Testing a Structural Model of Credit Constraints Using a Large-Scale Quasi-Experimental Microfinance Initiative

Joseph P. Kaboski and Robert M. Townsend[*,†]

April 3, 2007

Abstract

This paper develops a structural model of credit constrained household consumption, indivisible investment, and savings behavior, and tests the model using a major government microfinance program as an exogenous quasi-experimental injection of credit. After estimation of parameters using pre-program data, the estimated model is evaluated using the Thai Million Baht Village Fund Program. Simulated predictions from the model mirror actual data in reproducing a greater increase in consumption than credit, which is interpreted as evidence of credit contraints. A cost-benefit analysis using the model indicates that the program costs just 66 percent of a transfer program providing an equivalent benefit.

[*] This research is funded by NICHD grant R03 HD04776801. We would like to thank Sombat Sakuntasathien for collaboration and making the data acquisition possible. We would like to thank Aleena Adam, Francisco Buera, Xavier Gine, Donghoon Lee, Audrey Light, Masao Ogaki and Bruce Weinberg for helpful comments. We have also benefited from comments received on presentations of versions of this work at IUPUI, NIH, Ohio State, UW Milwaukee, NYU, and the NEUDC 2006 and 2006 Econometric Society winter meetings. Bin Yu, Taehyun Ahn, and Jungick Lee provided excellent research assistance on this project.

[†] Email: Kaboski <kaboski.1@osu.edu>, Townsend <rtownsen@uchicago.edu>

1 Introduction

A large literature has explored the importance of credit constraints, both their theoretical relevance to the consumption, saving and investment decisions of individuals, and the evidence for them in the real world. In theory, it is generally assumed that credit constraints are of greater importance in developing countries and among poorer populations (who have less assets to use as collateral), and many papers have posited that credit constraints act as a major obstacle to economic development and growth.[1] These theories provide strong motivation for empirical tests of credit constrained behavior.

Most of the empirical evidence for credit constraints can be divided into two types. The first type of work tests well-defined, often structural, models, but uses data where the exogeneity and identifying assumptions are *a priori* doubtful. For example, the cross-sectional correlation between investment decisions and some measure of current liquidity (such as income or wealth) is cited as evidence for credit constraints based on a model in which the distribution of current liquidity is assumed to be exogenous to other factors affecting investment.[2] Papers of the second type apply plausibly exogenous variation from a natural, policy, or controlled experiment, but lack the structure of a model to aid in identifying credit constraints and extrapolating results beyond the experiment itself. For example, the high estimated wage gains from instrumental variables methods have been cited as evidence for credit constraints in schooling investment.[3]

This paper follows the first type of research in testing a fully-specified structural model,

[1] See, for example, Banerjee and Newman (1993), Galor and Zeira (1993), Ghatak and Jiang (1999), Gine and Townsend (2001), Lloyd-Ellis and Bernhardt (2001), and Owen and Weil (1997), all of which incorporate indivisible investments.

[2] Related approaches that assume cross-sectional exogeneity have been used to test for credit constraints in the choice of consumption (e.g., Townsend, 1995, Zeldes, 1989, Attanasio and Davis, 1996), investment (e.g., Alem and Townsend, 2002, Sampantharak, 2003, Fazzari et al, 1988), entreprenuership (e.g., Townsend and Paulson, 2002, Evans and Jovanovic, 1989, and education (e.g., Blossfeld and Shavit,1993, Kane, 1994, Cameron and Heckman, 1998, Cameron and Heckman, 2000)

[3] Card (1999) and Krueger (2003) are two examples. Carneiro and Heckman (2001) and Cameron and Taber (2002) critique the interpretation of these results.

but follows the second line of research in using exogenous variation from a policy experiment to evaluate the model's predictions.[4] That is, we develop a model with strong theoretical predictions for household behavior in the presence of potential credit constraints, including predictions for the impact of counterfactual credit interventions on behavior. The most important of these predictions is that consumption increases on average more than one for one with credit, when the credit constraint is relaxed. We estimate the model using non-experimental panel data from rural and semi-urban households in Thailand, but then evaluate the model by comparing the actual data from a quasi-experimental government policy to the model's predictions for this policy.

The model we develop is based on the standard buffer stock model for savings behavior under income uncertainty studied by many authors[5]. Our model follows Aiyagari (1994) and Deaton (1991) in that households' levels of liquid assets are constrained below; they cannot borrow more than some multiple of their permanent income. Though their discount rate exceeds the rate of return on liquid savings, households often maintain savings above this bound as a precautionary buffer against income uncertainty.

In our model, households have the additional option of investing in an asset (or project) that pays a higher rate of return than liquid savings. This asset, however, is, indivisible and illiquid, paying out by increasing permanent income in the future. Furthermore, the size of the indivisible project is stochastic. We introduce investment into the bufferstock model because investment is important in the data. We will ultimately evaluate the same microfinance intervention that we use to test the model, and increasing investment is one of the major goals of microfinance. Investment is modeled as indivisible for two reasons. First, investments in the Thai data (e.g. agricultural equipment or small business start-ups), though small, are lumpy and sporadic.[6] Second, theoretically, credit constraints can

[4]In this way, the papers is related to Banerjee and Duflo (2002) and Todd and Wolpin (2003), who also use experimental or quasi-experimental data to test well-defined theoretical models.

[5]Important contributions include Aiyagari (1994), Carroll (1997), Carroll and Samwick (1997), Deaton (1991), Gourinchas and Parker (2003), and Zeldes (1989).

[6]In the data, nine percent of households invest in any year, but over forty percent of households invest in at least one year. In years when a household invests, the median investment as a fraction of income is

be more serious – creating poverty traps – when high-yield investments are indivisible.[7] Introducing this investment into the model produces strong, novel theoretical predictions. As an example, comparing households with the same permanent income, consumption is no longer monotonically increasing in current liquidity since at some threshold level of liquidity it may pay to reduce consumption and invest in the indivisible asset. Consequently, a household may appear *more* liquidity-constrained when receiving a relatively high transitory income shock than when receiving a lower transitory income shock. As a second example, the model predicts that households who borrow to consume will be more likely to default on credit than households who borrow to invest, since investments will increase future permanent income. Higher permanent income increases access to credit and allows households to repay old loans through the use of new credit.

The model is empirically estimated and tested using data from the Townsend Thai project[8], an ongoing (1997-present) panel survey of a stratified, clustered, random sample of institutions (256 in 2002), households (960 each year) and household businesses (658 in 2002), and village key informants (64, one in each village) in four provinces of Thailand. The model parameters are estimated via GMM using only the first five years of "pre-experiment" data.

The quasi-experiment we use to test the model is the Thai Million Baht[9] Village Fund Program, one of the largest scale government microfinance initiatives of its kind.[10] Started

---

fourteen percent, while the mean is seventy-nine percent.

[7]See, for example, Banerjee and Newman (1993), Lloyd-Ellis and Burnhardt (2000), or Buera (2006). When high yield investments are divisible, households can save there way out of poverty by starting with arbitrarily small investments.

[8]See Townsend, et al. (1997).

[9]In 2001, one million baht was about $25,000 based on exchange rate conversion.

[10]The 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 KUPEDES 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)

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 that lend to village members. The transfers themselves sum to about 1.5 percent of Thai GDP.

We view the program, an initiative of (then newly-elected, now) former Prime Minister Thaksin Shinawatra, as a quasi-experiment in that both the rapid introduction and the design of the program produced exogenous variation in available credit over time and across households. All sixty-four villages in the Townsend Thai panel data received the transfer and began lending between the 2001 and 2002 survey years.[11] More importantly, the total amount of funding given to each village was the same (one million baht) regardless of the size of the village, so village size gives us a potentially exogenous source of variation in treatment, i.e., credit injection, per household. Indeed, in astructural regressions, village size is only significantly related to the variables of interest *after* the introduction of the program (see Kaboski and Townsend, 2006).

We model the program as a reduction of borrowing constraints so as to increase the amount of (expected) credit in the village by one million baht. This entails larger relaxation of credit constraints – and a larger increase in average credit per household – in smaller villages. While the model easily fails statistical tests (i.e., Hansen's J-tests for overidentifying restrictions) in the pre-experimental data, for the program, it does reasonably well in reproducing the aggregate effects on consumption, investment and the probability of investing. That is, impact estimates from reduced form regressions using the model's simulated data are statistically indistinguishable from reduced form regression estimates using the actual data. Most importantly, the model is able to reproduce an impact on consumption that is greater than one million baht, so that consumption increases more than credit does. This is a prediction of a buffer stock model, since increased credit availability lowers households' optimal bufferstock, so that even non-borrowers may increase consumption. An increased supply of credit, even credit that did not need to be repaid, would not produce such results in a neoclassical world.

[11] In 2002, village bank loans amounted to twenty-four percent of loans and ten percent of total credit after amounting to less than one percent of loans and 0.1 percent of total credit in 2001.

After using the policy intervention to test the model's predictions, we use the model to quantify the impact of the policy intervention. A key advantage of a structural model is its potential for well-defined normative policy evaluation, and evaluation of microfinance interventions, especially such a large scale government program, is a matter of great importance. Proponents of microfinance argue that the unique policies of microfinance enable institutions to bring credit and savings services to underdeveloped areas and to people with otherwise insufficient or no access to contemporary financial systems. The hope and claim is that the provision of saving and credit is both effective in fighting poverty and more financially viable than other means. Detractors point to the frequent failure of microfinance institutions and reliance, implicitly or explicitly, on subsidies. Despite the prevalence of microcredit initiatives and the debate, there has been relatively little empirical examination of their impacts. The few efforts to evaluate the impacts of microfinance institutions using astructural methods and plausibly exogenous data have produced mixed or contradictory results.[12] To our knowledge, this is the first structural attempt at such an evaluation. That is, we use reduced form regressions to compare the model's simulated predictions to the data, but our estimate of the program's benefits comes directly from the structural model.

Our evaluation compares the costs of the microfinance program to the costs of a direct transfer program that is equivalent in the sense of providing the same utility benefit. The costs of the microfinance program are 33 percent less than the equivalent direct transfer program. This is because we assume that village fund monies have the effect of lowering borrowing constraints, while the transferred funds do not. In large villages, where the transfer per person was smaller, the cost savings is even greater. The cost savings comes from the fact that the village fund distributes liquidity only to those who most need it. That is, households who increase their borrowing are constrained households who have the highest marginal valuation of liquidity. The consumption increase that exceeds the cost of

[12] Pitt and Khandker (1998), Pitt et al (2003), Morduch (1998), Coleman (1999), Ravicz (2000), Woolcock (1999), Aportela (1998), Gertler, Levine and Moretti (2003), and Karlan and Zinman (2006) are examples. Kaboski and Townsend (2003) estimates positive impacts of microfinance in Thailand using non-experimental data.

the transfers/credit injection is evidence of these relaxed borrowing constraints.

This paper also contributes to the current literature on the uses and relative merits of structural models and exogenous instruments or natural experiments in the estimation of program impacts (e.g. Heckman, Urzua, and Vytlacil, 2004, Rosenzweig and Wolpin, 2000, Banerjee et al, 2005, Angrist and Imbens, 1994). We argue that one of the great benefits of having exogenous variation is in estimating and testing structural theoretical models. In turn, theoretical models add insight toward the interpretation of program impacts. In line with Heckman, Urzua, and Vytlacil, we show that heterogeneous impacts across individuals can confound the interpretation of reduced form impact coefficients.[13] That is, in the context of our model, an intervention's impact on a household depends on an observed state (liquidity) and two unobserved states (permanent income and the size of potential project) of that household.[14] The measurement and interpretation of impact coefficients from reduced form regressions is further complicated by (i) multiple dimensions of impact and (ii) non-linear (and even non-monotone) relationships between the sizes of impacts and treatments, even when impacts are properly measured.

Regarding the multiple dimensions, we have at least three observable outcomes in our simple model – consumption, investment, and credit – that might be used to evaluate the impact of microfinance. We highlight here as an example three households that each benefit from enhanced access to credit, but exhibit strikingly different "impacts" according to these outcome measures. Household I increases borrowing and also reduces consumption in order to make an indivisible investment that would not have been undertaken without the additional credit. Household II uses additional credit to increase current consumption and shows no change in investment. Household III takes no credit at all, but nonetheless is able to now invest and increase its consumption since the increased *availability* of credit has reduced its need for maintaining as high a buffer stock of liquidity. Moreover, these

[13] We use these reduced form impact coefficients as interesting moments for empirically comparing the model to the data, but our impact evaluation is done using the structural model.

[14] Both are partially observable in that observable income gives a signal of permanent income, and investment discloses project size, but only for those who choose to invest.

three households could appear identical in their *ex ante* observables, having the same *level* of savings and income. Household I has high current income *relative to* permanent income but also has a relatively large potential investment opportunity which requires it to reduce consumption in order to invest. Household II has unusually low current income *relative to* its permanent income, and so borrows to smooth consumption. Household III has a relatively small potential investment opportunity and so does not need to borrow in order to invest, but increases consumption because it no longer requires as large a buffer stock in the future.

With respect to the non-linearity of impacts on observables, consider a household that is constrained in both its consumption and its investment and faces a *one-period increase* in available credit. Availability of a small amount of credit may benefit the household greatly in allowing it to increase its consumption toward its desired level, and additional credit just beyond the desired consumption loan would provide no benefit since the household's consumption is no longer credit constrained. As available credit increases even more, however, the household will at some point invest in its project, but it will reduce its consumption in order to do this. At this point marginal increases in available credit will again provide a large current benefit, since the household is again constrained in consumption. As credit increases still more, the household is at some point again able to meet its desired consumption and is no longer credit constrained. Further increases in available credit beyond this point are again of no benefit to the household.[15]

The remainder of the paper is organized as follows. The model is presented in Section 2. Section 3 discusses the computational methods used to solve the model, and characterizes the resulting value and policy functions. Section 4 discusses the data and presents the GMM estimation procedure and resulting estimates. Section 5 incorporates the introduction of Thai Million Baht Village Fund policy intervention into the model, compares the model's

[15] In Thailand, in particular, Coleman (2001) finds that small loans (1500-1700 baht) made to ordinary members of village banks did not have measurable impacts, but larger effective loans available to committee members did produce measured impacts. Our model would predict this non-linearity for investment, but not consumption.

predictions to the actual post-intervention data, and presents a cost benefit analysis of the program. Section 6 concludes.

2 Model

This section develops the model of consumption, investment and buffer stock saving with income uncertainty and borrowing constraints. We first present the sequential infinite horizon problem, and then derive a normalized recursive expression of the problem with reduced dimensionality.

2.1 Sequential Problem

In period 0, the household begins with a potential investment project of size I_0^*, permanent income P_0, and liquid wealth L_0. Liquid wealth L_{t+1} includes the principal and interest on liquid savings from the previous period $(1+r)\,S_t$ and current income Y_{t+1} :

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

In turn, current income Y_{t+1} is the product of the permanent component of income P_{t+1} and U_{t+1}, a one-period transitory shock:

$$Y_{t+1} \equiv P_{t+1}U_{t+1} \tag{2}$$

The exogenous component of permanent income follows a random walk based on shock N_t with drift G, but permanent income can also be increased endogenously through investment. 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:

$$P_{t+1} = P_tGN_{t+1} + RD_{I,t}I_t^* \tag{3}$$

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 liquidity.[16] Project size is stochastic, governed by an exogenous shock i_t^* and proportional to permanent income:

$$I_t^* = i_t^* P_t \tag{4}$$

Since i_t^* is exogenous, investment opportunities I_t^* increase with permanent income P_t, high permanent income alone will not automatically eliminate credit constraints in investment.[17]

Liquid wealth is bounded below by a borrowing limit which is a multiple $\underline{s}$ of its permanent income:[18]

$$S_t \geq \underline{s} P_t \tag{5}$$

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 investment:

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

[16]Although investment I_t increases the permanent component of income in the following period P_{t+1} by a deterministic amount RI_t, it increases next period liquidity L_t stochastically, since income Y_t is the *product* of the permanent component and a transitory shock. In contrast, the net yield on savings increases next period liquidity by a deterministic rS_t. Furthermore, investment increases future permanent income P_{t+2} stochastically, since current permanent income is multiplied by a stochastic shock in contributing to future permanent income. See equations (1)-(3) above.

[17]The assumption that wealthy households are also constrained in investment is justified in the data. Households in the top income tercile invest more often than poorer households but still only 22 percent of the time. Related, investment is not concentrated among the same households each year. On average 9 percent of households invest in each year. If this percent were independent across years, one would predict that 51 percent of households would never invest over the seven-year panel. This is quite close to the 58 percent that is observed.

[18]This is a somewhat arbitrary borrowing constraint (not a natural constraint as in Aiyagari, 1994), but such a constraint can arise endogenously in models with limited commitment (see Wright, 2002).

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:

- 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). $\ln i_t^* \sim N(\mu_i, \sigma_i^2)$

The model also produces the possibility of default. As written, if $\underline{s} < 0$ and an agent has debt, $S_{t-1} < 0$, it is possible that no positive value of consumption satisfies equations (7) and (5). Essentially, given (5), 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 ($\underline{c}P_t$). This condition for default is therefore:

$$(\underline{s} + \underline{c})\, P_t > L_t = Y_t + S_{t-1}(1 + r) \tag{8}$$

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

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

In sum, the parameters of the model are $\{\sigma_N, \sigma_u, \mu_i, \sigma_i, G, r, R, \rho, \beta, \underline{s}, \underline{c}\}$. We have explicitly emphasized the value function's dependence on $\underline{s}$, since this will be the parameter of most interest in considering the microfinance intervention in Section 5. We drop this emphasis in the simplifying math that follows.

2.2 Normalized and Recursive Problem

Given the set-up, the problem is homogenous in P_t. Using lower case variables to indicate variables normalized[19] by permanent income, the recursive problem becomes:

$$
\begin{aligned}
V(L, I^*, P) &\equiv P^{1-\rho} v(l, i^*) \\
v(l, i^*) &= \max_{c, s', d_i} \frac{c^{1-\rho}}{1-\rho} + \beta E\left[(p')^{1-\rho} v\left(l', i^{*\prime}\right)\right] && (9) \\
& \quad s.t. \\
\lambda &: c + s + d_i i^* \leq l && (10) \\
\phi &: s \geq \underline{s} && (11) \\
p' &= GN' + R d_i i^* && (12) \\
l' &= y' + \frac{s(1+r)}{p'} && (13) \\
y' &= U' && (14)
\end{aligned}
$$

We further simplify by substituting in for l' and y' into the continuation value using (13) and (14), and substituting out s using the liquidity budget constraint (10), which will hold with equality, to yield:

$$
\begin{aligned}
v(l, i^*) &= \max_{c,\, d_i} \frac{c^{1-\rho}}{1-\rho} \\
& \quad + \beta E\left[(p')^{1-\rho} v\left(U' + \frac{(1+r)(l - c - d_i i^*)}{p'}, i^{*\prime}\right)\right] && (15) \\
& \quad s.t. \\
\phi &: (l - c - d_i i^*) \geq \underline{s} && (16) \\
p' &= GN' + R d_i i^* && (17)
\end{aligned}
$$

The normalized form of the problem has two advantages. First, it lowers the dimension-

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

ality of the state variable. Second, it allows the problem to have a steady state solution.[20] The necessary conditions for optimal consumption c_* and investment decisions d_{i*} are:[21]

$$(c_*)^{-\rho} = \beta(1+r)E\left[(p')^{-\rho}\frac{\partial v}{\partial l}(U' + \frac{(1+r)(l - c_* - d_{i*}i^*)}{p'}, i^{*\prime})\right] + \phi \tag{18}$$

$$\begin{aligned} \frac{c_*^{1-\rho}}{1-\rho} + \beta E\left[(p')^{1-\rho}\, v\left(U' + \frac{(1+r)(l - c_* - d_{i*}i^*)}{p'}, i^{*\prime}\right)\right] \quad \geq \\ \frac{c_{**}^{1-\rho}}{1-\rho} + \beta E\left[(p')^{1-\rho}\, v\left(U' + \frac{(1+r)\left[l - c_{**} - (1-d_{i*})i^*\right]}{p'}, i^{*\prime}\right)\right] \end{aligned} \tag{19}$$

where c_{**} indicates the optimal consumption given the alternative investment decision (i.e., c_{**} satisfies the analog to (18) for $1 - d_{i*}$). The constraint ϕ is only non-zero when (16) binds, i.e., $c_* = l - \underline{s} - d_{i*}i^*$.

3 Computed Optimal Policy Functions

The value function and policy functions are solved computationally via recursion on the value function. Standard numerical approximations techniques are used to approximate (i) the value function and (ii) the expectations.

The approximation of the value function involves both states l and i^*. We use the method of collocation to approximate a two-dimensional linear spline along these two dimensions. The collocation method (see Miranda and Fackler, 2002) yields an approximation $\tilde{v}$ that solves the functional equation defined by (15) exactly at $m*n$ values (or "collocation nodes") of l and i^*. The number of nodes are chosen in order to ensure that the $\tilde{v}$ approximately satisfies (15) along all values over the relevant range of liquidity and investment. Since the value function is a continuous function, the collocation approach is able to lower

[20] 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.

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

the computational burden relative to discretization, since fewer nodes are needed relative to discrete points in the discretization approximation.[22]

Expectations over N',U', and $i^{*'}$ are computed numerically using a Gaussian quadrature method for numerical integration. We use 10 points each for values of N' and U', respectively. For integration with respect to $i^{*'}$, we use the 7 discrete collocation node values whose values are chosen to be the Gaussian-Hermite nodes for numerical integration. The computation is done in Matlab and relies heavily on the techniques and codes of Miranda and Fackler (2002).

Figure 1 presents a three-dimensional graph of the 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. 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. Consumption actually *falls* as liquidity increases beyond this threshold.

Figure 2 illustrates this more clearly by showing a cross-section of the optimal consumption policy function at a given value of i^*. 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. At this point, the borrowing constraint again holds with equality, and marginal increases in liquidity are used for consumption.[23]

[22]For the estimation, we use 8 nodes (including zero) along the i^* dimension times 24 along the l dimension in the estimation. Simulation draws actual values of i^*, instead of only using the Gaussian quadrature points, so a finer grid of 20 and 40 nodes, respectively, is used. For the sake of visual clarity, we use still more nodes in Figures 1-4.

[23]Using 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

Figure 1: Value Function vs. Liquidity Ratio & Project Size

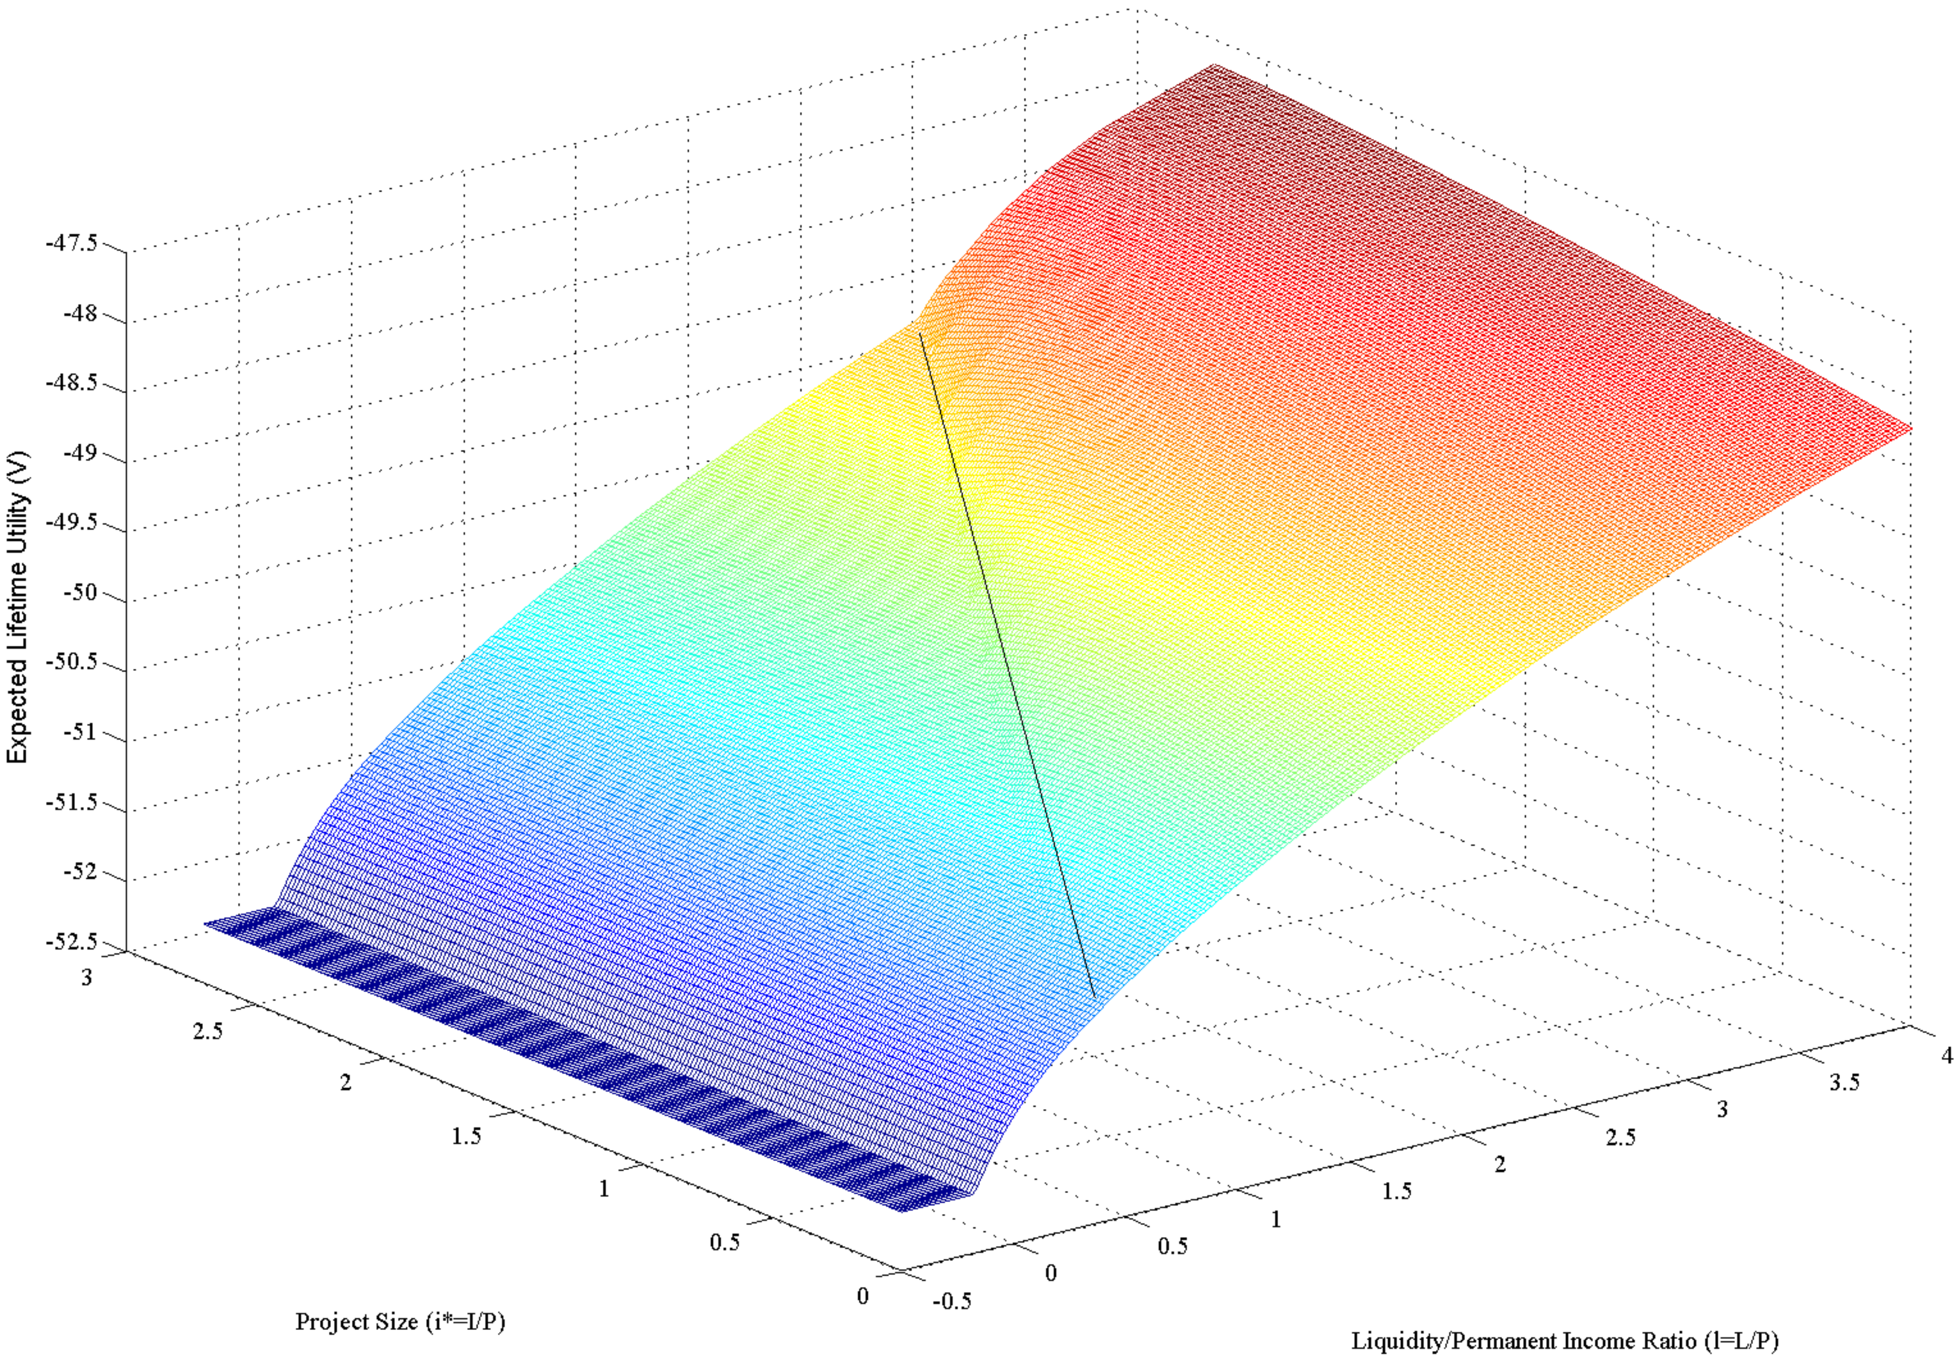

Figure 2: Consumption Policy for Given i*

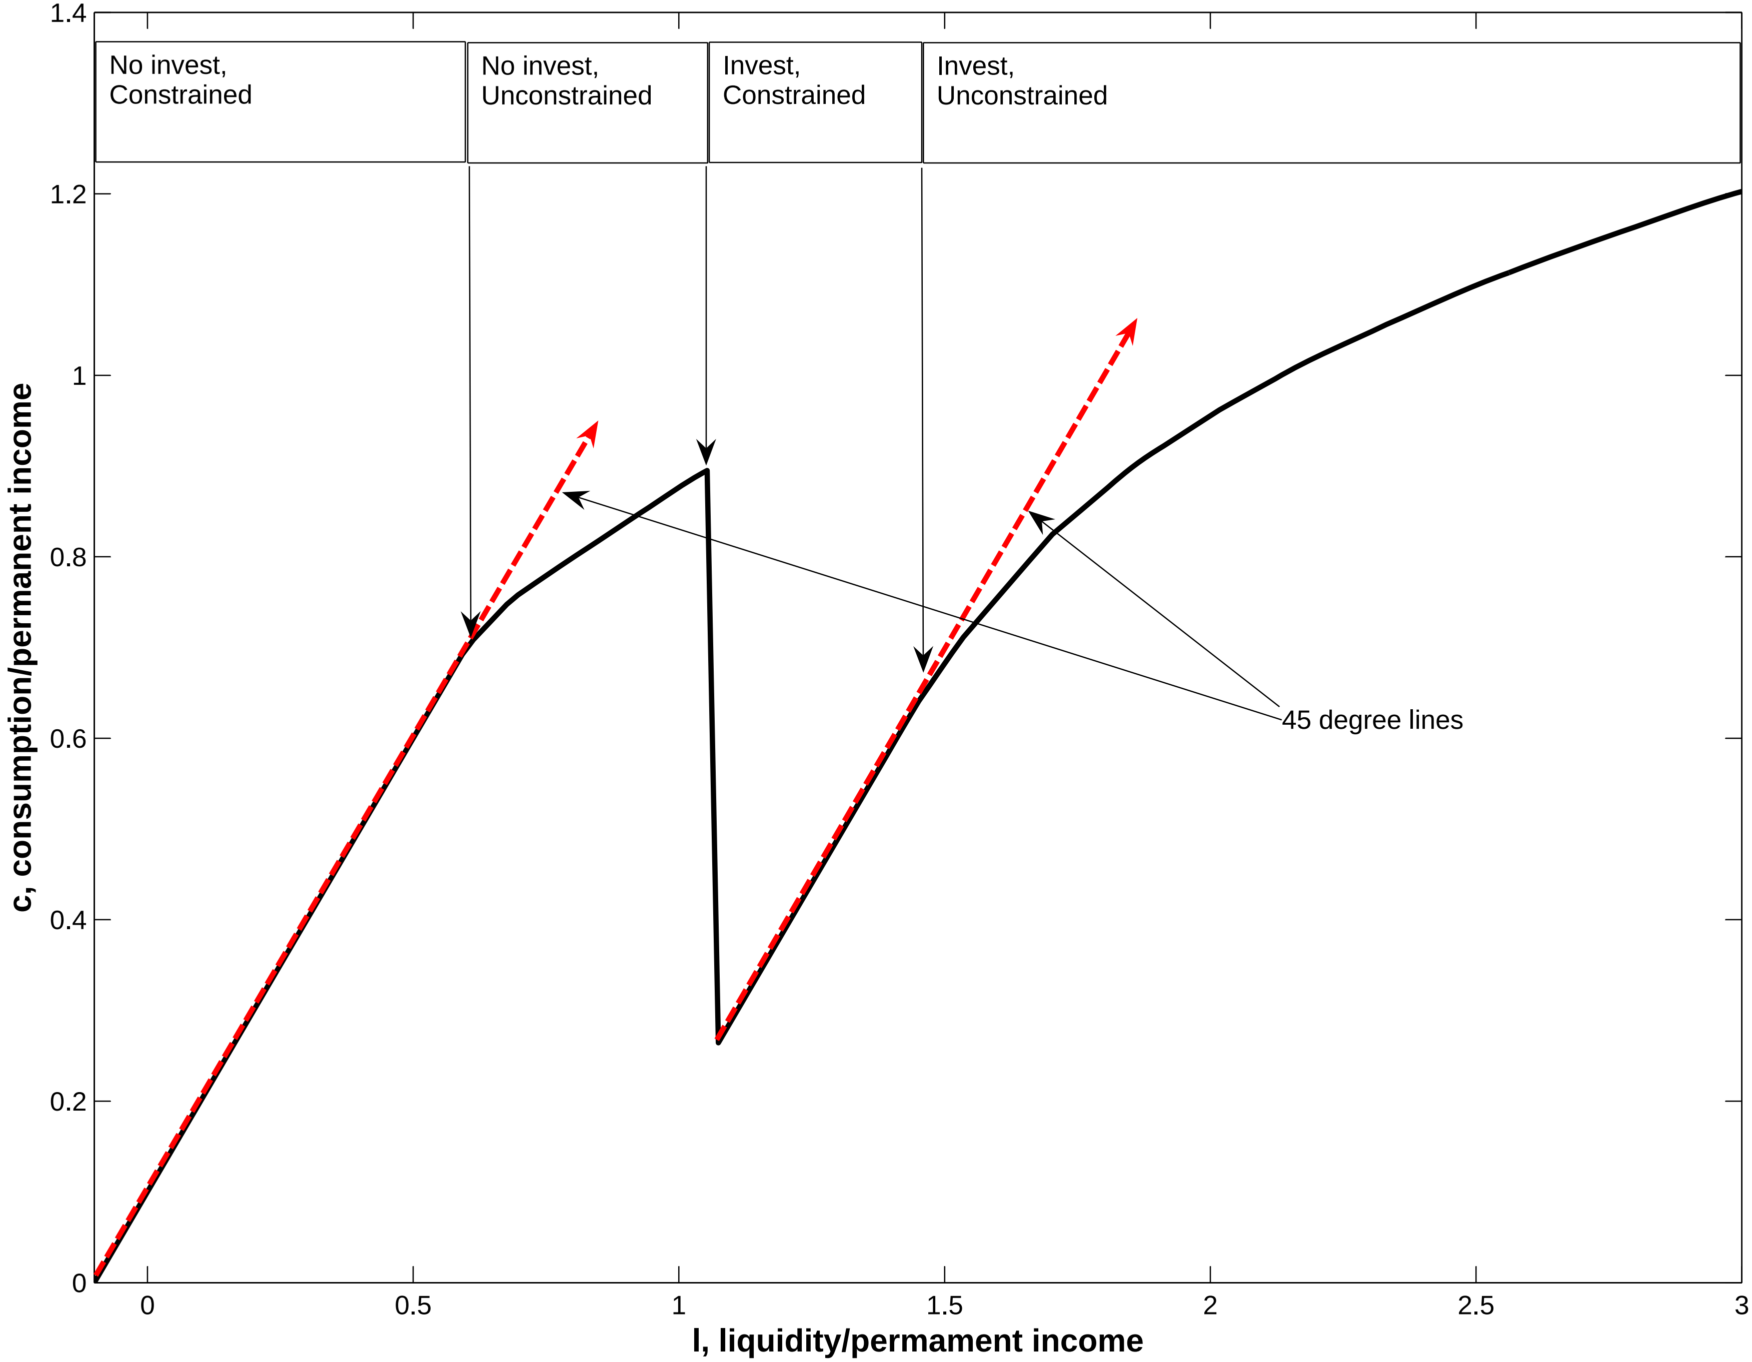

Figure 3 further displays this in a three-dimensional graph of consumption as a function of the liquidity ratio and project size and shows that l_* is increasing in i^*. The positions of Households I, II, and III from the example in the introduction are plotted on the figure. Recall each has the same realized absolute income Y and liquidity level L, and so could appear *ex ante* identical. In the graph, Household II has a lower normalized liquidity ratio l because it has higher permanent income P than the other two households, and Household III has a smaller project size i^* than the others (relative to its permanent income). Both permanent income and project size are unobservable *ex ante.*

Figure 4 shows the effect of a surprise permanent decrease in $\underline{s}$ on the optimal consumption and saving/borrowing policies for these three households. In the top panel, Household I decreases consumption because of investment, and Household II increases consumption through borrowing. In the bottom panel, Household III invests and consumes more without borrowing by simply decreasing its saving, which it does because of the decreased need to have as large a savings bufferstock in the future. This predicted increase in consumption even for non-borrowers is one way of distinguishing the bufferstock model from a neoclassical model, and will be important in the empirical section.

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. 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 further 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 increases the borrowing limit next period, and therefore reduces the probability of a "margin call" on

that also contain investment as an explanatory variable:

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

and 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).

Figure 3: Consumption Policy as a Function of Liquidity and Project Size

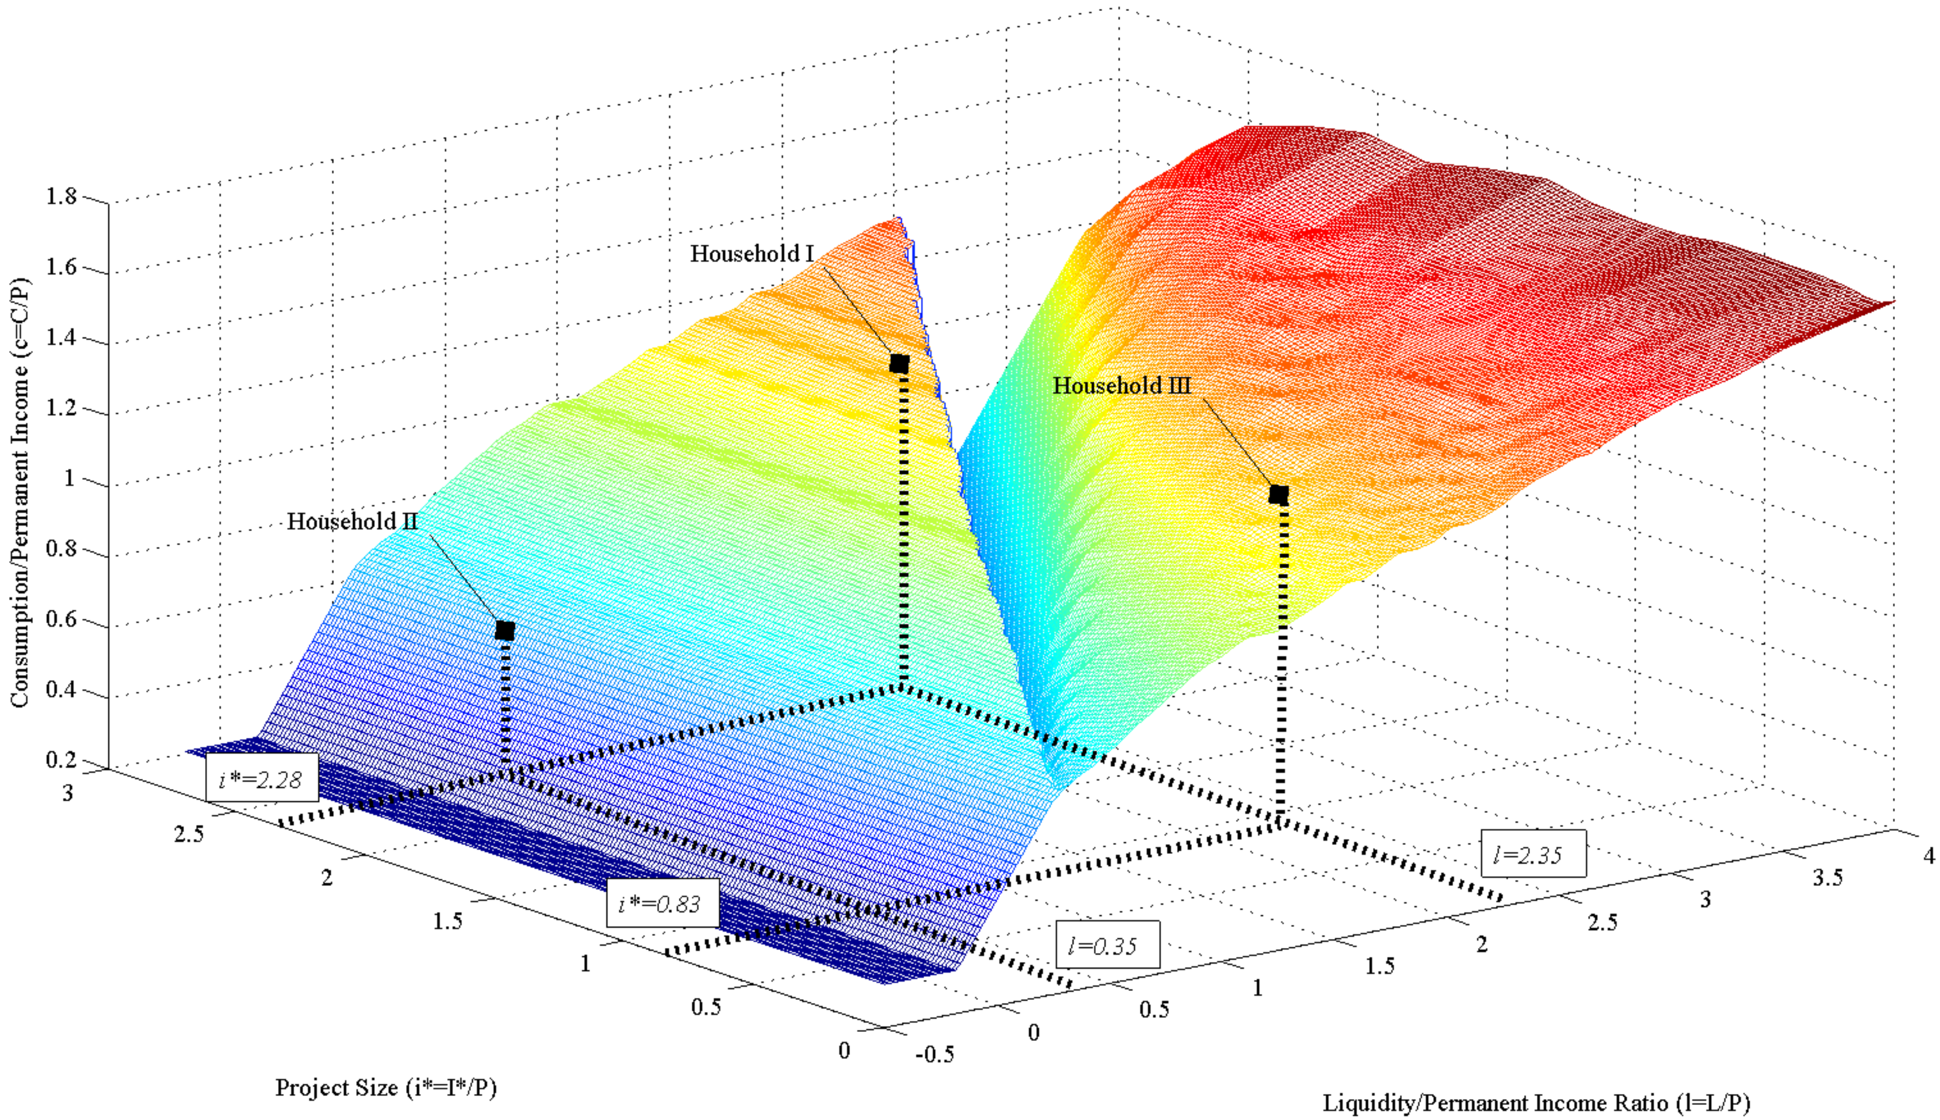

Figure 4: Examples of Consumption and Savings Policy Changes with Increased Borrowing Limit

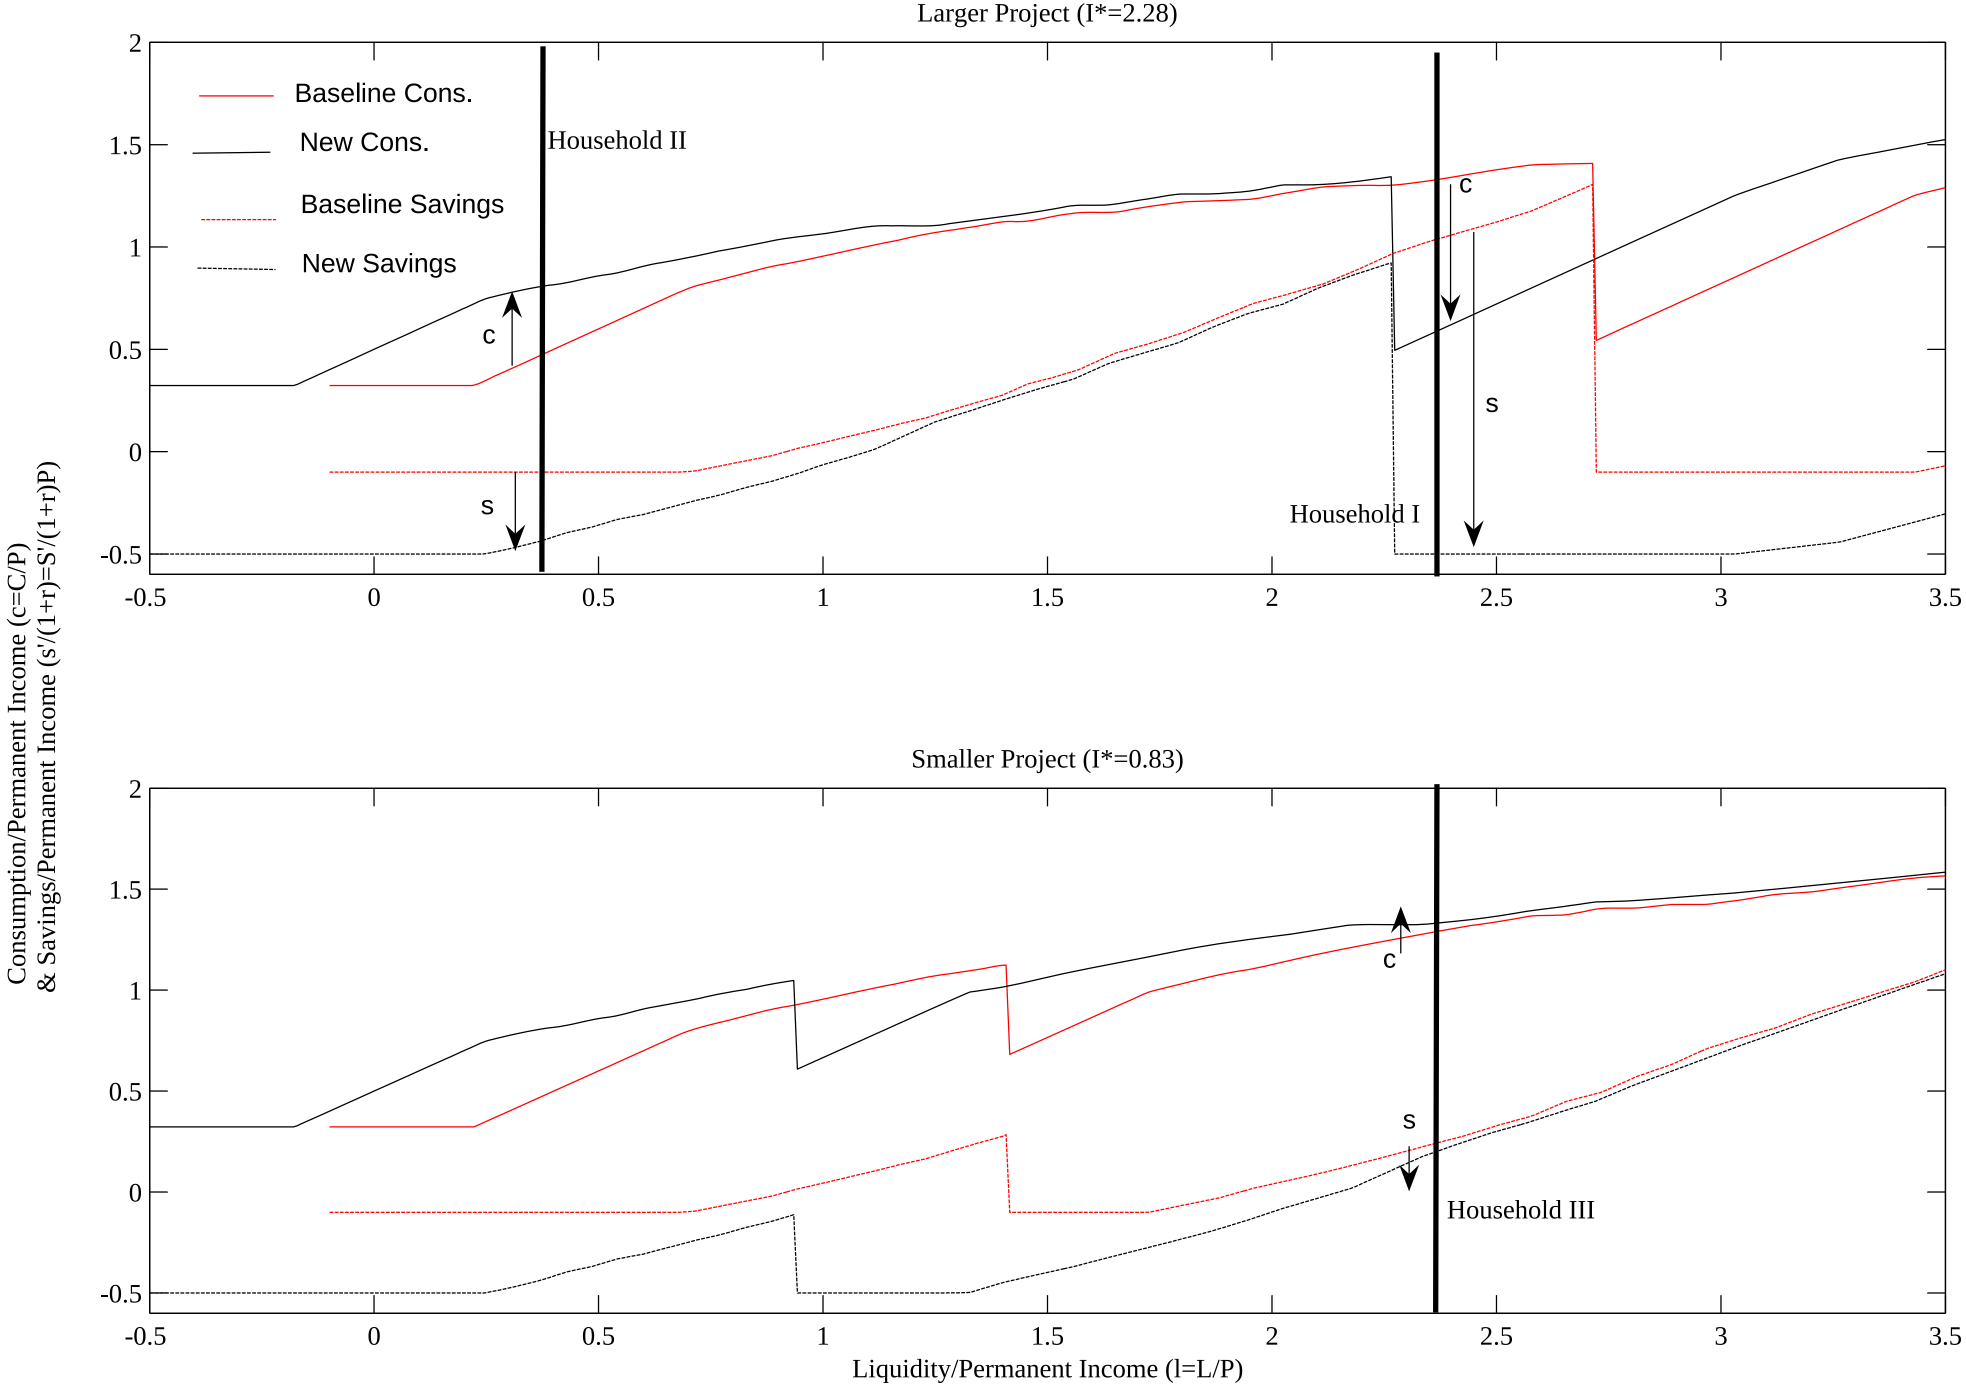

outstanding debt.

One can see this formally by substituting the definitions of liquidity (1) and income (2), and the law of motion for permanent income (3) into the condition for default (8) to yield:

$$\Pr\left(\text{Default at } t+1\right) = \Pr\left[U_{t+1} < (\underline{s} + \underline{c}) - \frac{S_t}{(P_t N_{t+1} G + R D_{I,t} I_t^*)}\right] \tag{20}$$

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

As in Gourinchas and Parker (2002), we divide the parameters to be estimated into two sets, $\chi = \{R, \sigma_N, \sigma_u, G, r, \underline{s} + \underline{c}\}$ and $\theta = \{\beta, \rho, \mu_i, \sigma_i, \underline{s}\}$, and estimate them using a two-stage GMM. The first set of parameters can be estimated using moments that involve only the data and no optimization from the model. We will only be able to identify the sum $\underline{s} + \underline{c}$ using default data in the first stage (see equation (20) above). The second step will involve solving the model's optimal policy functions. Both steps use the limiting optimal weighting matrix from multiple recursion (see Ogaki, 1993). We will be estimating the model using five years (1997-2001) of pre-intervention data, so that $t = 1$ corresponds to the year 1997.

4.1 First Stage GMM Estimation

Unfortunately, in estimating the income process parameters, σ_N, σ_u, and G, the techniques of Carroll and Samwick (1997) cannot be directly applied, since income is now endogenous to investment decisions. We borrow from their basic approach, however. Intuitively, the drift G can be identified by looking at average log growth rates in income, and in separately identifying and estimating the variances of permanent and transitory income shocks by looking at the variance of income growth across different time spans. The contribution of the variance of permanent shocks (i.e., σ_N^2) to the variance of income growth increases as the time span increases, while the contribution of the variance of transitory shocks (i.e., σ_u^2) stays constant. We do not estimate the return to investment R using GMM. Although it is theoretically identified, the low level of investment in the pre-experiment data does not allow us to estimate with any reasonable level of accuracy. Instead, we calibrate $R = 0.1$, which is consistent with an after-depreciation return calculated from households in the longitudinal monthly Thai data (Samphantarak and Townsend, 2006). Together we have 26 first-stage moment conditions to estimate the remaining six parameters in χ. All of these moments are defined on five years of pre-intervention data, which we denote X.

The four g_1 moments are derived from equations (2) and (3) and match the average income growth in the model to average income growth in the data:

$$g_{1t}\left(X,\chi\right) = \left\{\ln\left(Y_{t+1}/Y_t\right) - E\left[\ln(GN' + RI_t)|\sigma_N^2\right]\right\} \quad \forall t = 1, ..., 4$$

Moments g_2 through g_5 (ten in total) are derived from the same equations, but instead match the variance of growth rates over different horizons (k =1...4-year growth rates, respectively). We have four moments of one-year growth rates ($k = 1$, $t = 1, ..., 4$), three moments of two-year growth rates ($k = 2$, $t = 1, ..., 3$), and so on, for a total of 10 moments. The interior expectation in these moment conditions is an integration over values of N and is a function of σ_N^2 :

$$\begin{aligned} g_{k+1,t}(X,\chi) \quad = \quad & \left\{\ln\left(Y_{t+k}/Y_t\right) - \sum_{j=1}^{k} E_N\left[\ln(GN' + RI_{t+j-1})\right]\right\}^2 \\ & -E\left\{\left[\sum_{j=1}^{k}\left\{\begin{array}{c}\ln(GN' + RI_{t+j-1}) \\ -E\left[\ln(GN' + RI_{t+j-1})\right]\end{array}\right\}\right]^2 |\sigma_N^2\right\} - 2\sigma_u^2 \\ & \text{for } k = 1, ..., 4 \text{ and } t = 1, ..., 5-k \end{aligned}$$

Moments g_6 through g_8 are based on the assumption that default and interest (both paid and received) are measured with error. We discuss the reasons for this measurement error in the Data subsection. The g_6 moment condition is derived directly from the condition for default, by substituting $P_t = Y_t/U_t$ into the left-hand side inequality in equation (8). It helps identify the sum, $(\underline{s} + \underline{c})$, reducing the dimensionality of the second stage estimation. $DEFAULT_{t+1}$ is an indicator of whether a household defaulted on outstanding credit from year t.

$$\begin{aligned} g_{6t}(X,\chi) &= DEFAULT_{t+1} - \Pr\left[U_{t+1} < (\underline{s} + \underline{c})\,\frac{Y_t}{L_t}\right] \\ & \quad t = 1, ..., 4 \end{aligned}$$

The g_7 and g_8 moment conditions (eight in total) identify the interest rate on liquidity. In g_7, S_{t-1} is liquid savings while $EARNED\_INT_t$ is interest income received on this savings in the following year. Likewise, in g_8, CR_{t-1} is short-term credit owed at time $t-1$, and $PAID\_INT_t$ is interest paid on this short-term credit during the following year.[24]

$$\begin{aligned} g_{7t}(X,\chi) &= EARNED\_INT_{t+1} - rS_t \quad t = 1, ..., 4 \\ g_{8t}(X,\chi) &= PAID\_INT_{t+1} - rCR_t \quad t = 1, ..., 4 \end{aligned}$$

4.2 Second Stage GMM Estimation

The second stage estimates the set of parameters $\theta = \{\beta, \rho, \mu_i, \sigma_i, \underline{s}\}$. We have policy functions that are functions of θ, $C(L_t, P_t, I_t^*; \theta) = P_t c(l_t, i_t^*; \theta)$ and $D_i(L_t, P_t, I_t^*; \theta) = d_i(l_t, i_t^*; \theta)$, and observed data on decisions C_t and I_t, and states Y_t and L_t. We assume that the data are observed without measurement error, though the moments would be identical if we allowed for multiplicative classical measurement error in C_t or I_t.[25].We define moments based on the deviation of observed consumption and investment choices in the data from predictions conditional on the observables, liquidity L_t and income Y_t.

[24]In the data there are many low interest loans, and the difference between the average interest rate on short term borrowing and saving is small, just 2 percent.

[25]If there were measurement error in income, and it were classical, we would be interpreting measurement error and transitory income shocks, our estimates would be overestimates of the true σ_u from the first stage, and our estimate of the intertemporal elasticity of substitution/risk aversion parameter would also likely be an underestimate of the true ρ.

For log consumption, the deviation is defined below:

$$\begin{aligned}
\ln C_t &= \ln C(L_t, P_t, i_t^*) \\
\ln C_t - E\left[\ln C(L_t, P_t, I_t^*)|L_t, Y_t\right] &= \ln C(L_t, P_t, I_t^*) - E\left[\ln C(L_t, P_t, I_t^*)|L_t, Y_t\right] \\
&\equiv \varepsilon_{C,t}(C_t, L_t, Y_t; \theta)
\end{aligned}$$

where we have omitted the dependence of the policy function on θ until the last term. The above expectations are taken over realizations U_t and $i_t^{*,}$ which are both unobservable to the econometrician.

We can similarly define a deviation for the decision of whether to invest or not:

$$\begin{aligned}
D_{i,t} &= d_i(l_t, i_t^*) \\
D_{i,t} - E[d_i(l_t, i_t^*)|L_t, Y_t] &= d_i(l_t, i_t^*) - E[d_i(l_t, i_t^*)|L_t, Y_t] \\
&\equiv \varepsilon_{D,t}(D_{i,t}, L_t, Y_t; \theta)
\end{aligned}$$

and log investment when investments are made:

$$\begin{aligned}
D_{i,t} \ln I_t &= d_i(l_t, i_t^*) \ln I_t^* \\
\left\{ \begin{array}{l} D_{i,t} \ln I_t - \\ \quad E[d_i(l_t, i_t^*) \ln I_t^* | L_t, Y_t,] \end{array} \right\} &= \left\{ \begin{array}{l} d_i(l_t, i_t^*) \ln I_t^* - \\ \quad E[d_i(l_t, i_t^*) \ln I_t^* | L_t, Y_t] \end{array} \right\} \\
&\equiv \varepsilon_{I,t}(D_{i,t}, I_t, L_t, Y_t; \theta)
\end{aligned}$$

The conditional expectation required to compute the model's mean preidctions, hence the deviations, can be easily expressed in terms of the normalized policy functions, $c(l_t, i_t^*)$ and $d_i(l_t, i_t^*)$, and two shocks of known distribution, U_t and i_t^*.[26]

[26] The conditional expectation formulas are:

$$\begin{aligned}
E[C(L_t, P_t, i_t^*)|L_t, Y_t] &= \ln Y_t + \int \ln c(L_t U/Y_t, i^*) f(U) g(i^*) dU di^* \\
E[D_i(L_t, P_t, i_t^*)|L_t, Y_t] &= \int d_i(L_t U/Y_t, i^*) f(U) g(i^*) dU di^* \\
E[D_i L_t, P_t, i_t^*) \ln I^* | L_t, Y_t] &= \int \left[d_i(L_t U/Y_t, i^*) \ln\left(i^* Y_t/U\right)\right] f(U) g(i^*) dU di^*
\end{aligned}$$

where $f(U)$ and $g(i^*)$ represent the estimated lognormal distributions. Since we make no assumptions on the prior distribution of P_t, conditioning on Y_t and L_t does not affect the distribution of U.

By the Law of Iterated Expectations, the expected deviations are zero. Since they are conditional on income and liquidity, their interaction with functions of income and liquidity should also be zero in expectation. 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 overpredicting at low income or liquidity levels, while underpredicting at high levels.

We omit the functional dependence of these deviations in expressing below the nine valid moment conditions that we use:

$$\begin{array}{ccc} E\left[\varepsilon_{C,t}\right]=0 & E\left[\varepsilon_{D,t}\right]=0 & E\left[\varepsilon_{I,t}\right]=0 \\ E[\varepsilon_{C,t}\ln Y_t]=0 & E[\varepsilon_{D,t}\ln Y_t]=0 & E[\varepsilon_{I,t}\ln Y_t]=0 \\ E[\varepsilon_{C,t}\left(L_t/Y_t\right)]=0 & E[\varepsilon_{D,t}\left(L_t/Y_t\right)]=0 & E[\varepsilon_{I,t}\left(L_t/Y_t\right)]=0 \end{array}$$

Using five years of pre-experiment data, we can construct four years of moments ($t = 2, ..., 5$).[27] We therefore have thirty-six moments for the estimation of the five parameters in $\theta = \{\beta, \rho, \mu_i, \sigma_i, \underline{s}\}$.

4.3 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, 720 with complete data in the pre-experiment balanced panel used for estimation, and 719 and 705 in 2002 and 2003, respectively, which are used only for 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 to the middle of the pre-experiment data, 1999.

[27] We cannot use the first year, since we need s_{t-1} to construct L_t

The measure of household consumption we use (denoted $\tilde{C}_{j,t}$ for household j at time t) is calculated using detailed monthly expenditure data and weights derived from the Thai Socioeconomic Survey.[28] The income measure we use (denoted $\tilde{Y}_{j,t}$) 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. Our savings measure ($S_{j,t}$) includes not only savings deposits in formal and semi-formal financial institutions, but also the value of rice holdings in the household. The measure of liquid credit (CR_t) is short-term credit with loan durations of one year or less. The measure of investment ($\tilde{I}_{j,t}$) we use is total farm and business investments net of credit from the BAAC for such investment, the Thai agricultural bank.[29] Default is imputed each year for households who report short term loans the previous fifteen months that are outstanding at least three months. This only approximates default in the model, and it may underestimate default because of underreporting, but overestimate default as defined in the model to the extent that late loans are eventually repaid. The measurement of interest income on liquid savings ($EARNED\_INT_{j,t}$) is interest income in year t on savings in formal and semi-formal institutions. The interest paid on credit ($PAID\_INT_{j,t}$) is the reported interest owed on credit. A major source of measurement error on interest is that savings and borrowing may fluctuate within the year, sot that both earned and paid interest may not accurately reflect interest on the measured end of year stocks.

Table 1 presents key summary statistics for the data.

[28] The tildes represent raw, unnormalized data. See section 4.3.1.

[29] We subtract BAAC credit for several reasons. First, access to investment credit from the BAAC is extremely heterogeneous across households, regions and types of investment. Second, some BAAC loans, especially those that finance the largest investment projects, are long term loans. These first two reasons make make it difficult to incorporate BAAC-financed investment into the model in a straightforward way. The third reason for excluding BAAC-financed investment is that it creates occassional large outlier, from, for example, purchase of a large amount of land, which if included would carry too much weight in the estimation and would swamp the little sources of investment that microfinance programs are designed to promote.

Table 1: Summary Statistics of Household Data

Variable	Obs	Mean	Std. Dev.	Min	Max
Primary Variables:					
Non-Interest Household Income*	4304	92400	200200	600	6260500
Household Consumption*	4303	68600	86500	750	1488800
Dummy Variable for Agr/Business Investment	4304	0.093	0.29	0	1
Value of Agr./Business Investment*	4304	4300	51500	0	2461800
Dummy Variable for Short-Term Default	4304	0.163	0.409	0	1
Short-Term Credit*	4304	9700	36500	0	994035
Interest Paid*	4304	560	2270	0	41920
Liquid Savings*	4304	22900	111858	0	4700600
Interest Earned*	4304	320	4200	0	238350
Number of Households in Village	4304	166	295	21	3194
Regressors for Demographic/Cylical Variation:					
Number of Male Adults	4304	1.46	0.88	0	7
Number of Female Adults	4304	1.56	0.75	0	5
Number of Children	4304	1.53	1.19	0	9
Dummy Variable for Male Head of Household	4304	0.72	0.45	0	1
Years of Education of Head of Household	4304	6.5	3	0	15
Birthyear of Head of Household	4304	1946	14	1909	1974
Exchange Rate	4304	41.4	2.1	37.8	44.5
Exchange Rate, Lagged	4304	39.7	4.3	31.4	44.5
Exchange Rate, Twice Lagged	4304	36.7	6.6	25	44.5

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

4.3.1 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. The data, however, contain important variation in household composition across time and households. The data also contain business cycle variation which is not included in the model. Ignoring either of 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 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.

We therefore follow Gourinchas and Parker (2002) in removing the business cycle and household composition variation from the data. In particular, 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.) The estimated equations are:

$$
\begin{aligned}
\ln \tilde{Y}_{nt} &= \alpha_Y X_{nt} + \delta_Y Z_t + e_{Y,nt} \\
\ln \tilde{C}_{nt} &= \alpha_C X_{nt} + \delta_C Z_t + e_{C,nt} \\
\tilde{L}_{nt}/\tilde{Y}_{nt} &= \alpha_L X_{nt} + \delta_L Z_t + e_{L,nt}
\end{aligned}
$$

where X_{nt} 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, birth year of head of household, and years education of head of household) for household n and Z_t is a common vector of variables capturing the business cycle (i.e., current value and two lags of the baht/dollar exchange rate).[30] These regressions are run using only the pre-program data, 1997-2001. The R^2 values for the three regressions are small: 0.09, 0.17, and 0.01, respectively.

For the full sample, 1997-2003, we construct the adjusted data for a household with mean values of the explanatory variables ($\overline{X}$ and $\overline{Z}$) using the estimated coefficients and residuals:

$$\begin{aligned}\ln Y_{nt} &= \hat{\alpha}_Y \overline{X} + \hat{\delta}_Y \overline{Z} + \hat{e}_{Y,nt} \\ \ln C_{nt} &= \hat{\alpha}_C \overline{X} + \hat{\delta}_C \overline{Z} + \hat{e}_{C,nt} \\ L_{nt}/Y_{nt} &= \hat{\alpha}_L \overline{X} + \hat{\delta}_L \overline{Z} + \hat{e}_{L,nt}\end{aligned}$$

Furthermore, we construct investment data I_{nt} by multiplying the actual measured values of i_{nt} (i.e., $i_{nt} = \tilde{I}_{nt}/\tilde{Y}_{nt}$) by the newly constructed income data Y_{nt}.

4.4 Estimation Results

Table 2 presents the estimation results for the structural model. The first stage estimates show the high level of income uncertainty – both permanent $\hat{\sigma}_N$ and transitory $\hat{\sigma}_u$ – that households face, the relatively small exogenous growth in income over the period $\hat{G}$, and the fact that households will default if liquidity falls below about 20% of permanent income ($\underline{s} + \underline{c}$).

For the second stage variables, the discount factor $\hat{\beta}$ of 0.961 and elasticity of substitution $\hat{\rho}$ are within range of the usual values. 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 (1.61 vs. -1.96), and there is a greater variance in project size $\hat{\sigma}_i$ than in investments in the data (4.08 vs. 1.50). In the model, these difference between the average sizes of realized investment and

[30] We could capture business cycle or secular variation more generally by adding year-specific fixed effects, but these could not be extrapolated to the post-program data.

Table 2: Parameter Estimates and Model Fit

First Stage GMM			Second Stage		
Parameter	Estimate	Std. Err.	Parameter	Estimate	Std. Err.
σ_N	0.128	0.019	β	0.961	0.006
σ_U	0.343	0.023	ρ	1.41	0.30
G	1.005	0.009	μ_i	1.61	0.10
r	0.015	0.002	σ_i	4.08	0.32
$\underline{s}+\underline{c}$	0.213	0.012	$\underline{s}$	-0.078	0.040
J-Value		249.5	J-Value		1318.2

Pre-Program Averages	$\mathbf{C_t}$	$\mathbf{D_t}$	$\mathbf{I_t}$
Data	50,900	0.095	2,400
Fitted Model	33,700	0.173	5,800

potential projects stem from the fact that larger potential projects are much less likely to be undertaken. The estimated borrowing constraint parameter $\underline{\hat{s}}$ indicates that agents could borrow up to about eight percent of their annual permanent income as short-term credit in the baseline period.

The standard errors are relatively small for the first-stage estimates, and larger in the second stage, particularly for the elasticity of substitution, and project size distribution. The Hansen J-tests easily reject the overidentifying restrictions of the model at both estimation states, but given the relatively parsimonious parametrization (5 parameters relative to 26 moments in the first stage, and 5 parameters and 36 moments in the second stage), this is not surprising.

In the second stage, the optimal weighting matrix puts the bulk of the weight on the direct levels of consumption, investment, and the probability of investing, and puts slightly more weight on the earlier years. Nonetheless, in absolute levels, rather than logs, the fitted model's predictions differ considerably from the pre-program data, with the model predicted too frequent investment and too little consumption. The difference for average investment is driven almost entirely by the higher frequency of investment, rather than investment size.

Again, the model is statistically rejected for the pre-program period, which tells us that the model is not the real world.[31] Still, 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 Thai Million Baht Credit Experiment

The exogenous 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

[31] One major dimension that it misses is in matching the moments on a year to year basis. Perhaps the exchange rate controls, Z_t, are not fully capturing the secular or business cycle variation in the data.

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, urban[32] 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 design and organization of the funds were intended to allow all existing villagers equal access to loans through competitive application and loan evaluation, with minimal membership requirements. The funds make short-term loans with a typical nominal interest rate of six percent (about 4 percent real). The funds and program are described in more detail in Kaboski and Townsend (2004).

As described in the introduction, the program design was beneficial in two ways. First, because it arose from a quick election, after the Thai parliament was dissolved in November, 2000, and was rapidly implemented in 2001, households would not have anticipated the program in earlier years. We therefore model the program as a surprise. Second, the same sized fund was established in each village, regardless of the size, so smaller villages received more funding per household and therefore we would predict more relaxation of credit constraints in small villages. Differences in village household population in 2001 therefore yield exogenous variation in credit constraint reduction. Indeed, Kaboski and Townsend (2004) show that in IV regressions, village population has no significant relationship on credit or outcome variables until after the program is implemented. Figure 5 shows the strong relationship in 2002, after the program.

The injection of credit is incorporated into the model as a reduction in $\underline{s}$. That is, for each village v – a subscript we now add to the notation – 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

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

Figure 5: Village Fund Credit/Household vs. Inverse Village Size in 2002

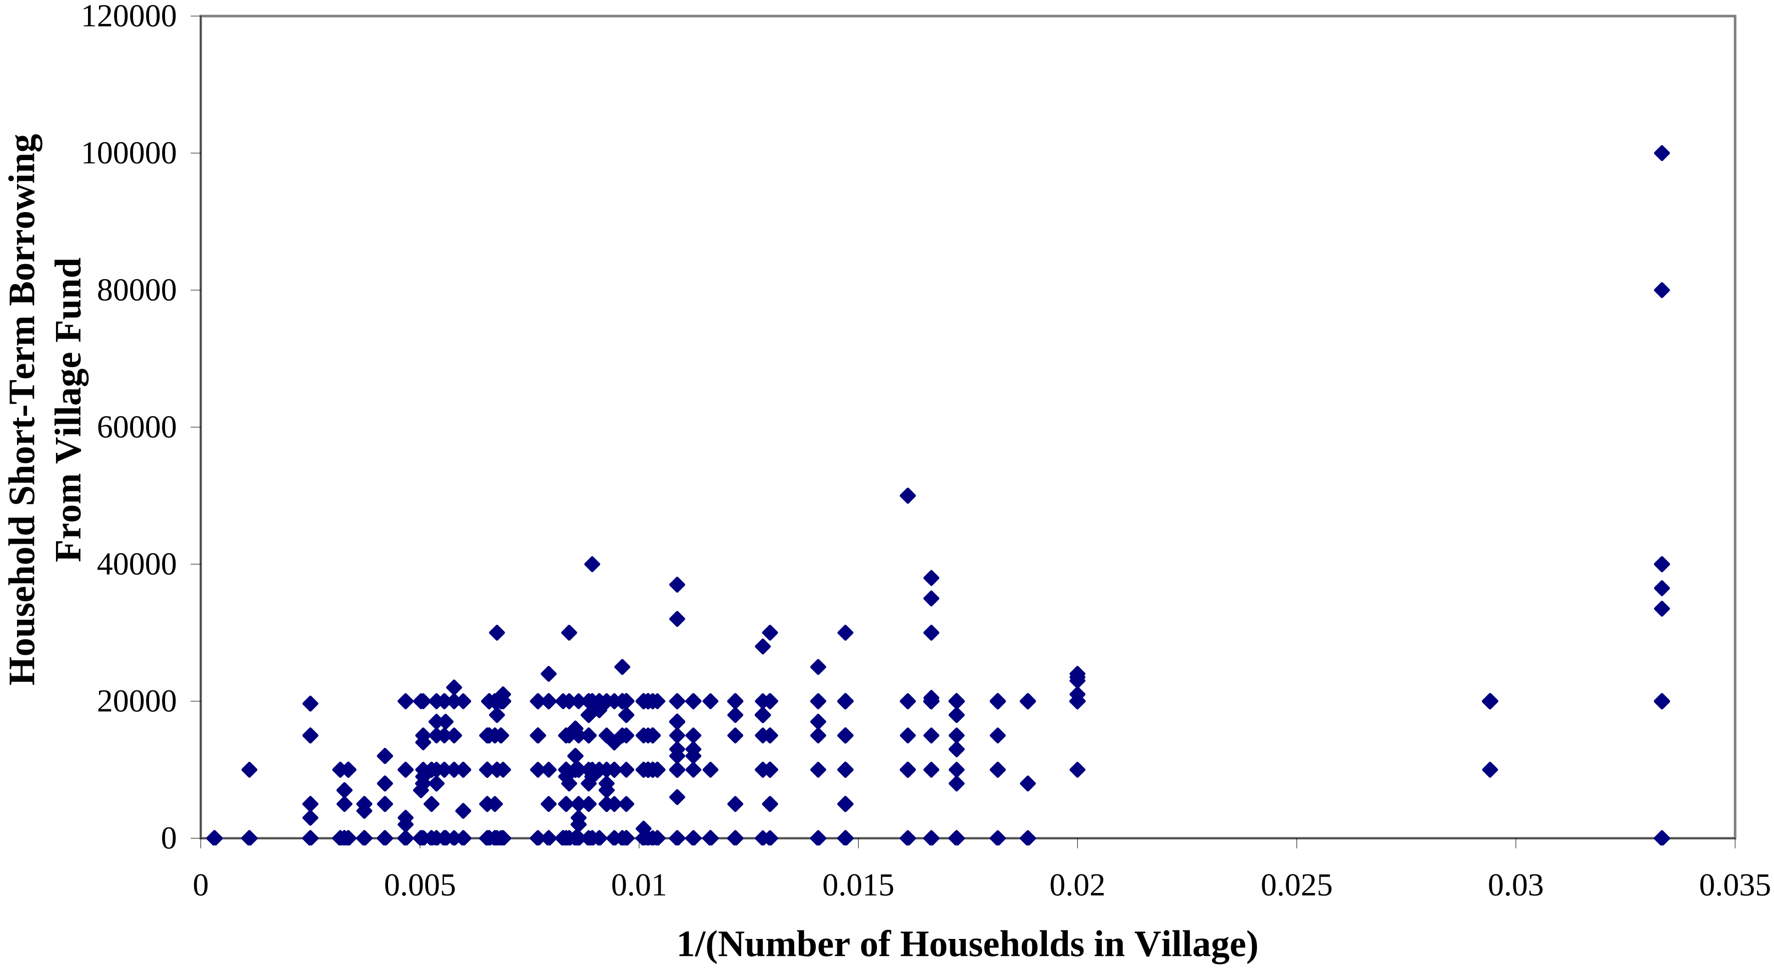

Notes: Each dot represents a household.

million baht fund intervention:

$$E\left[B^{mb}_{n,t,v}|L_{n,t,v},Y_{n,t,v};\underline{s}^{mb}_v\right]=E\left\{\mathcal{I}_{<0}\left[\begin{array}{c}L_t-C(L_t,P_t,I^*_t;\underline{s}^{mb}_v)\\-D_I(L_t,P_t,I^*_t;\underline{s}^{mb}_v)I^*_t\end{array}\right]|L_{n,t,v},Y_{n,t,v}\right\}$$

and in the baseline without the intervention:

$$E\left[B_{n,t,v}|L_{n,t,v},Y_{n,t,v};\underline{s}\right]=E\left\{\mathcal{I}_{<0}\left[\begin{array}{c}L_t-C(L_t,P_t,I^*_t;\underline{s})\\-D_I(L_t,P_t,I^*_t;\underline{s})I^*_t\end{array}\right]|L_{n,t,v},Y_{n,t,v}\right\}$$

where $\mathcal{I}_{<0}$ is shorthand notation for the indicator function that the bracketed expression is negative (i.e., borrowing and not savings). We choose $\underline{s}^{mb}_v$ so that additional predicted borrowing in the village averages one million baht in 2002 ($t=6$) and 2003 ($t=7$)[33]:

$$\frac{1}{2N_v}\sum_{t=6,7}\sum_{n=1}^{N_v}\left\{\begin{array}{c}E\left[B^{mb}_{n,t,v}|L_{n,t,v},Y_{n,t,v};\underline{s}^{mb}_v\right]\\-E\left[B_{n,t,v}|L_{n,t,v},Y_{n,t,v};\underline{s}\right]\end{array}\right\}=\frac{1,000,000}{\#\text{ HHs in village}_v}$$

Here N_v represents the number of surveyed households in village v, and is multiplied by two because of the two years of data used.

We should note it is not by construction that the calibration increases predicted expenditures in the village by one million baht. First, some of the credit may simply go to households who would otherwise have defaulted on existing credit. In this case, the borrowing households would increase expenditures by less than the amount borrowed (if at all). Second, non-borrowers will generally lower their savings in response to the relaxed borrowing constraint, so that expenditures could increase by more than what is borrowed.

The resulting $\underline{s}^{mb}_v$ values average -0.70, with a standard error of 0.31, a minimum of -1.67 and a maximum of -0.149. Hence, the post-program ability to borrow is substantial relative to the baseline ($\underline{s}=0.078$), averaging about two-thirds of permanent income after the introduction of the program.

5.1 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 three dimensions: log consumption, probability

[33] Since 1999 is the base year used, the million baht is deflated to 1999 values.

of investing, and log investment levels. We simulate 1000 sets of data for the two post program years by using the observed liquidity ($L_{n,t}$) and income data ($Y_{n,t}$) for those years, but drawing $U_{n,t}$ and $i^*_{n,t}$ shocks from the estimated distributions, and then using the model to predict consumption and investment behavior. For each dataset, we combine simulated consumption and investment data for the post-program years, with the actual data on consumption and investment in the pre-program years, and actual data on liquidity and income for the full sample.

We then ask whether reduced-form regressions would produce similar impact estimates using simulated data as they would using the actual data, even though statistically the model is rejected. The reduced form regressions regress the choice variables on time dummies, the observed states, and the average credit per household injected in the program years[34]:

$$\begin{aligned}
C_{n,t} &= \theta_{C,t} + \gamma_{C,1} Y_{n,t} + \gamma_{C,2} L_{n,t} + \gamma_{C,3} \frac{1,000,000}{\#\text{ HHs in village}_v} \mathcal{I}_{t>5} + e_{C,n,t} \\
D_{n,t} &= \theta_{D,t} + \gamma_{D,1} Y_{n,t} + \gamma_{D,2} L_{n,t} + \gamma_{D,3} \frac{1,000,000}{\#\text{ HHs in village}_v} \mathcal{I}_{t>5} + e_{D,n,t} \\
I_{n,t} &= \theta_{I,t} + \gamma_{I,1} Y_{n,t} + \gamma_{I,2} L_{n,t} + \gamma_{I,3} \frac{1,000,000}{\#\text{ HHs in village}_v} \mathcal{I}_{t>5} + e_{I,n,t}
\end{aligned}$$

Here $\hat{\gamma}_{C,3}$, $\hat{\gamma}_{D,3}$, and $\hat{\gamma}_{I,3}$ would be estimates of the impact of the program on consumption, investment probability, and average investment, respectively, since the village size during the program years is the exogenous intervention variable in the analysis. Although there is heterogeneity across households and non-linearity in the impact of credit, these coefficients capture (a linear approximation of) the relationship between the average change (e.g., in consumption) and the average credit per household injected in a village. We run these regressions on both the simulated and actual data and compare the estimates and standard errors. The year dummies in both regressions will pick up any differences in average levels between the simulated and actual data, and so identification of impact is from the cross-

[34] Recall that the consumption, investment, liquidity, and income data have already been purged of variation based on household composition, education, and business cycle fluctuations captured by the exchange rate and its lags.

sectional variation.

Table 3 shows that the model does quite well in replicating these reduced form impact estimates. Regressions using the actual data produce a highly significant estimate for consumption of 1.07. The point estimate of the probability of investing is positive but not significant, while the coefficient for the level of investment is actually negative, though again insignificant. (This negative sign is again driven by a large investment outlier in the data).

Regressions using the simulated data produce very similar results. See the bottom of Table 3. The first row shows the average (across 1000 samples) estimated coefficient, 1.07 for consumption, and the standard error of these 1000 estimates is given in parentheses, 0.28. The second row gives the average (across 1000 samples) estimated standard error, 0.19, and the standard error of these 1000 standard error estimates, 0.03. The main point is that the estimates in the data are typical of the estimates the model produces. The estimate in the data of 1.07 is quite similar (actually, identical) to the average estimated coefficient across the 1000 samples, while the standard error of this coefficient in the data of 0.21 is very close to 0.19, the average estimated standard error across the 1000 samples.

Again, the large impact on consumption distinguishes the bufferstock model from a neoclassical world without credit constraints. In the neoclassical world, increased credit from the village fund (at the same interest rate as existing credit) would crowd out borrowing from other sources and have no effect on consumption. If the village fund credit were at a lower interest rate, overall borrowing might increase as might consumption through a substitution effect. Still, even in the extreme counterfactual case in which credit was lent at a zero gross interest rate, which would amount to a one time gift of wealth, consumption would typically be smoothed intertemporally, and so consumption would not respond one for one. In the bufferstock model, consumption can respond more than one for one because even non-borrowers (like Household III, in the previous section) may increase consumption by reducing their bufferstocks. Furthermore, while current investment may reduce consumption for some (like Household II), the prospects of more future investment for more households increase lifetime income, hence current consumption.

Table 3: Reduced Form Regression Estimates:
Actual Data vs. "Million Baht" Simulated Data
(Standard Errors in Parentheses)

Actual Data	**Consumption**	**Investment Probability**	**Investment**
γ, "Impact" Coefficient*	**1.07**	1.14e-06	-0.064
Standard Error	**0.21**	1.29e-6	0.077
Simulated Data			
γ, Average "Impact" Coefficient*	**1.07** **(0.28)**	1.80e-6 (1.23e-6)	0.29 (1.19)
Average Standard Error	**0.19** **(0.03)**	1.25e-6 (1.49e-8)	0.50 (0.53)

*The impact coefficient is the coefficient on 1,000,000/number of households in the village, the credit injection per household.

Bold face represents significance at a 5 percent level.

The simulated estimates for investment probability and investment are positive, but insignificant. The average estimated standard error is large, but both the coefficient and standard error estimates show a great deal of variance over the 1000 samples, and the estimates in the actual data are within one standard error of the model averages. In theory, the effect of increased liquidity on investment is clearly positive. When the 1000 samples are pooled together, the pooled estimates of 1.87e-6 (standard error=3.96e-8) on the probability of investing and 0.317 (standard error=0.038) on investment are indeed highly significant. Multiplying by the average credit injection per household of about 9500 baht, the former amounts to just a 1.78 percent increase in the average probability of investing (relative to a pre-sample average of 9.48 percent), while the latter amounts to an average increase in investment of 3010 baht per household (relative to a pre-sample average of 2440 baht/household). In the model, of course, these effects are non-linear and heterogeneous across agents.

The model gives insight into why the effects on investment are difficult to find in reduced form regressions on the data.[35] First, the sizes of investments are extremely skewed, with a very few large investments amounting to a large fraction of total investment. Thus, although the increase in investment levels amounts to thirty percent (0.317) of credit in expectation, in the simulated data overall investment levels are driven by large outliers (just as in the actual data), and so extremely large data sets are needed to discern this effect. Second, the average project is much larger than the average investment. Since most projects are too large to undertake, the investment probability is low.[36] For households with very large projects, small increases in liquidity have no effect on their ability to invest. Third, households with very small projects may invest regardless of whether they have increased liquidity through microfinance. Hence, it is only a small group of households

[35] Kaboski and Townsend (2006) run similar astructural regressions but with household-specific fixed effects, which reduce the variance in the error term. In these regressions, the positive impact on the decision to invest, though small, is indeed significant.

[36] Another way of interpreting this 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.

with intermediate-sized projects, whose investment decision might be affected by small increases in available liquidity.

While households who increase investment benefit from the program, both households who were already investing and households who do not invest may also benefit from increased liquidity. Their benefit is realized through better consumption smoothing and/or an increase in current consumption. We formally quantify the program's benefits below.

5.2 Normative Analysis

The cost benefit analysis is done by comparing the cost of the program (the reduction in $\underline{s}$) to a transfer program (an increase in l) that is equivalent in terms of providing the same expected level of utility (given $L_{n,t}$ and $Y_{n,t}$ in 2002, when the program is introduced). That is, we solve the equivalent transfer T_n for each household using the following equation:

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

The average equivalent liquidity transfer per household in the sample equals 14,900 baht. The Million Baht Program costs just 9500 baht per household in the sample, or about 66 percent of the direct transfer program.[37]

We provide a heuristic interpretation of the advantages of the program over a simple transfer. Compare the Million Baht Program (the transfer of a million baht to the village in the form of a village fund) to a direct household level transfer program that costs one million baht. Both programs cost the same, and both constitute a wealth transfer of a million baht to the village as a whole. Hence the major difference is the value of the increased liquidity the programs create. There are three broad types of agents who can receive the funds. First, there are unconstrained households, who save the funds. Their

[37]This includes only the seed fund, and omits any administrative or monitoring costs of the village banks. Nonetheless, since the cost savings is substantial, administrative or monitoring costs would need to be quite large to overturn the result. It also omits potential distortions involved in raising tax revenues for the funds. These distortions would likely exist with both the Million Baht Program and the alternative transfer program.

marginal valuation of liquidity is therefore the market interest rate, $1+r$, times the marginal utility of consumption tomorrow. Second, non-defaulting constrained households. For these households, the marginal valuation of funds exceeds the market interest rate. Finally, there are defaulting constrained households. For these households, marginal liquidity simply goes to pay off credit which would have defaulted, so their marginal valuation of additional liquidity is zero. The transfer program gives funds to all households. The village fund distributes the funds by lowering the borrowing constraint, and increasing credit. Hence, only constrained households receive the liquidity. But default is relatively uncommon, even for constrained households, so it is the non-defaulting credit constrained, with the high marginal valuation of liquidity, that gets the vast majority of funds. Thus, the Million Baht program requires fewer funds to create the same benefit as the transfer program. Although current liquidity only flows to borrowers in the Million Baht Program, it is not only borrowers who benefit. Non-borrowers benefit by facing relaxed borrowing constraints in the future, which allows them to reduce their current buffer stocks.

Again, the village fund program is more cost effective than a simple transfer program because it directs funds toward those who have the highest marginal valuation of liquidity. As the size of the fund increases relative to the size of the village, consumption increases, and both the number of people who are currently credit constrained and the severity of their credit constraint fall. For this reason, the microfinance fund is even more cost effective, the smaller it is relative to the size of the village. This is shown in Figure 6, which shows the average cost of the Million Baht program in a village (relative to a transfer program) as a function of village size. Each point represents a village, and for all villages below the horizontal line at one, the Million Baht program is more cost-effective than the direct transfer.[38]

[38] The few exceptions in which the relative cost is above one are a result of the fact that credit may be used to simply pay off debt that would have defaulted (to no benefit of the household) and the within village inequality in permanent income across households. We assume that the Million Baht program increases ability to borrow in proportion to permanent income. Poor households receive smaller absolute increases in available liquidity than rich households, and so their equivalent transfer is less than average credit per household.

Figure 6: Relative Cost of Million Baht Fund as a Function of Cost/Household

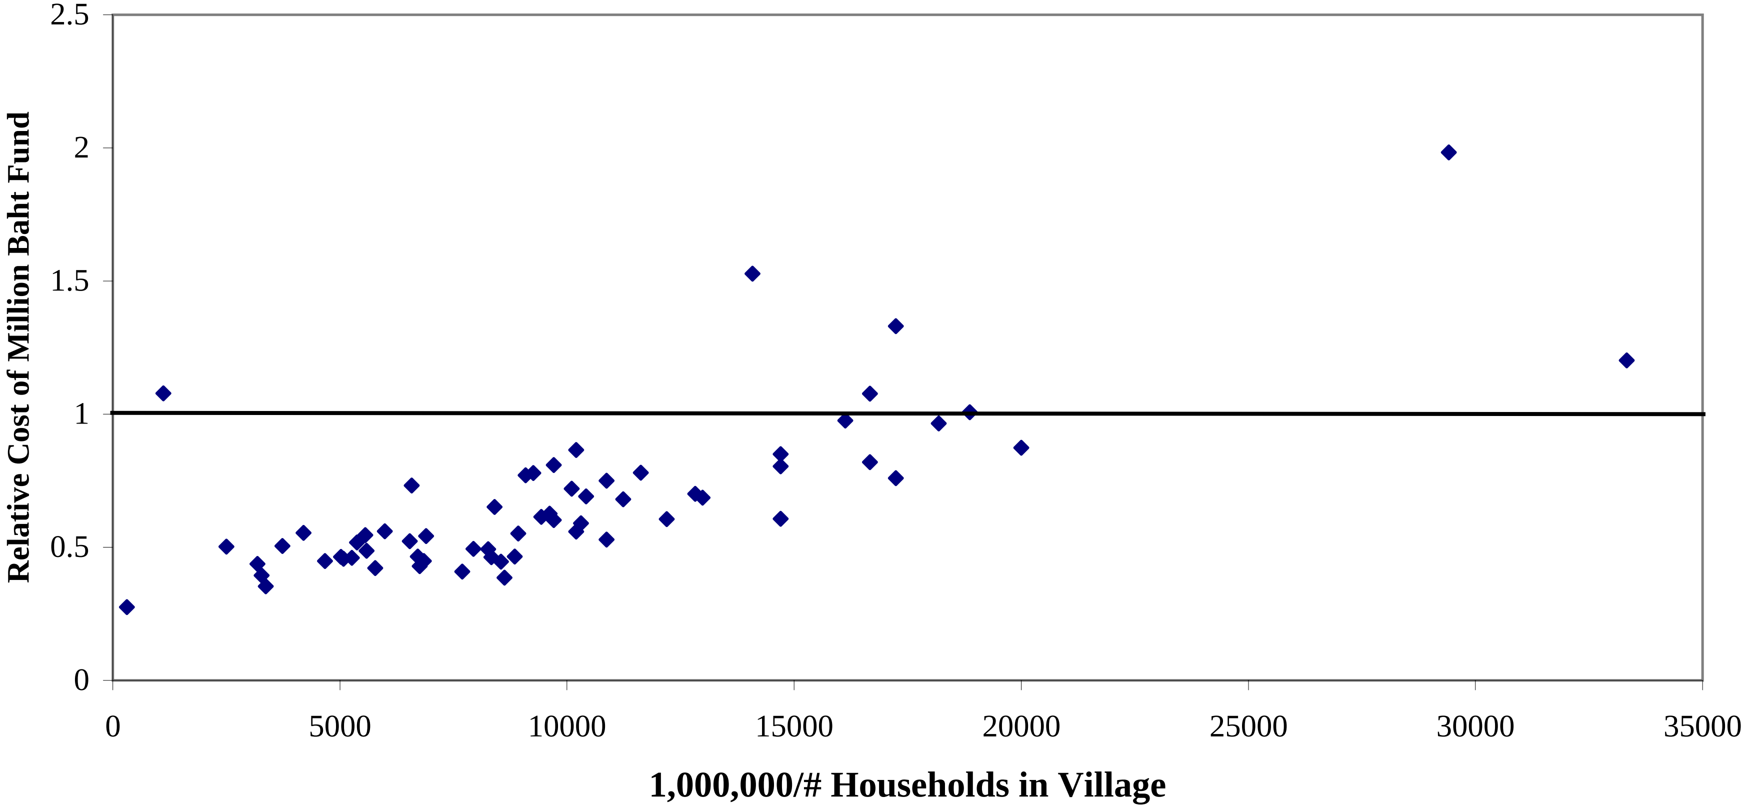

Notes: Each dot represents a village.

This impact and cost-effectiveness of the program is directly based on the assumption that the village fund is able to lower borrowing constraints, and, related, lower them permanently. If the borrowing constraint only held for outside funds, but households could borrow and lend unconstrained to one another in the village on their own (i.e., without the fund), then a direct transfer would also increase access to credit within the village. In contrast, we assume the village fund can lower the borrowing constraint, but direct transfers cannot. A household receiving a transfer in our scheme can lend to other households within the village or outside the village, but the borrowing households on the other side of such a transaction still face the borrowing constraint. Implicitly we assume the microfinance institution has a specific lending technology (perhaps multilateral enforcement) which allows it to lower the borrowing constraint beyond what would be available from lending between households in a village.

Similarly, if the fund only lowered the borrowing constraint for one period, it would have a much smaller impact on consumption and would yield a much smaller benefit.

5.3 Alternative Policy

Although the Million Baht Village Fund Program appears to be cost effective relative to a transfer in providing utility, the most discernible impacts are on consumption rather than investment. Although households may be quite happy with this, policy makers may not, since inducing investment is often a primary goal behind microfinance initiatives. 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. Nevertheless, such a policy would eliminate households like Household 1 in Figures 3 and 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 (averaging -0.87 vs. -0.70 in the actual policy) but only for those who borrow. The new range of borrowing constraints is from -0.19 to 1.85.

The policy does indeed lower the impact on consumption and increase the impact on investment. Pooling all 1000 simulated samples yields significant estimate for consumption that is lower than the actual million baht intervention (0.26 vs. 1.07). In contrast, it yields a much larger significant estimate for investment probability (4.98e-6 vs. and 1.87e-6), while yielding a comparably sized significant estimate for the level of investment. In a single sample of data only the positive impact on investment probability is typically significant, however. Clearly, the counterfactual policy shifts borrowing toward small investments rather than consumption by allowing credit only for those that invest. It offers less flexibility for constrained households who would rather not invest, but relaxes the borrowing constraint even more for investing households than the actual policy. In total, the counterfactual program is slightly more cost effective than the actual policy, costing on average 62 percent (vs. 66 percent for the actual policy) what a direct transfer program with equivalent benefit would cost.

6 Conclusions

We have developed a model of bufferstock saving and indivisible investment, and tested it using 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. The model predicts that intermediation, through microfinance, is a cost effective means of increasing available liquidity to constrained households. Finally, we have emphasized the relative strengths of a natural 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 constant across projects and households. Heterogeneity in project quality may be an important dimension for analysis, since microfinance may change the composition of project quality.

The analysis has also been purely partial equilibrium analysis of household behavior. For a large scale intervention, one might suspect that general equilibrium effects on income, rates of return to investment, and interest rates on liquidity may be important (see Kaboski and Townsend, 2006). Related, we did not consider the potential interactions between villagers or between villages, nor was the intermediation mechanism explicitly modeled. These are all avenues for future research.

References

[1] Ahlin, C. and Townsend, R. "Using Repayment Data to Test Across Models of Joint Liability Lending." mimeo, 2000

[2] Aiyagari, S. R., "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics,* 109 (August 1994): 659-684

[3] Alvarez, F. and Stokey, N. "Dynamic Programming with Homogeneous Functions", *Journal of Economic Theory*, 82 (1998): 167-189

[4] Aportela, F. "Effect of Financial Access on Savings by Low-Income People." mimeo, Massachusetts Institute of Technology, 1998

[5] Attanasio, O. and Davis, S. , "Relative wage movements and the distribution of consumption", *Journal of Political Economy* 104 (1996): 1227-1262

[6] Banerjee, A., Bardhan, P., Basu, K., Kanbur, R., and Mookherjee, D. "New Directions in Development Economics: Theory or Empirics: A Symposium." Institute for Economic Development Discussion Paper No. 153, August 2005

[7] Banerjee, A. and Newman, A. "Occupational Choice and the Process of Development." *Journal of Political Economy, (*April 1993*): 274-98*

[8] Banerjee, A. and Duflo, E. "Do Firms Want to Borrow? Testing Credit Constraints Using a Direct Lending Program." mimeo, Massachusetts Institute of Technology, May 2002

[9] Blossfeld, H. and Shavit,Y. *Persistent Inequality: Changing Educational Attainment in Thirteen Countries. (*Boulder, Co: Westview Press), 1993.

[10] Buera, F. J. "Persistency of Poverty, Financial Frictions, and Entrepreneurship." mimeo, Northwestern University, 2006

[11] Cameron, S. and Taber, C. "Borrowing Constraints and the Returns to Schooling." NBER Working Paper 7761, 2000

[12] Cameron, S. and Heckman, J. "The Dynamics of Educational Attainment for Black, Hispanic and White Males." *Journal of Political Economy,* 109 (June 2001): 455-499

[13] Cameron, S. and Heckman, J. "Life Cycle Schooling and Dynamic Selection Bias: Models and Evidence for Five Cohorts of American Males." *Journal of Political Economy,* 106 (April 1998): 262-333

[14] Card, D.(1999) "The Causal Effect of Education on Earnings.", in (O. Ashenfelter, and D. Card, eds.), *Handbook of Labor Economics*, Vol 3A. Amsterdam: Elsevier Science, North-Holland, pp.1801–63

[15] Carneiro, P. and Heckman, J. "The Evidence on Credit Constraints on Post-Secondary Schooling." *Economic Journal,* 112 (October 2001): 705-734

[16] Carroll, C. D. "Buffer Stock Saving and the Lifecycle Permanent Income Hypothesis." *Quarterly Journal of Economics, 107 (1) 1997*

[17] Carroll, C. D. "Buffer Stock Saving: Some Theory." Manuscript, The Johns Hopkins University, 2004

[18] Carroll, C. D. and Samwick, A. A. "The Nature of Precautionary Wealth." *Journal of Monetary Economics,* 40 (1) 1997

[19] Coleman, B. "Microfinance in Northeast Thailand: Who Benefits and How Much?" Asian Development Bank ERD Working Paper Series No. 9, 2001

[20] Coleman, B. "The Impact of Group Lending in Northeast Thailand." *Journal of Development Economics*, (October 1999): 105-141

[21] Conning, J. "Outreach, Sustainability and Leverage in Monitored and Peer-Monitored Lending." *Journal of Development Economics*, (October 1999): 105-141

[22] Evans, D. and Jovanovic, B. "An Estimated Model of Entrepreneurial Choice Under Liquidity Constraints." *Journal of Political Economy, (*August 1989*): 808-27*

[23] Fazzari, S., Hubbard, G. and Petersen, B. "Financing Constraints and Corporate Investment." *Brookings Papers on Economic Activity* (1988): 141-195

[24] Galor, O. and Zeira, J. "Income Distribution and Macroeconomics." *Review of Economic Studies*, 60 (1 1993): 35-52

[25] Gine, X. and R. Townsend "Evaluation of Financial Liberalization: A Model with Constrained Occupation Choice." mimeo, University of Chicago, 2001

[26] Gourinchas, P.-O., and Parker, J."Consumption over the Life Cycle." *Econometrica,* 71 (January 2002)*:*47-91.

[27] Heckman, J. J., Urzua, S., Vytlacil, E., 2004. "Understanding Instrumental Variables in Models with Essential Heterogeneity". Stanford University, Stanford, CA. http://www.stanford.edu/~vytlacil/

[28] Kaboski, J. and Townsend, R. "The Impacts of Credit on Village Economies." mimeo, 2007

[29] Kaboski, J. and Townsend, R. "Borrowing and Lending in Semi-Urban and Rural Thailand." mimeo, University of Chicago, 1998

[30] Kaboski, J. and Townsend, R. "Policies and Impact: An Analysis of Village Microfinance Institutions." mimeo, University of Chicago, 2003

[31] Kane, T. "College Entry by Blacks Since 1970: The Role of College Costs, Family Background, and the Returns to Education." *Journal of Political Economy,* 102 (October 2003): 878-911

[32] Karlan, D. and Zinman, J. "Expanding Credit Access: Using Randomized Supply Decisions To Estimate the Impacts." mimeo, August, 2006

[33] Keane, M. and Wolpin, K. "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment." *International Economic Review,* 42 (December 2001): 1051-1103

[34] Krueger, A. (2003). 'Inequality: too much of a good thing', in (J. Heckman and A. Krueger), *Inequality in America: What Role for Human Capital Policies*, Cambridge: MIT Press

[35] Lloyd-Ellis, H. and Bernhardt, D. "Enterprise, Inequality, and Economic Development." *Review of Economic Studies*, (January 2000): 147-68

[36] Miranda, M. and Fackler, P. (2002) *Applied Computational Economics and Finance* Cambridge: MIT Press, 512 pp.

[37] Ogaki, M. (1993) "Generalized Method of Moments: Econometric Applications," in *Handbook of Statistics, Vol. 11: Econometrics*, edited by G. S. Maddala, C. R. Rao, and H. D. Vinod, North-Holland, Amsterdam, 455-488.

[38] Owen, A. and Weil, D. "Intergenerational Earnings Mobility, Inequality, and Growth." NBER Working Paper 6070 (June 1997)

[39] Paulson, A. and Townsend, R. M. "Entrepreneurship and Liquidity Constraints in Rural Thailand." mimeo, 2000

[40] Rosenzweig, M. and Wolpin. K "Natural 'Natural Experiments' in Economics." *Journal of Economic Literature* 38 (December 2000): 827-874

[41] Samphantharak, K. *Internal Capital Markets in Business Groups.* Ph. D. thesis, University of Chicago: Chicago, 2003

[42] Samphantharak, K. and Townsend, R. ""Households as Corporate Firms: Constructing Financial Statements from Integrated Household Surveys" mimeo, 2006

[43] Todd, P. and Wolpin, K. "Using Experimental Data to Validate a Dynamic Behavioral Model of Child Schooling and Fertility: Assessing the Impact of a School Subsidy Program in Mexico." working paper, University of Pennsylvania, 2003

[44] Townsend, R., principal investigator with Paulson, A., Sakuntasathien, S., Lee, T. J., and Binford, M., Questionnaire Design and Data Collection for NICHD Grant *Risk, Insurance and the Family* and NSF grants, The University of Chicago, 1997

[45] Wright, M. "Investment and Growth with Limited Commitment." mimeo, October 2002

[46] Zeldes, S. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, 97 (March 1989): 305-46