thai Bank of Agriculture and Agricultural Cooperatives, and funds could only be withdrawn with a withdrawal slip from the village fund committee. Residence in the village was the only official eligibility requirement for membership, and so although migrating villagers or newcomers would likely not receive loans, there was no official targeting of any sub-population within villages. Loans were uncollateralized, though most funds required guarantors. Repayment rates were quite high; less than three percent of funds lent to households in the first year of the program were 90 days behind by the end of the second year. Indeed, based on the household level data, ten percent more credit was given out in the second year than in the first, presumably partially reflecting repaid interest plus principal. There were no firm rules regarding the use of funds, but reasons for borrowing, ability to repay, and the need for funds were the three most common loan criteria used. Indeed many households were openly granted loans for consumption. The funds make short-term loans – the vast majority of lending is annual – with an average nominal interest rate of seven percent. This was about a five percent real interest rate in 2001, and about five percent above the average money market rate in Bangkok.¹⁰

2.1 Quasi-Experimental Elements of the Program

As described in the introduction, the program design was beneficial for research in two ways. First, it arose from a quick election, after the Thai parliament was dissolved in November, 2000, and was rapidly implemented in 2001. None of the funds had been founded by our 2001 (May) survey date, but by our 2002 survey, each of our 64 villages had received and lent funds, lending 950,000 baht on average.¹¹ Households would not have anticipated the program in earlier years. We therefore model the program as a surprise. Second, the same amount was given to each village, regardless of the size, so villages with fewer households received more funding per household. Regressions below report a highly significant relationship between household's credit from a village fund and inverse village size in 2002 after the program.

Our policy intervention is not a clean randomized experiment, and so we cannot have the same level of certainty about the exogeneity of the program. Several potential problems could contaminate the results. First, variables of interest for households in small villages could differ from those in large villages even before the program. Second, different trends in these variables across small and large villages would also be problematic, since the program occurs in the last years of the sample. If large villages had faster growth rates, we would see level differences at the end of the period and attribute these to the intervention during those years. Third, other policies or economic conditions during the same years could have affected households in small and large villages differentially.¹²

Other issues and caveats arise from all of our variation coming at the village level. On the one hand, village-level variation has important benefits because, in many ways, each village is viewed as its own small economy. These village economies are open but not entirely integrated with one another and the rest of the broader economy (nearby provinces, regions, etc.) in terms of their labor, credit, and risk-sharing markets and institutions. This gives us confidence that program impacts are concentrated at the village level.¹³ On the other hand, one could certainly envision potential risks involved with our use of village size. For example, even if credit itself were exogenous, its impact could differ in small and large

¹⁰More details of the funds and program are presented in Kaboski and Townsend (2009).

¹¹We know the precise month that the funds were received, which varies across villages. This month was uncorrelated with the amount of credit disbursed, but may be an additional source of error in predicting the impacts of credit.

¹²Other major policies initiated by the Thaksin government included the "30 Baht Health Plan" (which set a price control at 30 baht per medical visit), and "One Tambon-One Product" (a marketing policy for local products). However, neither were operated at the village level, since the former is an individual level program while the latter is at the tambon (sub-district) level.

¹³GIS analysis including neighboring villages in Kaboski and Townsend (2009) support this claim.

villages. Small villages might be more closely connected, with better information or less corruption, and so might show larger impacts not only because they received more credit per household but because the credit was used more efficiently. Conversely, small villages might have smaller markets and so credit might have smaller impacts. Keeping this caveat in mind, our approach is to take a stand on a plausible structural model in Section 3. Within this structural model, village size will be fully excluded from all equations. So that when we introduce the policy in Section 4, the only role of village size will be in determining the expansion of credit. We are encouraged that the simple model does well in replicating the out-of-sample patterns in the data.

Despite the potential risks and caveats, there are both a priori and a posteriori reasons for pursuing our exclusion restriction and accepting inverse village size as exogenous with respect to important variables of interest.

First, villages are geopolitical units, and villages are divided and redistricted for administrative purposes. These decisions are fairly arbitrary and unpredictable, since the decision processes are driven by conflicting goals of multiple government agencies. (See, for example, Pugenier (2002) and Arghiros (2001)), and splitting of villages is not uncommon. Data for the relevant period (1997–2003 or even the years directly preceding this, which might perhaps be more relevant) are unavailable, but growth data is available for 1960–2007 and for 2002–2007, so we know that the number of villages grew on average by almost one percent a year both between 1960 and 2007 and during the more recent period. Clearly, overall trends in new village creation are driven in part by population growth, but the above literature indicates that the patterns of this creation are somewhat arbitrary.

Second, because inverse village size is the variable of interest, the most important variation comes from a comparison among small villages (e.g., between 50 and 250 households). Indeed, the companion paper focuses its baseline estimates on these villages, but show that results are robust to including the whole sample. That is, the analysis is not based on comparing urban areas with rural areas, and we are not picking up the effects of other policies biased toward rural areas and against Bangkok.

Third, village size is neither spatially autocorrelated, nor correlated with underlying geographic features like roads or rivers, which might arise if village size were larger near population centers or fertile areas. Using data from Community Development Department (CDD), Figure 1 shows the random geographical distribution of villages by decile of village size in the year 2001 over the four provinces for which we have Townsend Thai data (Chachoengsao, Lopburi, Buriram and Sisaket). The Moran spatial autocorrelation statistics in these provinces are 0.019 (standard error of 0.013), 0.001 (0.014), 0.002 (0.003), and 0.016 (0.003), respectively.¹⁴ Only the Sisaket autocorrelation is statistically significant, and the magnitudes of all of them are quite small. For comparison, the spatial autocorrelation of the daily wage in villages ranges from 0.12 to 0.21. We also checked whether village size was correlated to other underlying geographic features by running separate regressions of village size onto distance to nearest two-lane road or river

¹⁴The general formula for Moran's statistic is:

$$I = \frac{n}{\sum_{i=1}^n \sum_{j=1}^n w_{ij}} \left(\frac{\sum_{i=1}^n \sum_{j=1}^n w_{ij} (z_i - \bar{z})(z_j - \bar{z})}{\sum_{i=1}^n \sum_{j=1}^n (z_i - \bar{z})^2} \right)$$

where n is the number of observations (villages), z_i is the statistic for observation i (village size of village i), and w_{ij} is the weight given villages depending on their spatial distance. Here we use inverse cartesian distance between villages.

(conditioning on changwat dummies). The estimated coefficients were 0.26 (standard error of 0.32) and -0.25 (0.24), so neither was statistically significant. Small villages did tend to be located closer to forest areas however, where the coefficient of 0.35 (0.03) was highly significant, indicating that forest area may limit the size of villages.¹⁵ Nonetheless, these regressions explain at most five percent of the variation in village size, so the variation is not particularly well explained by geographic features. We have included roads, rivers, and forest in Figure 1.

Finally, the regression analysis in our companion paper, Kaboski and Townsend (2009), strengthens our a posteriori confidence in the exogeneity of village size. Specifically, we present reduced form regressions on a large set of potential outcome variables. Using seven years of data (1997–2003, so that $t=6$ is the first post-program year, 2002), we run a first-stage regression to predict village fund credit of household n in year t , $VFCR_{n,t}$:

$$VFCR_{n,t} = \sum_{j=6,7} \alpha_{1,VFCR,j} \frac{1,000,000}{\# \text{ HHs in village}_{v,6}} I_{t=j} + \alpha_{2,VFCR} \text{Control}_{n,t} + \gamma_{VFCR} \mathbf{X}_{nt} + \theta_{VFCR,t} + \theta_{VFCR,n} + \varepsilon_{VFCR,nt}$$

and second-stage outcome equation of the form:

$$Z_{nt} = \alpha_{1,Z} VFCR_{n,t} + \alpha_{2,Z} \text{Control}_{n,t} + \gamma_Z \mathbf{X}_{nt} + \theta_{Z,t} + \theta_{Z,n} + \varepsilon_{Z,nt} \quad (1)$$

where Z_{nt} represents an outcome variable of interest for household n in year t . Comparing the two equations, the crucial variable in the first stage is inverse village size in the post-intervention years (the latter captured by the indicator function $I_{t=j}$), since it creates variation in $VFCR_{n,t}$, but is excluded from the second stage outcome equation. Although there is heterogeneity across households and non-linearity in the impact of credit, $\hat{\alpha}_{1,Z}$ captures (a linear approximation of) the relationship between the average impact of a dollar of credit on the outcome of Z_{nt} .

The sets of controls in the above equations are $\mathbf{X}_{nt}$, a vector of demographic household controls, year fixed effects ($\theta_{VFCR,t}$ and $\theta_{Z,t}$), household fixed effects ($\theta_{VFCR,n}$ and $\theta_{Z,n}$), and $\text{Control}_{n,t}$, which captures the general role of village size, in order to emphasize that the impact identified is specific to the post-intervention years.

We used two alternative specifications for $\text{Control}_{n,t}$, $\frac{1,000,000}{\# \text{ HHs in village}_{v,6}}$ and $\frac{1,000,000}{\# \text{ HHs in village}_{v,6}} * t$. Given the first-specification, $\hat{\alpha}_{2,Z}$, “levels” would capture the relationship between village size and the level of the outcome that is common to both the pre- and post-intervention years. In the latter specification, $\hat{\alpha}_{2,Z}$, “trends” captures the relationship between village size and the trend in the outcome variable. The level specification is of less interest, since our results are unlikely to be contaminated by levels differences; Household fixed effects $\theta_{Z,t}$ already capture persistent level differences (across households and villages), and our analysis will utilize household fixed effects. Moreover, the $\hat{\alpha}_{2,Z}$, “levels” is only identified from within-village variation in village size (i.e., the sizes of given villages varying over the years of the panel), which constitutes only 5 percent of the total variation in village size, and our analysis will only use village size in one year, the first year of the intervention ($t = 6$). The trend specification is therefore of more relevance.

¹⁵Forest conservation efforts have driven some redistricting decisions but these decisions have been largely haphazard and unsystematic. For discussions, see Pugenier (2001) and Gine (2005).

Table I presents $\hat{\alpha}_{2,Z}^{\text{“trends”}}$ results for the 37 different outcome variables Z_{it} from Kaboski and Townsend (2009), and three additional variables relevant to this study: investment probability, default probability, and total consumption. Together, these regressions cover the details of household income, consumption, investment, and borrowing activities. Only two of these 40 estimates are significant, even at a conservative 10 percent level; smaller villages were associated with higher growth in the fraction of income coming from rice and faster growth in the amount of credit from commercial banks. In terms of economic significance, this would mean that for the average village the rice fraction would fall by 1 percentage point a year less than in the largest village. Similarly, the amount of commercial bank credit would rise by 500 baht (12 dollars) a year more than in the largest village.¹⁶ Though not presented in the table, the estimates for $\alpha_{2,Z}^{\text{“levels”}}$ also show few significant relationships.¹⁷ We also note that our results are robust to whether or not these controls are included. Thus, we have a measure of confidence that pre-existing differences in levels or trends associated with inverse village size are few and small.

2.2 Reduced Form Impacts

The above regressions produce several interesting “impact” estimates $\hat{\alpha}_{1,Z}$ as reported in detail in our companion paper, Kaboski and Townsend (2009).¹⁸ With regard to credit, the program expanded village fund credit roughly one for one, with the coefficient $\hat{\alpha}_{1,VFCR}$ close to one. Second, total credit overall appears to have had a similar expansion, with an $\hat{\alpha}_{1,Z}$ near one and there is no evidence of crowding out in the credit market. Finally, the expansion did not occur through a reduction in interest rates. Indeed the $\hat{\alpha}_{1,Z}$ is positive, though small for interest rates.

Household consumption was obviously and significantly affected by the program, with a $\hat{\alpha}_{1,Z}$ point estimate near one. The higher level of consumption was driven by non-durable consumption and services, rather than durable goods. While the frequency of agricultural investments did increase mildly, total investment showed no significant response to the program. The frequency of households in default increased mildly in the second year, but default rates remained less than 15 percent of loans. Asset levels (including savings) declined in response to the program, while income growth increased weakly.¹⁹

Together, these results are puzzling. In a perfect credit, permanent income model, with no changes in prices, unsubsidized credit should have no effect, while subsidized credit would simply have an income effect. If credit did not need to be repaid, this income effect would be bounded above by the amount of credit injected. Yet repayment rates were actually quite high, with only 3 percent of village fund credit in default in the last year of the survey. But again, even if credit were not repaid, an income effect would produce at most a coefficient of the market interest rate (less than 0.07), i.e., the household would keep the principle of the one-time wealth shock and consume the interest. The fact that households appear to have

¹⁶More generally, one million divided by the number of household averages roughly 10,000 in our sample, so the economic magnitude on a per year basis is the coefficient multiplied by 10,000.

Note that the coefficient on investment probability is positive an order of magnitude larger than our results in Section 5, but the standard deviation is two orders of magnitude larger and so it is insignificant.

¹⁷Again using a more conservative ten percent level of significance only 3 out of 39 coefficients (8 percent) were significant. Small villages tended to have higher levels of short-term credit in fertilizer ($\hat{\alpha}_{2,Z}^{\text{“levels”}} = 1.14$ with a standard error of 0.50) and higher shares of total income from rice ($8.3e-6$, std. error $3.0e-6$) and other crops ($4.1e-6$, std. error $2.2e-6$). Thus, as villages grow, they appear to become somewhat less agrarian.

¹⁸The sample in Kaboski and Townsend (2009) varies slightly from the sample in this paper. Here we necessarily exclude 118 households who did not have complete set of data for all seven years. To avoid confusion, we do not report the actual Kaboski and Townsend (2009) estimates here.

¹⁹Wage income also increased in response to the shock, which is a focus of Kaboski and Townsend (2009). The increase is quite small relative to the increase in consumption, however, and so this has little promise in explaining the puzzles. We abstract from general equilibrium effects on the wage and interest rate in the model we present.

simply increased their consumption by the value of the funds lent is therefore puzzling. Given the positive level of observed investment, the lack of a response to investment might point to well-functioning credit markets, but the large response of credit and consumption indicate the opposite. Thus, the coefficients overall require a theoretical and quantitative explanation.

2.3 Underlying Environment

Growth, savings/credit, default, and investment are key features in the Thai villages during the pre-intervention period (as well as afterward). Households income growth averages 7 percent over the panel, but both income levels and growth rates are stochastic. Savings and credit are important buffers against income shocks (Samphantharak and Townsend (2008)), but credit is limited (Puentes (2008)). Income shocks are neither fully insured nor fully smoothed (Chiappori et al (2008)), and Karaivanov and Townsend (2008) conclude that savings and borrowing models and savings only models fit the data better than alternative mechanism design models. High income households appear to have access to greater credit. That is, among borrowing households, regressions of log short-term credit on log current income yield a coefficient of 0.32 (std. err.=0.02).

Related, default occurs in equilibrium, and appears to be one way of smoothing against shocks. In any given year, 19 percent of households are over three months behind in their payments on short-term (less than on year) debts. Default is negatively related to current income, but household consumption is substantial during periods of default, averaging 164 percent of current income, and positively related to income. Using only years of default, regressions of log consumption on log income yield a coefficient of regression of 0.41 (std. error=0.03).

Finally, investment plays an important role in the data, averaging 10 percent of household's income. It is lumpy, however. On average only 12 percent of households invest in any given year. Investment is large in years when investment occurs and highly skewed with a mean of 79 percent of total income and a median of 15 percent. When they invest, high income households make larger investments; a regression of log investment on the (log) predictable component of income yields a significant regression coefficient of 0.57 (std. error=0.15).²⁰ High income households still invest infrequently, however, and indeed the correlation between investment and predictable income is 0.02 and insignificant. Related, investment is not concentrated among the same households each year. If the average probability of investing (0.12) were independent across years and households, one would predict that $(1 - 0.88^5) = 47$ percent of households would invest at least once over the five years of pre-intervention data. This is quite close to the 42 percent that is observed.

The next section develops a model broadly consistent with this underlying environment.

3 Model

We address these key features of the data by developing a model of a household facing permanent and transitory income shocks and making decisions about consumption, low yield liquid savings, high yield illiquid investment and default. The household is infinitely-lived, and, in order to allow for growth, tractability requires that we make strong functional form assumptions.²¹ In particular, the problem is written so that all constraints are linear in the permanent component of income, so that the value function and policy functions can all be normalized by permanent income. We do this to attain a stationary, recursive problem.

²⁰These predictions are based on a regression of a regression of log income on: age of head of household, squared age, number of males, number of females, number of kids, and log assets.

3.1 Sequential Problem

At $t + 1$, liquid wealth L_{t+1} includes the principle and interest on liquid savings from the previous period $(1 + r) S_t$ (negative for borrowing) and current realized income Y_{t+1} :

$$L_{t+1} \equiv Y_{t+1} + S_t(1+r) \quad (2)$$

Following the literature on precautionary savings (e.g., Zeldes, 1989, Carroll, 1997, Gourinchas and Parker, 2001), current income Y_{t+1} consists of a permanent component of income P_{t+1} and a transitory one-period shock, U_{t+1} , additive in logs:

$$Y_{t+1} \equiv P_{t+1} U_{t+1} \quad (3)$$

We follow the same literature in modeling an exogenous component of permanent income that follows a random walk (again in logs) based on shock N_t with drift G . Meghir and Pistaferri (2004) have presented strong evidence for the importance of permanent income shocks in the U.S., and we believe that the standard ideas of permanent income shocks (e.g., long term illness or disability, obsolescence of specialized human capital, shocks affecting the profitability of businesses or capital) are at least as important in a developing country context. Nonetheless, our innovation in this paper is to also allow for endogenous increases in permanent income through investment.²² Investment is indivisible – the household makes a choice $D_{I,t} \in \{0, 1\}$ of whether to undertake a lumpy investment project of size I_t^* or to not invest at all. In sum,

$$P_{t+1} = P_t G N_{t+1} + R D_{I,t} I_t^* \quad (4)$$

Investment is also illiquid and irreversible, but again it increases permanent income, at a rate R , higher than the interest rate on liquid savings, r , and sufficiently high to induce investment for households with high enough liquidity. Having investment increase the permanent component of future income simplifies the model by allowing us to track only P_t rather than multiple potential capital stocks.²³ While we have endogenized an important element of the income process, we abstract from potentially endogenous decisions such as labor supply, and the linearity in R abstracts from any diminishing returns that would follow from a non-linear production function.

Project size is stochastic, governed by an exogenous shock i_t^* and proportional to the permanent component of income:

²¹We model an infinitely-lived household for several reasons. Using a life-cycle approach in the U.S., Gourinchas and Parker (2001) show that life-cycle savings plays a relatively smaller role until the last ten years before retirement. In the rural Thai context, there is no set retirement age or pension system, and households often include family from multiple generations. Deaton and Paxson (2000) show that profiles of household head age vs. household savings do not fit the life cycle theory well.

²²Low et al (forthcoming) endogenize permanent income in the U.S. context through participation and occupational mobility decisions.

²³This approach ignores many issues of investment “portfolio” decisions and risk diversification. Still, the lumpy investment does capture the important portfolio decision between a riskless, low yield, liquid asset and a risky, illiquid asset, which is already beyond what is studied in a standard bufferstock model. We can show this by defining $A_t \equiv P_t/R$ and using (2), (3), (4), and (7) to write:

$$A_t + S_t = (R U_t + G N_t) A_{t-1} + S_{t-1}(1+r) - C_t]$$

Physical assets A_t pay a stochastic gross return of $(R U_t + G N_t)$, while liquid savings pay a fixed return of $(1 + r)$.

$$I_t^* = i_t^* P_t \quad (5)$$

We assume that investment opportunities I_t^* are increasing in permanent income P_t , which the data seem to support. A more flexible specification would be $I_t^* = i_t^* P_t^\omega$. A regression of log investment on the log of the component of income predicted by observables (a proxy for P_t) yields a coefficient of 0.57 indicating an $\omega < 1$. Still, our assumption of linearity ($\omega = 1$) will be necessary for analytical tractability, and it will yield results consistent with investment decisions being uncorrelated with the predictable component of income (as described in Section 2.3).²⁴ The linearity we assume is consistent with the empirical literature, where large firms invest higher amounts, and so investment is typically scaled by size.

Liquid savings can be negative, but borrowing is bounded by a limit which is a multiple $\underline{s}$ of the permanent component of income. That is, when $\underline{s}$ is negative, borrowing is allowed, and the more negative it is, the more can be borrowed. This is the key parameter that we calibrate to the intervention:

$$S_t \geq \underline{s} P_t \quad (6)$$

For the purposes of this partial equilibrium analysis, this borrowing constraint is exogenous. It is not a natural borrowing constraint as in Aiyagari (1994) and therefore somewhat ad hoc, but such a constraint can arise endogenously in models with limited commitment (see Wright (2002)) or where lenders have rights to garnish a fraction of future wages (e.g., Lochner and Monge-Naranjo (2008)). Most importantly, it allows for default (see below), which is observed in the data and of central interest to microfinance interventions.

In period 0, the household begins with a potential investment project of size I_0^* , a permanent component of income P_0 , and liquid wealth L_0 all as initial conditions. The household's problem is to maximize expected discounted utility by choosing a sequence of consumption $C_t > 0$, savings S_t , and decisions $D_{I,t} \in \{0, 1\}$ of whether or not to invest:

$$\begin{aligned} V(L_0, I_0^*, P_0; \underline{s}) = & \max_{\{C_t > 0\}} E_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\rho}}{1-\rho} \right] \\ & \{S_{t+1}\} \\ & \{D_{I,t}\} \\ & \text{subject to (2), (3), (4), (5), (6), and} \\ & C_t + S_t + D_{I,t} I_t^* \leq L_t \end{aligned} \quad (7)$$

The expectation is taken over sequences of permanent income shocks N_t , transitory income shocks U_t , and investment size shocks i_t^* . These shocks are each i.i.d. and orthogonal to one another:

²⁴Households policies will be to invest in all project below a threshold I_t^* , call it $\tilde{I}_t^*$. If investment opportunities did not increase with P_t , i.e., $\omega = 0$, then high P_t households would invest at a higher rate than poor households, since the threshold $\tilde{I}_t^*$ would be higher for high P_t . We cannot solve this case, but we conjecture that it would be quantitatively important, since given the relatively low frequency of investment (12 percent) the cutoff $\tilde{I}_t^*$ would typically fall on the left-tail of the log normal i_t^* distribution where the density and inverse Mills ratio are high.

- N_t is random walk shock to permanent income. $\ln N_t \sim N(0, \sigma_N^2)$.
- U_t is a temporary (one period) income shock. $u_t \equiv \ln U_t \sim N(0, \sigma_u^2)$.
- i_t^* is project size (relative to permanent income). $\ln i_t^* \sim N(\mu_i, \sigma_i^2)$

If $\underline{s} < 0$, an agent with debt, i.e., $S_{t-1} < 0$, and a sufficiently low income shock may need to default. That is, with $L_t = Y_t + S_{t-1}(1+r)$, even with zero consumption and investment, the liquid assets budget constraint (7) could imply $S_t < \underline{s}P_t$. Essentially, given (6), a bad enough shock to permanent income (i.e., a low N_t) can produce a “margin call” on credit that exceeds current liquidity.

In this case, we assume default allows for a minimum consumption level that is proportional to permanent income ($\underline{c}P_t$). Defining the default indicator, $D_{def,t} \in \{0, 1\}$, this condition for default is expressed:

$$D_{def,t} = \begin{cases} 1, & \text{if } (\underline{s} + \underline{c})P_t < L_t \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

and the defaulting household's policy for the period becomes:

$$\begin{aligned} C_t &= \underline{c}P_t \\ S_t &= \underline{s}P_t \\ D_{i,t} &= 0 \end{aligned}$$

This completes the model. The above modeling assumptions are strong and not without costs. Still, as we have seen, they are motivated by the data, and they do have analytical benefits beyond allowing us to deal easily with growth. First, the model is simple and has limited heterogeneity, but consequently has a low dimension, tractable state space $\{L, I^*, P\}$ and parameter space $\{r, \sigma_N, \sigma_u, G, \underline{c}, \beta, \rho, \mu_i, \sigma_i, \underline{s}\}$. Hence, the role of each state and parameter can be more easily understood. Furthermore, the linearity of the constraints in P_t reduces the dimensionality of the state space to two, which allows for graphical representation of policy functions (in Section 5.2). The next subsection derives the normalized, recursive representation.²⁵

3.2 Normalized and Recursive Problem

Above, we have explicitly emphasized the value function's dependence on s , since this will be the parameter of most interest in considering the microfinance intervention in Section 5. We drop this emphasis in the simplifying notation that follows. Using lower case variables to indicate variables normalized²⁶ by permanent income, the recursive problem becomes:

²⁵Since all conditions are linear P_t , we avoid the problems that unbounded returns lead to in representing infinite horizon models in recursive fashion (see Stokey, Lucas, and Prescott (1989)). In particular, the conditions for the equivalence of the recursive and sequential problems and existence of the steady state are straightforward extensions of conditions given in Alvarez and Stokey (1998) and Carroll (2004). In particular, for $\rho < 1$, G and $RE[i^*]$ must be sufficiently bounded.

²⁶Here the decision whether to invest d_i is not a normalized variable and is in fact identical to D_i in the earlier problem. We have denoted it in lower case to emphasize that it will depend only on the normalized states l and i^* .

$$\begin{aligned}
 V(L, I^*, P) &\equiv P^{1-\rho} v(l, i^*) \\
 v(l, i^*) &= \max_{c, s', d_l} \frac{c^{1-\rho}}{1-\rho} + \beta E \left[(p')^{1-\rho} v(l', i^{*'}) \right] \text{ subject to} \\
 \lambda: c + s + d_l i^* &\leq l \quad \text{from (7)}
 \end{aligned} \tag{9}$$

$$\begin{aligned}
 \varphi: s &\geq \underline{s} \quad \text{from (6)} \\
 p' &= GN' + Rd_l i^* \quad \text{from (4)} \\
 l' &= y' + \frac{s(1+r)}{p'} \quad \text{from (2)}
 \end{aligned} \tag{10}$$

$$y' = U' \quad \text{from (3)} \tag{11}$$

We further simplify by substituting l' and y' into the continuation value using (10) and (11), and substituting out s using the liquidity budget constraint (9), which will hold with equality, to yield:

$$v(l, i^*) = \max_{c, d_l} \frac{c^{1-\rho}}{1-\rho} + \beta E \left[(p')^{1-\rho} v \left(U' + \frac{(1+r)(l - c - d_l i^*)}{p'}, i^{*'} \right) \right] s.t. \tag{12}$$

$$\varphi: (l - c - d_l i^*) \geq \underline{s} \tag{13}$$

$$p' = GN' + Rd_l i^* \tag{14}$$

The normalized form of the problem has two advantages. First, it lowers the dimensionality of the state variable to two. Second, it allows the problem to have a steady state solution. Using $*$ to signify optimal decision rules, the necessary conditions for optimal consumption c^* and investment decisions d_{l^*} are:²⁷

$$(c^*)^{-\rho} = \beta(1+r)E \left[(p')^{-\rho} \frac{\partial v}{\partial l} \left(U' + \frac{(1+r)(l - c^* - d_{l^*} i^*)}{p'}, i^{*'} \right) \right] + \varphi \tag{15}$$

$$\begin{aligned}
 \frac{c^{1-\rho}}{1-\rho} + \beta E \left[(p')^{1-\rho} v \left(U' + \frac{(1+r)(l - c^* - d_{l^*} i^*)}{p'}, i^{*'} \right) \right] &\geq \\
 \frac{c^{1-\rho}}{1-\rho} + \beta E \left[(p')^{1-\rho} v \left(U' + \frac{(1+r)[l - c^{**} - (1 - d_{l^*}) i^*]}{p'}, i^{*'} \right) \right] &
 \end{aligned} \tag{16}$$

²⁷Although the value function is kinked, it is differentiable almost everywhere, and the smooth expectation removes any kink in the continuation value.

Equation (15) is the usual credit constrained Euler equation. The constraint φ is only non-zero when the credit constraint (13) binds, i.e., $c^* = l - \underline{s} - d_I^* i^*$. Equation (16) ensures that the value given the optimal investment decision d_I^* , exceeds the maximum value given the alternative, $1 - d_I^*$, c^{**} indicates the optimal consumption under this alternative investment decision (i.e., c^{**} satisfies the analog to (15) for $1 - d_I^*$).

In practice, the value function and optimal policy functions must be solved numerically, and indeed the indivisible investment decision complicates the computation.²⁸

Figure 2 presents a three-dimensional graph of a computed value function. The flat portion at very low levels of liquidity l comes from the minimum consumption and default option. The dark line highlights a groove going through the middle of the value function surfaces along the critical values at which households first decide to invest in the lumpy project. Naturally, these threshold levels of liquidity are increasing in the size of the project. The slope of the value function with respect to l increases at this point because the marginal utility of consumption increases at the point of investment.²⁹ Consumption actually *falls* as liquidity increases beyond this threshold.

Figure 3, panel A illustrates this more clearly by showing a cross-section of the optimal consumption policy as a function of normalized liquidity for a given value of i^* . At the lowest values, households are in default. At low values of liquidity, no investment is made, households consume as much as possible given the borrowing constraint, and hence the borrowing constraint holds with equality. At higher liquidity levels, this constraint is no longer binding as savings levels s exceed the lower bound $\underline{s}$. At some crucial level of liquidity l^* , the household chooses to invest in the lumpy project, at which point consumption falls and the marginal propensity to consume out of additional liquidity increases. Although not pictured, for some parameter values (e.g., very high R), the borrowing constraint can again hold with equality, and marginal increases in liquidity are used for purely for consumption.³⁰

Panel B of Figure 3 shows the effect of a surprise permanent decrease in $\underline{s}$ on the optimal consumption policy for the same given value of i^* . Consumption increases for liquidity levels in every region, except for the region that is induced into investing by more access to borrowing.

An additional interesting prediction of the model is that for a given level of borrowing ($s_t < 0$), a household that invests ($d_{I,t} = 1$) has a lower probability of default next period. Conditional on investing, the default probability is further decreasing in the size of investment. Thus, other things equal borrowing to invest leads to less default than borrowing to consume because investment increases future income and therefore ability to repay. The maximum amount of debt that can be carried over into next period (i.e., $-\underline{s}P_t$) is proportionate to permanent income. Because investment increases permanent income, it

²⁸Details of the computational approach and codes are available from the authors upon request.

²⁹Given the convex kink in the value function, households at or near the kink would benefit from lotteries, which we rule out consistent with the idea that borrowing and lending subject to limits is the only form of intermediation.

³⁰Using a bufferstock model, Zeldes (1989) derived reduced form equations for consumption growth, and found that consumption growth was significantly related to current income, but only for low wealth households, interpreted as evidence of credit constraints. We run similar consumption growth equations that also contain investment as an explanatory variable:

$$\ln C_{n,t+1} / C_{n,t} = X_{n,t} \beta_1 + \beta_2 Y_{n,t} + \beta_3 I_{n,t} + \varepsilon_{n,t}$$

For the low wealth sample, we find significant estimates $\hat{\beta}_2 < 0$ and $\hat{\beta}_3 > 0$, which is consistent with the prediction of investment lowering current consumption (thereby raising future consumption growth).

increases the borrowing limit *next period*, and therefore reduces the probability of a “margin call” on outstanding debt.

One can see this formally by substituting the definitions of liquidity (2) and income (3), and the law of motion for permanent income (4) into the condition for default (8) to yield:

$$E(D_{def,t+1}|S_t, P_t, D_{I,t}, I_t^*) = \Pr \left[U_{t+1} < (\underline{s} + \underline{c}) - \frac{S_t}{(P_t N_{t+1} G + R D_{I,t} I_t^*)} \right] \quad (17)$$

Since S_t is negative and R is positive, the right-hand side of the inequality is decreasing in both $D_{I,t}$ and I_t^* . Since both N_{t+1} and U_{t+1} are independent of investment, the probability is therefore decreasing in $D_{I,t}$ and I_t^* .

4 Estimation

This section addresses the data used and then the estimation approach. The model is quite parsimonious with a total of 11 parameters. Due to poor identification, we calibrate the return on investment parameter, R , using a separate data source. After adding classical measurement error on income with log variance σ_E , we estimate the remaining parameters, $\theta = \{r, \sigma_N, \sigma_u, \sigma_E, G, \underline{c}, \beta, \rho, \mu_i, \sigma_i, \underline{s}\}$ via MSM using the optimal weighting matrix. This estimation is performed using five years (1997–2001) of pre-intervention data, so that $t = 1$ corresponds to the year 1997.

4.1 Data

The data come from the Townsend Thai data project, an ongoing panel dataset of a stratified, clustered, random sample of institutions (256 in 2002), households (960 each year, 715 with complete data in the pre-experiment balanced panel used for estimation, and 700 in 2002 and 2003, respectively, which are used to evaluate the model’s prediction), and key informants for the village (64, one in each village). The data are collected from sixty-four villages in four provinces: Buriram and Srisaket in the Northeast region, and Lopburi and Chachoengsao in the Central region. The components used in this study include detailed data from households and household businesses on their consumption, income, investment, credit, liquid assets and the interest income from these assets, as well as village population data from the village key informants. All data has been deflated using the Thai consumer price index to the middle of the pre-experiment data, 1999.

The measure of household consumption we use (denoted $\tilde{C}_{n,t}$ for household n at time t) is calculated using detailed data on monthly expenditure data for thirteen key items, and scaled up using weights derived from the Thai Socioeconomic Survey.³¹ In addition, we include household durables in consumption, though durables play no role in the observed increases in consumption. The measure of investment ($\tilde{I}_{n,t}$) we use is total farm and business investments, including livestock and shrimp/fish farm purchases.

We impute default each year for households who report one or more loans due in the previous 15 months that are outstanding at least three months. Note that (i) this includes all loans, and not just short-term, since any (non-voluntary) default indicates a lack of available liquidity, and (ii) due dates are based on the original terms of the loan, since changes in duration are generally a result of default.³² This only approximates default in the model,

³¹The tildes represent raw data which will be normalized in Section 4.3.1.

and it may underestimate default because of underreporting, but overestimate default as defined in the model or to the extent that late loans are eventually repaid.

The income measure we use (denoted $\tilde{Y}_{n,t}$) includes all agricultural, wage, business and financial income (net of agricultural and business expenses) but excludes interest income on liquid assets such as savings deposits as in the model. Our savings measure ($S_{n,t}$) includes not only savings deposits in formal and semi-formal financial institutions, but also the value of rice holdings in the household. Cash holdings are unfortunately not available. The measure of liquid credit ($CR_{n,t}$) is short-term credit with loan durations of one year or less. The measurement of interest income on liquid savings ($EARNED\_INT_{n,t}$) is interest income in year t on savings in formal and semi-formal institutions. The interest owed on credit ($OW\_ED\_INT_{n,t}$) is the reported interest owed on short-term credit.

While the data is high quality and detailed, measurement error is an important concern. Net income measures are complicated when expenditures and corresponding income do not coincide in the same year, for example. If income is measured with error, the amount of true income fluctuations will be overstated in the data, and household decisions may appear to be less closely tied to transitory income shocks, hence credit constraints may not appear to be important. Consumption and investment may also suffer from measurement error, but classical measurement error will just add additional variation to these endogenous variables will not effect the moments, only the weighting matrix. A major source of measurement error for interest is that savings and borrowing may fluctuate within the year, so that the annual flow of both earned and paid interest may not accurately reflect interest on the end-of-year stocks contained in the data. This measurement error will assist in the estimation.

Table II presents key summary statistics for the data.

4.1.1 Adjusting the Data for Demographic and Cyclical Variation—The model is of infinitely lived dynasties that are heterogeneous only in their liquidity, permanent income, and potential investment. That is, in the model, the exogenous sources of variation among households come from given differences in initial liquidity or permanent income, and histories of shocks to permanent income, transitory income, and project size. Clearly, the data, however, contain important variation due to heterogeneity in household composition, business cycle and regional variation, and unmodeled aspects of unobserved household heterogeneity. Ignoring these sources of variation would be problematic. For household composition, to the extent that changes in household composition are predictable, the variance in income changes may not be capturing uncertainty but also predictable changes in household composition. Likewise, consumption variation may not be capturing household responses to income shocks but rather predictable responses to changes in household composition. Failure to account for this would likely exaggerate both the size of income shocks and the response of household consumption to these shocks. In the data, the business cycle (notably the financial crisis in 1997 and subsequent recovery) also plays an important role in household behavior, investment and savings behavior in particular. Although our post-program analysis will focus on the across-village differential impacts of the village fund program, and not merely the time-changes, we do not want to confound the impacts with business cycle movements. Finally, differences in consumption, for example, across households may tell us less about past and current income shocks, and more about unobserved differences in preferences or consumption needs.

³²According to this definition, default probability is about 19 percent, but alternative definitions can produce different results. The probability for short-term loan alone is just 12 percent, for example. On the other hand, relabeling all loans from non-family sources that have no duration data whatsoever as in default yields a default probability of 23 percent. Our results for consumption and default hold for the higher rates of default.

A common approach in structural modeling is to account for these sources of heterogeneity and predictable variation across households explicitly in the model and estimation (see Keane and Wolpin (1994, 1997, 2001)). These methods have the advantage of incorporating this heterogeneity into the household decision making process, but they typically require finite horizons and discretizing the choice variables (e.g., consumption or savings). Within the bufferstock literature, a common approach has been to instead purge business cycle and household composition variation from the data (e.g., Gourinchas and Parker (2001), Carroll and Samwick, (1998)). Though the former approach is certainly of interest, given the continuity of consumption, our infinite horizon, and the precedent within the buffer stock literature, we follow the latter approach. We return to the issue of heterogeneity in the concluding section.

Specifically, we run linear regressions of log income, log consumption, and liquidity over income. (We do not take logs of liquidity, since it takes both positive and negative values, but instead normalize by income so that high values do not carry disproportionate weight.)³³ The estimated equations are:

$$\begin{aligned}\ln \tilde{Y}_{n,t} &= \gamma_Y \mathbf{X}_{n,t} + \theta_{Y,j,t} + e_{Y,n,t} \\ \tilde{L}_{n,t} / \tilde{Y}_{n,t} &= \gamma_L \mathbf{X}_{n,t} + \theta_{L,j,t} + e_{L,n,t} \\ \ln \tilde{C}_{n,t} &= \gamma_C \mathbf{X}_{n,t} + \theta_{C,j,t} + e_{C,n,t} \\ \ln \tilde{D}_{n,t} &= \gamma_D \mathbf{X}_{n,t} + \theta_{D,j,t} + e_{D,n,t}\end{aligned}$$

where $\mathbf{X}_{n,t}$ is a vector of household composition variables (i.e., number of adult males, number of adult females, number of children, male head of household dummy, linear and squared terms of age of head of household, years education of head of household, and a household-specific fixed effect) for household n at time t and $\theta_{\cdot,j,t}$ is a time t -specific effect that varies by region j and captures the business cycle. These regressions are run using only the pre-program data, 1997–2001, which ensures that we do filter out the effects of the program itself. Unfortunately, the pre-program, time-specific effects cannot be extrapolated for the post-program data, so we rely on across village, within-year variation to evaluate the model's predictions. The R^2 values for the four regressions are 0.63, 0.34, 0.76, and 0.31, respectively, so the regressions are indeed accounting for a great deal of heterogeneity and variation.

For the full sample, 1997–2003, we construct the adjusted data for a household with mean values of the explanatory variables ($\bar{\mathbf{X}}$ and $\bar{\theta}_{\cdot,j}$) using the estimated coefficients and residuals:

$$\begin{aligned}\ln Y_{nt} &= \hat{\gamma}_Y \bar{\mathbf{X}} + \bar{\theta}_{Y,j} + g_Y(t - 1999) + \hat{e}_{Y,n,t} \\ L_{nt} / Y_{nt} &= \hat{\gamma}_L \bar{\mathbf{X}} + \bar{\theta}_{L,j} + \hat{e}_{L,n,t} \\ \ln C_{nt} &= \hat{\gamma}_C \bar{\mathbf{X}} + \bar{\theta}_{C,j} + g_C(t - 1999) + \hat{e}_{C,n,t} \\ D_{nt} &= \hat{\gamma}_D \bar{\mathbf{X}} + \bar{\theta}_{D,j} + \hat{e}_{D,n,t}\end{aligned}$$

where g_Y and g_C are the average growth rates of the trending variables, income and consumption, respectively, in the pre-program data. Next, we use a multiplicative scaling term to ensure that average income, liquidity ratios, consumption, and default are equal in

³³As noted before, 79 of the original 960 households realized negative income (net of business and agricultural income) at some point in the pre-intervention sample. The model yields only positive income, and so these households were dropped.

the raw and adjusted data. Finally, we construct investment data $I_{n,t}$ by multiplying the measured investment/income ratios ($\tilde{I}_{nt}/\tilde{Y}_{nt}$) by the newly constructed income data $Y_{n,t}$.

4.2 Returns on Investment

In principle, income growth and investment data should tell us something about the return on investment, R . In practice, however, the parameter cannot be well estimated because investment data itself is endogenous to current income, and also because investment occurs relatively infrequently. We instead use data on physical assets rather than investment, and we calibrate R to match cross-sectional relationship between assets and income.

To separate the effect of assets and labor quality on income, we assume that all human capital investments are made prior to investments in physical assets. Let $t - J$, indicate the first year of investing in physical assets. That is, substituting the law of motion for permanent income, equation (4), J times recursively into the definition of actual income, equation (3), yields:

$$Y_t = \underbrace{\left[P_{t-J} G^J \prod_{j=1}^J N_{t+1-j} \right] U_t}_{\text{income of investment prior to } t-J} + R \underbrace{\left[\sum_{j=1}^J I_{t-j} G^{j-1} \prod_{k=1}^j N_{t+1-k} \right] U_t}_{\text{income from investment after } t-J}$$

The first term captures income from the early human capital investments, which we measure by imputing wage income from linear regressions of wages on household characteristics (sex, age, education, region). The second term involves the return R multiplied by the some of the past J years of investments (weighted by the deterministic and random components of growth.) We measure this term using current physical assets. That is, R is calibrated using the following operational formula:

$$\varepsilon_R = Y_t - \text{imputed labor income}_t - R(\text{physical assets}_t)$$

We have the additional issue of how to deal with the value of housing and unused land. Neither source of assets contributes to Y_t , so we would ideally exclude them from the stock of assets.³⁴ Using data on the (i) value of the home, (ii) value of the plot of land including the home, and (iii) the value of unused or community use land, we construct three variants of physical assets.

We use a separate data set, the Townsend Thai Monthly Survey, to calibrate this return. The data is obtained from different villages, but the same overall survey area, and the monthly has the advantage of including wage data used to impute the labor income portion of total income.

We use a procedure which is analogous to GMM. We choose R to set the average ε_R to zero in the sample of households. The baseline value (which excludes categories (i)–(iii) from assets) yields $R = 0.11$, while including (iii), or (ii) and (iii), yield $R = 0.08$ and $R = 0.04$, respectively. If we choose R to solve $\varepsilon_R = 0$ for each household, then the median R values are identical to our estimates. Not surprisingly, R substantially varies across households, however. This is likely due in part because permanent shock histories and current transitory shocks differ across households, but also in part because households face different ex ante returns to investment.

³⁴Our measure of Y_t does not include imputed owner occupied rent.

4.3 Method of Simulated Moments

In estimating, we introduce multiplicative measurement error in income which we assume is log normally distributed with zero log mean and standard deviation σ_E . Since liquidity L_t is calculated using current income, measurement error will also produce measurement error in liquidity.

We therefore have eleven remaining parameters $\theta = \{r, G, \sigma_N, \sigma_u, \sigma_E, \underline{c}, \beta, \rho, \mu_i, \sigma_i, \underline{s}\}$, which are estimated using a Method of Simulated Moments. The model parameters are identified jointly by the full set of moments. We include, however, an intuitive discussion of the specific moments that are particularly important for identifying each parameter.

The first two types of moments help identify the return to liquid savings, r :

$$\begin{aligned}\varepsilon_s(X, r) &= EARNED\_INT_t - rS_{t-1} \\ \varepsilon_{cr}(X, r) &= OWED\_INT_t - rCR_{t-1}\end{aligned}$$

In ε_s , S_{t-1} is liquid savings in the previous year, while $EARNED\_INT_t$ is interest income received on this savings. Likewise, in ε_{cr} , CR is outstanding short-term credit in the previous year, and $OWED\_INT$ is the subsequent interest owed on this short-term credit in the following year.³⁵

The remaining moments require solving for consumption, $C(L_t, P_t, I_t^*; \theta) = P_t c(l_t, i_t^*; \theta)$, investment decisions, $D_i(L_t, P_t, I_t^*; \theta) = d_i(l_t, i_t^*; \theta)$, and default decisions, $D_{def}(L_t, P_t; \theta) = d_{def}(l_t; \theta)$, where we have now explicitly denoted the dependence of policy functions on the parameter set θ . We observe data on decisions, $C_t, I_t, D_{def,t}$, and states L_t and Y_t . Our strategy is to use these policy functions to define deviations of actual variables (policy decisions and income growth) from the corresponding expectations of these variables conditional on L_t and Y_t .³⁶ By the Law of Iterated Expectations, these deviations are zero in expectation and therefore valid moment conditions. With simulated moments, we calculate these conditional expectations by drawing series of shocks for U_t, N_t, i_t^* , and measurement error for a large sample, simulating, and taking sample averages. Details are available upon request.

The income growth moments help to identify the income process parameters and are derived from the definition of income and the law of motion for permanent income, equations (3) and (4).³⁷ Average income growth helps identify the drift component of growth income growth, G :

$$\varepsilon_g(L_t, Y_t, Y_{t+1}; \theta) = \ln(Y_{t+1}/Y_t) - E[\ln(Y_{t+1}/Y_t) | L_t, Y_t]$$

The variance of income growth over different horizons ($k=1 \dots 3$ -year growth rates, respectively) helps identify standard deviation of transitory and permanent income shocks, σ_u and σ_N , since transitory income shocks add the same amount of variance to income growth regardless of horizon k , whereas the variance contributed by permanent income shocks increases with k . The standard deviation of measurement error σ_E will also play a strong role in measured income growth. The deviations are defined as:

³⁵In the data there are many low interest loans, and the average difference between households interest rates on short term borrowing and saving is small, just 2 percent.

³⁶Since L_t requires the previous years savings S_{t-1} , these moments are not available in the first year.

³⁷Carroll and Samwick (1997) provide techniques for estimating the income process parameters G, σ_N , and σ_u without solving the policy function. These techniques cannot be directly applied in our case, however, since income is depends on endogenous investment decisions.

$$\varepsilon_{v,k}(L_t, Y_t, Y_{t+k}; \theta) = \left[\begin{array}{c} \ln(Y_{t+k}/Y_t) \\ -E[\ln(Y_{t+k}/Y_t)|L_t, Y_t] \end{array} \right]^2 - E \left[\left[\begin{array}{c} \ln(Y_{t+k}/Y_t) \\ -E[\ln(Y_{t+k}/Y_t)|L_t, Y_t] \end{array} \right]^2 | L_t, Y_t \right]$$

for $k=1, 2, 3$

We identify minimum consumption, $\underline{c}$; the investment project size distribution parameters μ_i and σ_i ; the preference parameters β and ρ , and the variance of measurement error σ_E using moments on consumption decisions, investment decisions, and the size of investments. Focusing on both investment probability and investment size should help in separately identifying the mean (μ_i) and standard deviation (σ_i) of the project size distribution. Focusing on deviations in log consumption, investment decisions, and log investments (when investments are made):

$$\begin{aligned} \varepsilon_c(C_t, L_t, Y_t; \theta) &= C_t - E[C_t | L_t, Y_t] \\ \varepsilon_D(D_{I,t}, L_t, Y_t; \theta) &= D_{I,t} - E[D_{I,t} | L_t, Y_t] \\ \varepsilon_I(D_{I,t}, I_t, L_t, Y_t; \theta) &= D_{I,t} I_t - E[D_{I,t} I_t^* | L_t, Y_t], \end{aligned}$$

we are left with essentially three moment conditions for five parameters:

$$E[\varepsilon_c] = 0 \quad E[\varepsilon_D] = 0 \quad E[\varepsilon_I] = 0$$

However, we gain additional moment conditions by realizing that since these deviations are conditional on income and liquidity, their interaction with functions of income and liquidity should also be zero in expectation. Omitting the functional dependence of these deviations, we express below the remaining six valid moment conditions:

$$\begin{aligned} E[\varepsilon_c \ln Y_t] &= 0 & E[\varepsilon_D \ln Y_t] &= 0 & E[\varepsilon_I \ln Y_t] &= 0 \\ E[\varepsilon_c (L_t/Y_t)] &= 0 & E[\varepsilon_D (L_t/Y_t)] &= 0 & E[\varepsilon_I (L_t/Y_t)] &= 0 \end{aligned}$$

Intuitively, in expectation, the model should match average log consumption, probability of investing, and log investment across all income and liquidity levels, e.g., not over-predicting at low income or liquidity levels, while underpredicting at high levels. These moments play particular roles in identifying measurement error shocks σ_E and $\underline{c}$, in particular. If the data shows less response of these policy variables to income than predicted, that could be due to a high level of measurement error in income. Similarly, high consumption at low levels of income and liquidity in the data would indicate a high level of minimum consumption $\underline{c}$.

Finally, given $\underline{c}$, default decision moments are used to identify the borrowing constraint $\underline{g}$, which can be clearly seen from equation (8):

$$\varepsilon_{def}(L_t, Y_t, D_{def,t}) = D_{def,t} - E[D_{def,t} | L_t, Y_t]$$

In total, we have 16 moments to estimate 11 parameters.

4.4 Estimation Results

Table III presents the estimation results for the structural model as well as some measures of model fit. The interest rate $\hat{r}$ (0.054) is midway between the average rates on credit (0.073) and savings (0.035), and is quite similar to the six percent interest rate typically charged by

village funds. The estimated discount factor $\hat{\beta}$ (0.915) and elasticity of substitution $\hat{\rho}$ (1.16) are within the range of usual values for bufferstock models. The estimated standard deviations of permanent $\hat{\sigma}_N$ (0.31) and transitory $\hat{\sigma}_U$ (0.42) income shocks are about twice those for wage earners in the United States (see Gourinchas and Parker, 2002), but reflect the higher level of income uncertainty of predominantly self-employed households in a rural, developing economy. In contrast, the standard deviation of measurement error $\hat{\sigma}_E$ (0.15) is much smaller than that of actual transitory income shocks, and is the only estimated parameter that is not significantly different from zero. The average log project size $\hat{\mu}_i$ greatly exceeds the average size of actual investments (i.e., $\log I_t/Y_t$) in the data (1.47 vs. -1.96), and there is a greater standard deviation in project size $\hat{\sigma}_i$ than in investments in the data (2.50 vs. 1.22). In the model, these difference between the average sizes of realized investment and potential projects stem from the fact that larger potential projects are much less likely to be undertaken.³⁸ The estimated borrowing constraint parameter $\hat{g}$ indicates that agents could borrow up to about 8 percent of their annual permanent income as short-term credit in the baseline period. (In the summary statistics of Table II, credit averages about 20 percent of annual income, but liquid savings net of credit, the relevant measure, is actually positive and averages 9 percent of income.) The value of $\hat{c}$ indicates consumption in default is roughly half of the permanent component of income.

Standard errors on the model are relatively small. We attempt to shed light on the importance of each of the 16 moments to identification of each the 11 parameters, but this is not trivial to show. Let ε be the (16-by-1) vector of moments and W , the (16-by-16) symmetric weighting matrix, then the criterion function is $\varepsilon'W\varepsilon$ and the variance-covariance matrix is $[\varepsilon'W\varepsilon]^{-1}$. The minimization condition for the derivative of the criterion function is

then $2\varepsilon'W\frac{\partial \varepsilon}{\partial \theta}=0$. Table IV presents $\frac{\partial \varepsilon}{\partial \theta}$, a 16-by-11 matrix showing the sensitivity of each moment to any given parameter. The influence of the parameter on the criterion function involves $2\varepsilon'W$, which has both positive and negative elements, however. Hence, the magnitudes of the elements in Table IV vary substantially across parameters and moments. W is also not a simple diagonal matrix so that the parameters are jointly identified. Some moments are strongly affected by many parameters (e.g., income growth and variances), while some parameters have strong effects on many moments (e.g., r , G , and β).

Still, the partial derivatives confirm the intuition above, in that the moments play a role in pinning down the parameters we associate with them. In particular, the interest rate r is the only parameter in the interest moments (rows ε_S and ε_{CR}). While σ_N is relatively more important for the variance of two and three-year growth rates (rows $\varepsilon_{V,2}$ and $\varepsilon_{V,3}$), σ_U is important for the variance of one-year growth rates (row $\varepsilon_{V,1}$). σ_E has important effects on the variance of income growth (rows $\varepsilon_{V,1}$, $\varepsilon_{V,2}$ and $\varepsilon_{V,3}$), but also the interaction of consumption and investment decisions with Y ($\varepsilon_C * \ln Y$, $\varepsilon_D * \ln Y$, and $\varepsilon_I * \ln Y$) and L/Y (rows $\varepsilon_C * L/Y$, $\varepsilon_D * L/Y$, and $\varepsilon_I * L/Y$). (These moments are even more strongly affected by r , σ_N , G , β , and ρ , however.) The utility function parameters β and ρ have the most important effect on consumption and investment moments (rows $\varepsilon_C - \varepsilon_I * L/Y$). Also, while μ_i and σ_i also affect income growth variance (rows $\varepsilon_{V,1}$, $\varepsilon_{V,2}$ and $\varepsilon_{V,3}$), the investment probability and investment level moments (rows $\varepsilon_D - \varepsilon_I * L/Y$) also help identify them. Finally, both $\underline{g}$ and $\underline{c}$ affect default similarly, but have opposite-signed effects on the interaction of measured income and liquidity ratios with investment (rows $\varepsilon_D * \ln Y$, $\varepsilon_D * L/Y$, $\varepsilon_I * \ln Y$, and $\varepsilon_I * L/Y$) and, especially, consumption (rows $\varepsilon_C * \ln Y$ and $\varepsilon_C * L/Y$) decisions.

In terms of fit, the model does well in reproducing average default probability, consumption, investment probability and investment levels (presented in Table III), and indeed deviations

³⁸In the model, the average standard deviation of log investment (when investment occurs) is 1.37, close to the 1.22 in the data.

are uncorrelated with log income or liquidity ratios. Still, we can easily reject the overidentifying restrictions in the model, which tells us that the model is not the real world. The large J-statistic in the bottom-right of Table III is driven by two sets of moments.³⁹ First, the estimation rejects that the savings and borrowing rates are equal.⁴⁰ Second, the model does poorly in replicating the volatility of the income growth process, yielding too little volatility.

We suspect this is the result of the income process and our statistical procedures failing to adequately capture cyclical effects of income growth, in particular the Thai financial crisis and recovery of 1997 and 1998 (survey years 1998 and 1999, respectively). Only *mean* time-varying volatility is extracted from the data using our regression techniques, but the crisis presumably affected the variance as well.⁴¹ Excluding the crisis from the pre-sample is not possible, since it would leave us just one year of income growth to identify both transitory and permanent income shocks. An alternative estimation that uses only data from 2000 and 2001, except for 1999 data used to create two-year income growth variance moments, produced estimates with wide standard errors that were not statistically different from the estimates above. The only economically significant difference was a much lower borrowing constraint ($\bar{g} = -0.25$), which is consistent with an expansion of credit observed in the Thai villages even pre-intervention. Recall that this trend is not related to village size, however.

Another way of evaluating the within-sample fit of the model is to notice that it is comparable to what could be obtained using a series of simple linear regressions estimating 11 coefficients (rather than 11 parameters estimated by the structural model). By construction, the nine moments defined on consumption, investment probability, and investment levels could be set equal to zero by simply regressing each on a constant, log income, and liquidity ratios. This would use nine coefficients. The two remaining coefficients could simply be linear regressions of growth and default on constant terms (i.e., simple averages). These linear regressions would exactly match the eleven moments that we only nearly fit. On the other hand, these linear predictors would predict no income growth volatility, and would have nothing to say about the interest on savings and credit.

So the result on the fit of the model are mixed. However, we view the model's ability to make policy predictions on the impact of credit as a stronger basis for evaluating its usefulness. We consider this in the next section.

5 Million Baht Fund Analysis

This section introduces the Million Baht fund intervention into the model, examines the model's predictions relative to the data, presents a normative evaluation of the program, and then presents alternative analyses using the structural model.

³⁹The J-statistic is the number of households (720) times $\varepsilon'W\varepsilon$. Since, W is symmetric, we can rewrite this as $\tilde{\varepsilon}'\tilde{\varepsilon}$. The major elements of the summation $\varepsilon'\varepsilon$ are 0.02 (ε_S), 0.02 (ε_{CP}), 0.03 ($\varepsilon_{V,1}$), 0.04 ($\varepsilon_{V,2}$), while the others are all less than 0.01.

⁴⁰It would be straightforward to allow for different borrowing and saving rates. This would lead to a kink in the budget constraint, however. The effect would that one would never observe simultaneous borrowing and saving and there would be a region where households neither save nor borrow. In the data, simultaneous short-term borrowing and saving is observed in 45 percent of observations, while having neither savings nor credit is observed in only 12 percent.

⁴¹We know from alternative estimation techniques that the model does poorly in matching year-to-year fluctuations in variables. In the estimation we pursue, we construct moments for consumption, investment, etc., that are based only on averages across the four years. For income growth volatility, the moments necessarily have a year-specific component.

5.1 Relaxation of Borrowing Constraints

We incorporate the injection of credit into the model as a surprise decrease in $\underline{s}$.⁴² That is, for each of sixty four villages, indexed by v , we calibrate the new, reduced constraint under the million baht fund intervention $\underline{s}_v^{mb}$ as the level for which our model would predict one million baht of additional credit relative to the baseline at $\underline{s}$. We explain this mathematically below.

Define first the expected borrowing of a household n with the Million Baht Fund intervention:

$$E[B_{n,t,v}|L_{n,t}, Y_{n,t}; \underline{s}_v^{mb}] = E\left\{I_{<0}\left[L_t - C(L_t, P_t, I_t^*; \underline{s}_v^{mb}) - D_t(L_t, P_t, I_t^*; \underline{s}_v^{mb})I_t^*\right] | L_{n,t}, Y_{n,t}\right\}$$

and in the baseline without the intervention:

$$E[B_{n,t,v}|L_{n,t}, Y_{n,t}; \underline{s}] = E\left\{I_{<0}\left[L_t - C(L_t, P_t, I_t^*; \underline{s}) - D_t(L_t, P_t, I_t^*; \underline{s})I_t^*\right] | L_{n,t}, Y_{n,t}\right\}$$

where $I_{<0}$ is shorthand notation for the indicator function that the bracketed expression is negative (i.e., borrowing and not savings). On average, village funds lent out 950,000 baht in the first year, so we choose $\underline{s}_v^{mb}$ so that we would have hypothetically predicted an additional 950,000 baht of borrowing in each village in the pre-intervention data:⁴³

$$\frac{1}{N} \sum_{n=1}^N \left\{ E[B_{n,t,v}|L_{n,t}, Y_{n,t}; \underline{s}_v^{mb}] - E[B_{n,t,v}|L_{n,t}, Y_{n,t}; \underline{s}] \right\} = \frac{950,000}{\# \text{ HHs in village}_v}$$

Here N represents the number of surveyed households in the pre-intervention data.

The resulting $\underline{s}_v^{mb}$ values average -0.28 across the villages, with a standard deviation of 0.14 , a minimum of -0.91 and a maximum of -0.09 . Hence, for most villages, the post-program ability to borrow is substantial relative to the baseline ($\underline{s} = -0.08$), averaging about one-fifth of permanent income after the introduction of the program.⁴⁴

5.2 Predictive Power

Using the calibrated values of borrowing limits, we evaluate the model's predictions for 2002 and 2003 (i.e., $t = 6$ and 7) on five dimensions: log consumption, probability of investing, log investment levels, default probability, and income growth. Using the observed liquidity ($L_{n,5}$) and income data ($Y_{n,5}$) for year five (i.e., 2001), the last pre-intervention year, we draw series of $U_{n,t}$, $N_{n,t}$, $i_{n,t}^*$, and measurement error shocks from the estimated distributions, and simulate the model for 2002 and 2003. We do this 500 times, and combine the data with the actual pre-intervention data, in order to create 500 artificial datasets.

⁴²Microfinance is often viewed as a lending technology innovation which is consistent with the reduction in $\underline{s}$. An alternative would be to model the expansion of credit through a decrease in the interest rate on borrowing, but recall that we did not measure a decline in short-term interest rates in response to the program.

⁴³Since 1999 is the base year used, the 950,000 baht is deflated to 1999 values. Predicted results are similar if we use the one million baht which might have been predicted ex ante.

⁴⁴These large changes are in line with the size of the intervention, however. In the smallest village, the ratio of program funds to village income in 2001 is 0.42 . If half the households borrow, this would account for the 0.83 drop in $\underline{s}$.

We then ask whether reduced-form regressions would produce similar impact estimates using simulated data as they would using the actual post-intervention data, even though statistically the model is rejected. We do not have a theory of actual borrowing from the village fund, so rather than using a first-stage equation for village fund credit, we put

$\frac{950,000}{\# \text{ HHs in village}_v}$, the average injection per household, directly into the outcome equations in place of predicted village fund credit. The following reduced form regressions are then:

$$\begin{aligned} C_{n,t} &= \sum_{j=6,7} \alpha_{C,j} \frac{950,000}{\# \text{ HHs in village}_v} I_{t=j} + \theta_{C,t} + e_{C,n,t} \\ D_{n,t} &= \sum_{j=6,7} \alpha_{D,j} \frac{950,000}{\# \text{ HHs in village}_v} I_{t=j} + \theta_{D,t} + e_{D,n,t} \\ I_{n,t} &= \sum_{j=6,7} \alpha_{I,j} \frac{950,000}{\# \text{ HHs in village}_v} I_{t=j} + \theta_{I,t} + e_{I,n,t} \\ DEF_{n,t} &= \sum_{j=6,7} \alpha_{DEF,j} \frac{950,000}{\# \text{ HHs in village}_v} I_{t=j} + \theta_{DEF,t} + e_{DEF,n,t} \\ \ln(Y_{n,t}/Y_{n,t-1}) &= \sum_{j=6,7} \alpha_{\Delta \ln Y,j} \frac{950,000}{\# \text{ HHs in village}_v} I_{t=j} + \theta_{\Delta \ln Y,t} + e_{\Delta \ln Y,n,t} \end{aligned}$$

Here $\hat{\alpha}_{C,j}$, $\hat{\alpha}_{D,j}$, $\hat{\alpha}_{I,j}$, $\hat{\alpha}_{DEF,j}$, and $\hat{\alpha}_{\Delta \ln Y,j}$ would be estimates of the year j impact of the program on consumption, investment probability, average investment, default probability, and log income growth, respectively. Beyond replacing village fund credit ($VFCR_{n,t}$) and

$\frac{950,000}{\# \text{ HHs in village}_v}$, its first-stage regression with $\frac{950,000}{\# \text{ HHs in village}_v}$, the above equations differ from the motivating regressions, equation (1), in two other ways. First, impact coefficients $\alpha_{Z,j}$ are now vary by year j . Second, the regressions above omit the household level controls and household fixed-effects, but recall Section 4.1.1, where we filtered the data of variation correlated with household level demographic data. We also filtered year-to-year variation out of the pre-program data, so the year fixed effects will be zero for the pre-program years. For the post-program years, however, the year fixed-effects will capture the aggregate effect of the program as well as any cyclical component not filtered out of the the actual post-program data. We run these regressions on both the simulated and actual data and compare the estimates and standard errors.

Table V compares the regression results of the model to the data, and shows that the model does generally quite well in replicating the results, particularly for consumption, investment probability, and investment.

The top panel presents the estimates from the actual data. These regressions yield the surprisingly high, and highly significant, estimates for consumption of 1.39 and 0.90 in the first year and second year, respectively. The estimate on investment probability is significant and positive, but only in the first year. For a village, with the average village fund credit per household of 9600, the point estimate of 6.3e-6 would translate into an increase in investment probability of six percentage points. Nonetheless, and perhaps surprising in a world without lumpy investment, the regressions find no significant impact on investment, and very large standard errors on the estimates. The impact effects on default are significant, but negative in the first year and positive in the second year reflecting transitional dynamics. Finally, the impact of the program on log income growth is positive and significant, but only in the second year. Again, given the average village fund credit per household, this coefficient would translate into a ten percentage point higher growth rate in the second year.

The second panel of Table V presents the regressions using the simulated data. The first row shows the average (across 500 samples) estimated coefficient and the second row shows the

average standard error on these estimates. The main point is that the estimates in the data are typical of the estimates the model produces for consumption, investment probability, and investment. In particular, the model yields a large and significant estimate of the coefficient on consumption that is close to one in the first year, and a smaller though still large estimate in the second year. The standard errors are also quite similar to what is observed. The model also finds a comparably sized significant coefficient on the investment probabilities, although its average coefficients are more similar in both the first and second years, whereas the data show a steep drop off in the magnitude and significance after the first year.

The model's predictions for default and income volatility growth are less aligned with the data. For default, both the model and data show a marked and significant decrease in default in the first year, though the model's is much larger. While the data show a significant increase in default in the second year, the model produces no effect.⁴⁵ The data also shows a significant increase in income growth in the second year, whereas regressions from the model measure no impact on income growth. Perhaps, both of these shortcomings are results of the model's inability to fully capture year to year fluctuations in the volatility of the income growth process in the estimation.

The final panel shows formally that the estimates from the model are statistically similar to those in the data. It shows the significance level of a Chow test on the combined sample between the actual post-program data and the simulated post-program data (from all simulations), where the null is no structural break between the actual and simulated data. Using a five (or even ten) percent level of significance, the Chow test would not detect a structural break in any of the regressions.

One further note is that while the impact coefficients in the data are quite similar to those in the simulated structural model, they differ substantially from what would be predicted using reduced form regressions. For example, if we added credit ($CR_{n,t}$) as a right-hand side variable in a regression on consumption, a reduced form approach might use the coefficient (say δ_1) on credit to predict the per baht impact of the village fund credit injection. That is,

we might predict a change in consumption of $\delta_1 \frac{950,000}{\# \text{ HHs in village}_v}$. However, in the following regression:

$$C_{n,t} = \delta_1 CR_{n,t} + \delta_2 \frac{950,000}{\# \text{ HHs in village}_v} I_{t=j} + \theta_{c,t} + e_{c,n,t}$$

an F-test does indeed reject that $\delta_1 = \delta_2$. Parallel regressions that replace credit with consumption, investment probability, or default also reject this restriction, and these restrictions are also rejected if credit is replaced with liquidity or income.

In sum, we measure large average effects on consumption and insignificant effects on investment, but the structural model helps us in quantitatively interpreting these impacts. First, these average coefficients mask a great deal of unobserved heterogeneity. Consider Figure 4 which shows the estimated policy function for consumption (normalized by permanent income) c as a function of (normalized) project size i^* and (normalized) liquidity l . Again, the cliff-like drop in consumption running diagonally through the middle of the graph represents the threshold level of liquidity that induces investment. In the simulations,

⁴⁵For the alternative definition of default, where all loans not from relatives with an unstated duration are considered in default, the data actually show a small decrease in the second year.

households in a village are distributed along this graph, and the distribution depends on the observables (Y and L), and stochastic draws of the shocks (i^* and U , since $p = \frac{Y}{U}$).

We have plotted examples of five potential households, all of whom could appear ex ante identical in terms of their observables, Y and L . (i.e., their state) constant, but resembles a leftward shift in the graphed decision (recall Figure 3, panel (b)). A small decrease in $\underline{s}$ can yield qualitatively different responses to the five households labeled. Household (i)'s income is lower than expected, and so would respond to small decrease in $\underline{s}$ by borrowing to the limit and increasing consumption. Household (ii) is a household that had higher than expected income. Without the intervention, the household invests and is not constrained in its consumption. Given the lower $\underline{s}$, it does not borrow, but nevertheless increases its consumption. Given the lower borrowing constraint in the future, it no longer requires as large a bufferstock today. Household (iii), though not investing, will similarly increase consumption without borrowing by reducing its bufferstock given a small decrease in $\underline{s}$. Thus, in terms of consumption, Household (i)–(iii) would increase consumption, and Households (ii) and (iii) would do so without borrowing. If these households were the only households, the model would deliver the surprising result that consumption increases more than credit, but Households (iv) and (v) work against this. Household (iv) is a household in default. A small decrease in $\underline{s}$ would have no affect on its consumption or investment, but simply increase the indebtedness of the household and reduce the amount of credit that would have been defaulted. Finally, Household (v) is perhaps the target household of microcredit rhetoric a small increase in credit would induce the household to invest. But if (as drawn) the household would invest in a sizable project, it would finance this by not only increasing its borrowing but also by reducing its current consumption. One can also see that the effects of changes in $\underline{s}$ are not only heterogeneous, but also nonlinear. For example, if the decrease in $\underline{s}$ were large enough relative to i^* , Household (v) would not only invest but also increase consumption.

Quantitatively, draws from the distributions of i^* and U (together with the empirical distribution of L/Y) determine the scattering of households in each village across Figure 3. The high level of transitory income growth volatility lead to a high variance in U , hence a di use distribution in the L/P dimension (given L/Y). We know that in the baseline distribution the model calibrates that 19 percent of households are in default (like Household (iv)), and an additional 26 percent are hand-to-mouth consumers (like Household (i), though 3 of the 26 percent are investing).⁴⁶ Based on the pre-sample years, the relaxation of $\underline{s}$ would lead to fewer defaulters (12 percent of households) but the same number of hand-to-mouth consumers (26 percent total, 4 percent of which are investing). Hence, the large share of hand-to-mouth consumers, together with the large share (51 percent) of unconstrained households (like Households (ii) and (iii)) who drive down their buffer stocks, explains the big increase in consumption.

Similarly, the low investment probability but sizable average investment levels in the data lead to high estimated mean and variance of the i^* distribution. Given these estimates, most households in the model have very large projects (with a log mean of 6.26), but investment is relatively infrequent (11.6 percent of observations in the model and data). The median investment is 14 percent (22 percent) of annual income in the data (model), so that most investments are relatively small, but these constitute only 4 percent (8 percent) of all investment in the data (model).⁴⁷ In contrast, a few very large i^* investments (e.g., a large

⁴⁶Many bufferstock models (e.g., Aiyagari (1994)) yield very low level of constrained households in equilibrium. Relative to these models, our model has three important differences. First, we allow for default with minimum consumption, which is empirically observed, so the costs of being liquidity constrained are much lower. Second, investment also causes households to be constrained. Third, we are not modeling a stationary, general equilibrium, but estimating parameters in a partial equilibrium model.

truck or a warehouse) have large effects on overall investment levels. For example, the top percentile of investments accounts for 36 percent (24 percent) of all investment in the data (model). Hence, while some households lie close enough to the threshold that changes in $\underline{s}$ induce investment (4 percent of households in the pre-sample years), the vast majority of these investments are small. That is, the density of households resembling Household (v) is low, especially for large investments (high levels of i^*).

Since a lower $\underline{s}$ can never reduce investment, the theoretical effect of increased liquidity on investment levels is clear. It is simply that the samples are too small to measure it. Given enough households, a small amounts of credit available will eventually decide whether a very large investment is made or not, and this will occurs more often the larger the decrease in $\underline{s}$. Indeed, when the 500 samples are pooled together, the pooled estimates of 0.40 (standard error=0.04) for $\gamma_{I,2002}$ is highly significant. The estimate is also sizable. Given the average credit injection per household, this would be an increase in investment of 3800 baht per household (relative to a pre-sample average of 4600 baht/household).

5.3 Normative Analysis

We evaluate the benefits of the Million Baht program by comparing its benefits to a simple liquidity transfer. As our analysis of Figure 4 indicates reductions in $\underline{s}$ (leftward shifts in the policy function from the Million Baht program) are similar to increases in liquidity (rightward shifts in the households from the transfer). Both provide additional liquidity.

The advantage of the Million Baht program is that it provides more than a million baht in potential liquidity ($-(\underline{s}_v^{mb} - \underline{s})P$). That is, (by construction) borrowers choose to increase their credit by roughly a million baht, but non-borrowers also benefit from the increased potential liquidity from the relaxed borrowing constraint in the future. More generally, those that borrow have access to a disproportionate amount of liquidity relative to what they would get if the money were distributed *equally* as transfers.

The disadvantage of the Million Baht program is that it provides this liquidity as credit, and hence there are interest costs which are substantial given $r = 0.054$. A household that receives a transfer of, say, 10,000 baht earns interest on that transfer relative to a household that has access to 10,000 baht in credit, even if it can be borrowed indefinitely.

The relative importance of these two differences depends on household's need for liquidity. Consider again the household in Figure 3. Household (ii) and (iii), who are not locally constrained (i.e., their marginal propensity to consumer is less than one), benefit little from a marginal decrease in $\underline{s}$, since they have no need for it in the current period, and may not need it for quite some time. Households (iv), who is defaulting, is actually hurt by a marginal reduction in $\underline{s}$, since the household will now hold more debt, and be forced to pay more interest next period. On the other hand, Households (i) and (v) benefit greatly from the reduction in $\underline{s}$, since both are locally constrained, in consumption and investment respectively.

A quantitative cost benefit analysis is done by comparing the cost of the program (the reduction in $\underline{s}$) to a transfer program (an increase in I) that is equivalent in terms of providing the same expected level of utility (given $L_{n,t}$ and $Y_{n,t}$ in 2001, just before the program is introduced). That is, we solve the equivalent transfer T_n for each household using the following equation:

⁴⁷ An alternative interpretation of the data is that most households do not have potential projects that are of the relevant scale for microfinance. Households with unrealistically large projects may correspond, in the real world, to households that simply have no potential project in which to invest.

$$E[V(L, P, I^*, \underline{s}_v^{mb})|Y_{n,5,v}, L_{n,5,v}] = E[V(L+T_n, P, I^*, \underline{s})|Y_{n,5,v}, L_{n,5,v}]$$

The *average* equivalent liquidity transfer per household in the sample is just 8200 baht which is about twenty percent less than the 10,100 baht per household that the Million Baht program cost.⁴⁸ Again, this average masks a great deal of heterogeneity across households, even in expectation. Ten percent of households value the program at 19,500 baht or more, while another ten percent value the program at 500 baht or less. 28 percent of households value the program at more than its cost (10,100 baht), but the median equivalent transfer is just 5900 baht. Thus, many households benefit disproportionately from the program because of the increased availability of liquidity, but most benefit much less. Although the Million Baht program is able to offer the typical household more liquidity (e.g., in the median village, $(-\underline{s}_v^{mb} - \underline{s})P = 13,400$ baht for a household with average income, while the average cost per household in that village is 9100 baht), this benefit is swamped by the interest costs to households.

5.4 Alternative Structural Analyses

The structural model allows for several alternative analyses including comparison with reduced form predictions, robustness checks with respect to the return on investment R , estimation using post-intervention data, long run predictions and policy counterfactuals. We briefly summarize the results here, but details are available upon request.

5.4.1 Return on Investment—Our baseline value of R was 0.11. Recall that two alternative calibrations of the return on assets were calculated based on the whether our measure of productive assets included uncultivated or community use land ($R = 0.08$) or the value the plot of land containing the home ($R = 0.04$). We redid both the estimation and simulation using these alternative values. For $R = 0.08$, the estimates were quite similar; only a higher β (0.94), a lower r (0.032); and a lower risk aversion (1.12) were statistically different than the baseline. The model had even more difficulty matching income growth and volatility, so that the overall fit was substantially worse (J-statistic=200 vs. 113 in the baseline). The simulation regression estimates were nearly identical. For the low value of $R = 0.04$, the estimation required that the return on liquidity be substantially lower than in the data ($r = 0.018$), and that β be substantially higher (0.97) than typical for bufferstock models. The fit was also substantially worse (J-statistic=324). Finally, the regression estimates on the simulated data were qualitatively similar but smaller (e.g., a consumption coefficient of 0.68 in the first year.) Indeed, only the reduction of default in the first year was statistically significant at a 0.05 percent level.

5.4.2 Estimation Using Ex Post Data—In this analysis, rather than use the post-intervention data to test the model using calibrated borrowing constraints, we use it to estimate the new borrowing constraints and better identify the other parameters in the model. We proceed by specifying a reasonably flexible but parametric function for $\underline{s}_{mb}$ in the post-program years:

$$\underline{s}_{mb,v} = \underline{s}_1 + \underline{s}_2 \left(\frac{1}{\# \text{ HHs in village}_v} \right)^{\underline{s}_3}$$

⁴⁸This includes only the seed fund, and omits any administrative or monitoring costs of the village banks.

where $\underline{s}_1$, $\underline{s}_2$, and $\underline{s}_3$ are the parameters of interest.⁴⁹ The moments for the post-program years cover: interest on savings and borrowing (two moments); income growth (two) and income growth volatility (three); consumption (two), investment probability (two), investment (two), and their interactions with measured income and liquidity ratios (twelve); and default (two). All but the interest moments are year-specific, and the only use of pre-program data is to construct the four income growth moments that require income in 2001. In total, the estimation now includes 27 moments and 14 parameters.

The estimated results from the sample are strikingly similar to the baseline estimates from the pre-program sample and the calibration from the post-program sample, all with two standard deviation bands.⁵⁰ The resulting estimates are $\hat{s}_1 = -0.07$, $\hat{s}_2 = -47$, and $\hat{s}_3 = -1.20$. The model fit is comparable to the baseline, performing well along the same dimensions and not well at all along the same dimensions. Finally, the average, standard deviation, minimum and maximum of $s_{mb,v}$ implied by the estimates are -0.26 (-0.28 in baseline calibration), -0.14 (-0.14), -0.86 (-0.91), and -0.07 (-0.09) respectively. The correlation between the two is very close to one by construction, since both increase monotonically with village size. That is, the estimated $\underline{s}_{mb,v}$ are quite similar to the calibrated values. The fact that the estimates and calibrated values are quite close indicates that cross-sectionally the simulated predictions of the model on average approximate a best fit to the variation in the actual data.

5.4.3 Long Run Predictions—The differences between $\hat{\alpha}_{Z,j}$ estimates in the first and second year (i.e., $j = 1, 2$) of the program indicate that impacts are time-varying, since there are transitional dynamics as households approach desired bufferstocks. The structural model allows for simulation and longer run horizon estimates of impact. We therefore simulate datasets that include five additional years of data and run the analogous regressions. Seven years out, none of the $\hat{\alpha}_{Z,7}$ estimates are statistically significant on average. While the average point estimates are quite small for investment probability (0.23), investment (0.10), and default probability (0.01) relative to the first year, the average $\hat{\alpha}_{Z,7}$ for consumption remains substantial (0.58) and close to the estimate in the second year (0.73). In the model, the impacts on consumption fall somewhat after the first year, but there remains a substantial persistent effect. Still, alternative regression estimates that simply measure a single (common for all post-program years j) coefficient α_Z do not capture any statistically significant impact on consumption when seven years of long run data are used. This shows the importance of considering the potential time-varying nature of impacts in evaluation.

5.4.4 Policy Counterfactual—From the perspective of policymakers, the Million Baht Village Fund Program may appear problematic along two fronts. Its most discernible impacts are on consumption rather than investment, and it appears less cost-effective than a simple transfer mainly because funds may simply go to prevent default and the increased borrowing limit actually hurts defaulting households. An alternative policy that one might attempt to implement would be to only allow borrowing for investment. We would assume that the village can observe investment, but since money is fungible, it would be unclear whether these investments would have been undertaken even without the loans, in which case the loans are really consumption loans. Since defaulting households cannot undertake

⁴⁹If all households borrowed every period and had identical permanent income, then the extra borrowing per household ($950,000/\#$ HHs in village_{*v*}) would translate into borrowing constraints with $\underline{s}_1 = \underline{s}$ (the pre-intervention borrowing constraint), $\underline{s}_2 = \frac{950,000}{P}$, and $\underline{s}_3 = 1$.

⁵⁰For comparison, the point estimates of the full-sample (baseline) estimation are $\hat{r} = 0.060$ (0.054), $\hat{\sigma}_N = 0.35$ (0.31), $\hat{\sigma}_U = 0.51$ (0.42), $\hat{\sigma}_E = 0.28$ (0.15), $\hat{G} = 1.052$ (1.047), $\hat{c} = 0.53$ (0.52), $\hat{\beta} = 0.926$ (0.926), $\hat{\rho} = 1.21$ (1.20), $\hat{\mu}_i = 1.24$ (1.47), $\hat{\sigma}_i = 2.56$ (2.50), and $\hat{s} = -0.12$ (-0.08).

investments, it would prevent households in default from borrowing. Nevertheless, such a policy would also eliminate households like Household (i) in Figure 4 from borrowing.

The ability to model policy counterfactuals is another strength of a structural model. In a model with this particular policy, households face the constraint $\underline{s}_v^{mb,alternative}$ in any period in which they decide to invest, while facing the baseline $\underline{s}$ if they decide not to invest. The default threshold is also moved to $\underline{s}_v^{mb,alternative}$, however, to prevent households from investing and borrowing in one period, and then purposely not investing in the next period in order to default. Under this policy, the new borrowing constraints are even lower with wider variation (a maximum, minimum, and mean of -0.16 , -4.78 , and -0.67 , respectively, vs. -0.09 , -0.91 and -0.28 for the actual policy) but only for those who borrow.

The policy increases both the impact on consumption and increase the impact on investment. Pooling all 500 simulated samples yields a significant estimate for consumption that is similar to the actual million baht intervention (1.40 vs. 1.38 in the first year). It also yields a much larger and significant estimate for investment levels (0.62 in the first year), which is expected since the borrowing constraints of investors are much lower under this policy. Naturally, this policy offers less flexibility for constrained households who would rather not invest, but the relatively larger benefits to defaulters and investors help outweigh this loss. There is much more variation in the benefits across households (e.g., the standard deviation of the equivalent transfer is $14,000$ baht in this counterfactual vs. $11,000$ in the baseline policy), but the average equivalent transfer is actually lower (7500 vs. 8200).

6 Conclusions

We have developed a model of bufferstock saving and indivisible investment, and used it to evaluate the impacts of the Million Baht program as a quasi-experiment. The correct prediction of consumption increasing more than one for one with the credit injection is a “smoking gun” for the existence of credit constraints, and is strong support for the importance of bufferstock savings behavior. Nevertheless, the microfinance intervention appears to be less cost effective on average than a simpler transfer program because it saddles households with interest payments. This masks considerable heterogeneity, however, including some households that gain substantially. Finally, we have emphasized the relative strengths of a quasi-experiment, a structural model, and reduced form regressions.

One limitation of the model is that although project size is stochastic, the quality of investments, modeled through R , is assumed constant across projects and households. In the data, R varies substantially across households. Heterogeneity in project quality may be an important dimension for analysis, especially since microfinance may change the composition of project quality. Ongoing research by Banerjee, Breza, and Townsend find that high return households do borrow more from the funds, but they also invest less often, which indicates that the data may call for a deeper model of heterogeneity and, related, a less stylized model of the process for projects sizes. Potential projects may not arrive each year, they may be less transient (which allows for important anticipatory savings behavior as in Buera, 2008), or households might hold multiple projects ordered by their profitability. Such extensions might help explain the investment probability results in the second year of the program: a positive impact in the model but no impact in the data.

Related, the analysis has also been purely partial equilibrium analysis of household behavior. In a large scale intervention, one might suspect that general equilibrium effects on income, wage rates, rates of return to investment, and interest rates on liquidity may be important (see Kaboski and Townsend, 2009). Finally, we did not consider the potential

interactions between villagers or between villages, nor were the intermediation mechanism or default contracting explicitly modeled. These are all avenues for future research.

References

1. Aaronsen, D.; Agarwal, S.; French, E. The Consumption Response to Minimum Wage Increases. FRB Chicago working paper; 2007-23, May; 2008.
2. Ahlin C, Townsend R. Using Repayment Data to Test Across Models of Joint Liability Lending. mimeo. 2000
3. Aiyagari SR. Uninsured Idiosyncratic Risk and Aggregate Saving. Quarterly Journal of Economics. 1994 August.109 :659–684.
4. Alvarez F, Stokey N. Dynamic Programming with Homogeneous Functions. Journal of Economic Theory. 1998; 82:167–189.
5. Aportela, F. mimeo. Massachusetts Institute of Technology; 1998. Effect of Financial Access on Savings by Low-Income People.
6. Arghiros, D. Democracy, Development and Decentralization in Provincial Thai-land. (Democracy in Asia, 8). Curzon Press; Richmond, Surrey, UK: 2001. p. 308
7. Banerjee A, Newman A. Occupational Choice and the Process of Development. Journal of Political Economy. 1993 April.:274–98.
8. Banerjee, A.; Dufllo, E. mimeo. Massachusetts Institute of Technology; 2002 May. Do Firms Want to Borrow? Testing Credit Constraints Using a Direct Lending Program.
9. Banerjee, A.; Dufllo, E.; Glennerster, R.; Kinnan, C. mimeo. 2009 May. The miracle of microfinance? Evidence from a randomized evaluation.
10. Blossfeld, H.; Shavit, Y. Persistent Inequality: Changing Educational Attainment in Thirteen Countries. Boulder, Co: Westview Press; 1993.
11. Buera, FJ. mimeo. Northwestern University; 2008. Persistency of Poverty, Financial Frictions, and Entrepreneurship. 2008
12. Burgess R, Pande R. Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment. American Economic Review. 2005 June.95 :780–795.
13. Carroll CD. Buffer Stock Saving and the Lifecycle Permanent Income Hypothesis. Quarterly Journal of Economics. 1997; 107(1)
14. Carroll, CD. Manuscript. The Johns Hopkins University; 2004. Buffer Stock Saving: Some Theory.
15. Chiappori, P.; Schulhofer-Wohl, S.; Samphantharak, K.; Townsend, R. manuscript. 2008 January. On Homogeneity and Risk Sharing in Thai Villages.
16. Carroll CD, Samwick AA. The Nature of Precautionary Wealth. Journal of Monetary Economics. 1997; 40(1)
17. Coleman, B. Asian Development Bank ERD Working Paper Series No. 9. 2001. Microfinance in Northeast Thailand: Who Benefits and How Much?.
18. Coleman B. The Impact of Group Lending in Northeast Thailand. Journal of Development Economics. 1999 October.:105–141.
19. Deaton A, Paxson C. Growth and Saving Among Individuals and Households. Review of Economics and Statistics. 2000 May.82:212–225.
20. Galor O, Zeira J. Income Distribution and Macroeconomics. Review of Economic Studies. 1993; 60 (1):35–52.
21. Gertler, P.; Levine, D.; Moretti, E. Center for International and Development Economics Research Paper C03-129. 2003 February. Do Microfinance Programs Help Insure Families Against Illness?.
22. Gine, X.; Townsend, R. Journal of Development Economics. Vol. 74. Elsevier; 2004 August. Evaluation of Financial Liberalization: A General Equilibrium Model with Constrained Occupation Choice; p. 269-307.
23. Gourinchas PO, Parker J. Consumption over the Life Cycle. Econometrica. 2002 January.71 :47–91.

24. Heckman, JJ.; Urzua, S.; Vytlacil, E. Understanding Instrumental Variables in Models with Essential Heterogeneity. Stanford University; Stanford, CA: 2004.
<http://www.stanford.edu/~vytlacil/>
25. Kaboski, J.; Townsend, R. The Impacts of Credit on Village Economies. mimeo. 2008.
<http://kaboski.econ.ohio-state.edu/impactofcredit.pdf>
26. Kaboski, J.; Townsend, R. mimeo. University of Chicago; 1998. Borrowing and Lending in Semi-Urban and Rural Thailand; p. 1998
27. Kaboski, J.; Townsend, R. Journal of the European Economic Association. Vol. 3. MIT Press; 2005. Policies and Impact: An Analysis of Village Microfinance Institutions; p. 1-50.p. 03
28. Karlan, D.; Zinman, J. Expanding Credit Access: Using Randomized Supply Decisions To Estimate the Impacts. Yale University Working Paper 68. 2009.
<http://www.econ.yale.edu/ddp/ddp50/ddp0068.pdf>
29. Keane M, Wolpin K. The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence. Review of Economics and Statistics. 1994 November.76:648–72.
30. Keane M, Wolpin K. The Career Decisions of Young Men. Journal of Political Economy. 1997 June.105:473–522.
31. Keane M, Wolpin K. The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment. International Economic Review. 2001 November.42:1051–1103.
32. Lloyd-Ellis H, Bernhardt D. Enterprise, Inequality, and Economic Development. Review of Economic Studies. 2000 January.:147–68.
33. Lise, J.; Seitz, S.; Smith, J. mimeo. 2005 April. Equilibrium Policy Experiments and the Evaluation of Social Programs.
34. Lise, J.; Seitz, S.; Smith, J. mimeo. 2005 April. Evaluating Search and Matching Models Using Experimental Data.
35. Lochner, L.; Monge-Naranjo, A. NBER working paper W13912. 2008 April. The Nature of Credit Constraints and Human Capital.
36. Low H, Meghir C, Pistaferri L. Wage Risk and Employment Risk over the Life Cycle. American Economic Review. forthcoming.
37. Meghir C, Pistaferri L. Income Variance Dynamics and Heterogeneity. Econometrica. 2004 January.72:1–32.
38. Owen, A.; Weil, D. NBER Working Paper 6070. 1997 June. Intergenerational Earnings Mobility, Inequality, and Growth.
39. Pitt M, Khandker S. The Impact of Group-Based Credit Programs on Poor Households in Bangladesh: Does the Gender of Participants Matter? Journal of Political Economy. 1998 October.:958–96.
40. Pitt M, Khandker S, Chowdury O, Millimet D. Credit Programs for the Poor and the Health Status of Children in Rural Bangladesh. International Economic Review. 2003 February.44 :87–118.
41. Puentes, E. Thai Project Note. 2008.
42. Puginier, O. Doctoral Dissertatoin. Humboldt University; 2001. Hill tribes struggling for a land deal: Participatory land use planning in northern Thailand amid controversial policies.
43. Rosenzweig M, Wolpin K. Natural ‘Natural Experiments’ in Economics. Journal of Economic Literature. 2000 December.38 :827–874.
44. Samphantharak, K. PhD thesis. University of Chicago; Chicago: 2003. Internal Capital Markets in Business Groups.
45. Samphantharak, K.; Townsend, R. Households as Corporate Firms: An Analysis of Household Finance Using Integrated Constructing Financial Statements from Integrated Household Surveys and Corporate Financial Accounting. Cambridge University Press; 2009.
46. Todd P, Wolpin K. Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility. American Economic Review. 2006 December.96 :1384–1417.

47. Townsend, R., principal investigator; Paulson, A.; Sakuntasathien, S.; Lee, T.J.; Binford, M. Risk, Insurance and the Family and NSF grants. The University of Chicago; 1997. Questionnaire Design and Data Collection for NICHD Grant.
48. World Bank. OED Precip. 1996 February. KUPEDS: Indonesia's Model Small Credit Program; p. 4
49. Wright, M. mimeo. 2002 October. Investment and Growth with Limited Commitment.
50. Zeldes S. Consumption and Liquidity Constraints: An Empirical Investigation. Journal of Political Economy. 1989 March.97 :305–46.

Figure 1.
Number of Households per Villages, Four Provinces, Thailand

Figure 2.
Value Function vs. Liquidity Ratio & Project Size

Figure 3.
Consumption Policy for Fixed i^* , Baseline and Reduced Borrowing Constraint

Figure 4.
Consumption Policy as a Function of Liquidity and Project Size

TABLE I
PRE-EXISTING TRENDS BY INVERSE VILLAGE SIZE

Outcome Variable, Z	$\alpha_{2,Z}, \text{"trends"}$	Outcome Variable, Z	$\alpha_{2,Z}, \text{"trends"}$
Village fund short-term credit	0.01 (0.02)	Business investment	0.03 -0.19
Total short-term credit	0.09 (0.15)	Agricultural investment	0.04 (0.13)
BAAC credit	0.04 (0.10)	Investment probability	5.1e-5 (2.1e-4)
Commercial bank credit	0.05 ^b (0.03)	Fertilizer expenditures	-0.04 (0.06)
Agricultural credit	-0.07 (0.04)	Total wages paid to laborers	0.19 (0.12)
Business credit	0.04 (0.10)	Consumption	0.19 (0.27)
Fertilizer credit	0.14 (0.10)	Nondurable Consumption	0.09 (0.21)
Consumption credit	0.05 (0.08)	Grain consumption	-0.03 (0.04)
Short-term interest rate	-1.6e-7 (5.3e-7)	Milk consumption	0 (0.01)
Probability in default	-9.8e-7 (1.3e-6)	Meat consumption	0.01 (0.01)
Credit in default	-1.1e-6 (1.5e-6)	Alcohol cons. in the house	-0.01 (0.01)
Informal credit	0.00 (0.09)	Alcohol cons. outside of the house	-0.01 (0.01)
Income growth	-7.2e-6 (4.5e-6)	Fuel consumption	-0.02 (0.03)
Fraction of net income from business	2.0e-7 (3.9e-7)	Tobacco consumption	-0.01 (0.01)
Fraction of income from wages	9.1e-7 (6.1e-7)	Education Expenditures	0.03 (0.02)
Fraction of income from rice	1.0e-6 ^a (5.6e-7)	Ceremony expenditures	-0.01 (0.03)
Fraction of income from other crops	7.9e-8 (4.1e-7)	Housing repair expenditures	0.06 (0.14)
Fraction of income from livestock	6.2e-8 (3.8e-7)	Vehicle repair expenditures	0.00 (0.03)
Log Asset growth	6.0e-7 (2.9e-6)	Clothes expenditures	0.00 (0.01)
Number of new businesses	4.9e-7 (1.1e-6)	Meals expenditures away from home	0.00 (0.01)

^a significant at a 10% level

TABLE II
SUMMARY STATISTICS OF PRE-INTERVENTION HOUSEHOLD DATA

Variable	Obs	Mean	Std. Dev.	Min	Median	Max
<i>Primary Variables:</i>						
Non-Interest Household Income ^a	3575	87200	202000	500	50300	6255500
Log Growth of Income ^a	2860	0.04	0.98	-4.94	0.01	10.28
Household Consumption ^a	3575	75200	93000	750	49800	1370300
Dummy Variable for Agr/Business Investment	3575	0.12	0.34	0	0	1
Value of Agr./Business Investment ^a	3575	4760	30200	0	0	715700
Dummy Variable for Short-Term Default	2860	0.194	0.395	0	0	1
Short-Term Credit ^a	3575	17900	51100	0	0	1021000
Interest Paid ^a	3575	1300	3900	0	0	108400
Liquid Savings ^a	2860	25000	132000	0	5100	4701600
Interest Earned ^a	3575	700	7200	0	0	18000
Number of Households in Village	3575	166	295	21	110	3194
<i>Regressors for Demographic/Cyclical Variation:</i>						
Number of Male Adults	3575	1.46	0.9	0	1	7
Number of Female Adults	3575	1.56	0.75	0	1	6
Number of Children	3575	1.59	1.21	0	1	9
Dummy Variable for Male Head of Household	3575	0.74	0.44	0	1	1
Years of Education of Head of Household	3575	6	3	0	7	15
Age of Head of Household	3575	41	15	22	40	84

^a All values are in baht deflated to 1999. The 1999 PPP conversion rate is 31.6 baht/dollar.

PARAMETER ESTIMATES AND MODEL FIT

TABLE III

Parameter Estimates		Pre-Intervention Averages			
Parameter	Estimate	Std. Err.	Variable	Data	Model
Borrowing/savings interest rate - r	0.054	0.003	C_t	75,200	75,800
Deviation of log permanent income shock - σ_N	0.31	0.11	D_t	0.116	0.116
Deviation of log transitory income shock - σ_U	0.42	0.07	I_t	4600	4600
Deviation of log measurement error shock - σ_E	0.15	0.09	DEF_t	0.194	0.189
Exogenous income growth - G	1.047	0.006	$\ln(Y_{t+1}/Y_t)$	0.044	0.049
Minimum consumption - $\bar{c}$	0.52	0.01			
Discount factor - β	0.926	0.006	Test for Overidentifying Restrictions		
Intertemporal elasticity - ρ	1.20	0.01		Actual Value	0.05% Value
Mean log project size - μ_i	1.47	0.09			
Deviation of log project size - σ_i	6.26	0.72			
Borrowing limit - $\bar{s}$	-0.08	0.03	J-Statistic	113.5	12.6

TABLE IV
IDENTIFICATION - PARTIAL DERIVATIVES OF MOMENTS WITH RESPECT TO PARAMETERS

Parameters											
	r	σ_N	σ_U	σ_E	G	ξ	β	ρ	μ_i	σ_i	$\bar{s}$
Savings interest - ε_S	-6.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Credit interest - ε_{CR}	-10.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Income growth - ε_g	-0.6	2.5	-0.5	-0.8	-4.4	-0.2	1.5	0.2	-1.8	-1.5	-1.6
1-year variance - $\varepsilon_{V,1}$	1.3	-0.1	1.1	-0.8	0.3	2.0	-2.4	-1.1	-2.4	0.8	-0.4
2-year variance - $\varepsilon_{V,2}$	1.0	0.3	-1.1	-1.4	-0.2	0.6	-2.3	1.0	-1.0	-0.6	0.7
3-year variance - $\varepsilon_{V,3}$	0.5	-0.3	0.8	-2.1	0.8	2.2	-0.3	-2.3	-2.4	-2.5	-0.1
Consumption - ε_C	1.4	0.8	0.1	-0.3	-2.0	-0.5	13.3	-7.1	0.0	0.0	0.7
Income covariance - $\varepsilon_C^* \ln Y$	-3.7	-2.2	-0.2	0.9	5.3	1.4	-35.0	18.6	0.0	0.0	-1.9
Liquidity covariance - $\varepsilon_C^* L/Y$	-0.1	-0.1	0.0	0.0	0.2	0.1	-1.4	0.8	0.0	0.0	-0.1
Investment probability - ε_P	51.5	18.7	-0.8	-0.2	-40.0	-0.7	-16.1	1.4	-0.5	0.1	0.2
Income covariance - $\varepsilon_P^* \ln Y$	-155.2	-55.9	2.5	0.5	120.0	1.9	47.9	-4.1	1.3	-0.3	-0.6
Liquidity covariance - $\varepsilon_P^* L/Y$	-23.2	-11.0	2.1	0.0	18.1	0.4	9.4	-0.6	0.3	-0.1	-0.1
Investment level - ε_I	28.0	10.0	0.1	-0.2	-22.1	-0.3	-8.6	0.6	-0.7	0.1	0.1
Income covariance - $\varepsilon_I^* \ln Y$	-80.0	-28.5	-0.2	0.6	63.1	0.8	24.5	-1.7	1.9	-0.4	-0.2
Liquidity covariance - $\varepsilon_I^* L/Y$	-9.9	-2.8	0.1	0.5	8.2	0.1	4.3	-1.6	0.1	0.0	0.0
Default - ε_{DEF}	0.0	0.0	-1.6	0.2	0.0	-3.6	0.0	0.0	0.0	0.0	-3.6
Moments											

REDUCED FORM REGRESSION ESTIMATES: ACTUAL DATA VS. “MILLION BAHT” SIMULATED DATA

TABLE V

Actual Data	Consumption		Investment Probability		Investment		Default Probability		Income Growth	
	$\gamma_{C,2002}$	$\gamma_{C,2003}$	$\gamma_{D,2002}$	$\gamma_{D,2003}$	$\gamma_{I,2002}$	$\gamma_{I,2003}$	$\gamma_{DEF,2002}$	$\gamma_{DEF,2003}$	$\gamma_{INR,2002}$	$\gamma_{INR,2003}$
“Impact” Coefficient ^a	1.39	0.90	6.3e-6	-0.2e-6	-0.04	-0.17	-5.0e-6	6.4e-6	-9.4e-6	12.6e-6
Standard Error	0.39	0.39	2.4e-6	2.4e-6	0.19	0.19	2.4e-6	2.4e-6	6.1e-6	6.1e-6
Simulated Data										
Average “Impact” Coefficient ^a	1.10	0.73	5.6e-6	3.6e-6	0.41	0.35	-9.0e-6	-0.2e-6	0.3e-6	0.3e-6
Average Standard Error	0.48	0.48	2.5e-6	2.5e-6	0.23	0.23	2.3e-6	2.3e-6	5.9e-6	5.9e-6
Chow Test Significance Level ^b	0.55		0.51		0.99		0.27		0.30	

^aThe impact coefficient is the coefficient on 950,000/number of households in the village interacted with a year dummy, the credit injection per household.

^bThis is the significance level of a Chow test on the actual post-intervention data and the pooled simulated data, where the null hypothesis is no structural break in the impact coefficients.

Bold face represents significance at a 5 percent level.