Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes*

Alexander Karaivanov
Simon Fraser University

Robert M. Townsend
MIT

February, 2010

Abstract

We formulate and solve a range of dynamic models of information-constrained credit that allow for moral hazard and unobservable investment. We compare them to full insurance and exogenously incomplete financial regimes (autarky, saving only, and borrowing and lending in a single asset). We develop computational methods based on mechanism design, linear programming, and maximum likelihood to estimate, compare, and statistically test these alternative theoretical models of financial constraints. Our methods work with both cross-sectional and panel data and allow for measurement error and unobserved heterogeneity. Empirically, we find that using consumption, income, and investment data jointly, or using intertemporal data, improves the researcher's ability to distinguish across the financial regimes relative to using consumption or investment data alone, especially in the presence of high measurement error. We estimate the models using data on Thai households running small businesses. We find that, overall, the borrowing and saving only regimes provide the best fit using joint data on consumption, investment, and income. However, there is evidence that family networks are helpful in consumption smoothing as in a moral hazard constrained regime and that there are regional differences in the best fitting regime (higher-wealth, near Bangkok vs. rural agricultural areas).

Keywords: financial constraints, mechanism design, structural estimation and testing

JEL Classifications: C61, D82

*We thank H. Cole, L. Hansen, A. Hortacsu, P. Lavergne, E. Ligon, T. Magnac, D. Neal, W. Rogerson, S. Schulhofer-Wohl, K. Wolpin, Q. Vuong, and audiences at Berkeley, Chicago, Duke, MIT, Northwestern, Penn, Santa Barbara, Toulouse, Stanford, the SED Conference, the European Meeting of the Econometric Society, and the Vienna Macroeconomics Workshop for many insightful comments and suggestions. Financial support from the NSF, NICHD, the John Templeton Foundation, SSHRC, and the Bill and Melinda Gates Foundation through a grant to the Consortium on Financial Systems and Poverty at the University of Chicago is gratefully acknowledged.

1 Introduction

We compute, estimate, and contrast the consumption and investment behavior of risk averse households running small nonfarm and farm businesses under various financial market environments, including exogenously incomplete settings (autarky, savings only, non-contingent borrowing) and endogenously information-constrained settings (moral hazard with observed or unobserved investment), both relative to full insurance. We analyze in what circumstances these financial regimes can be distinguished in consumption and income or investment and income data or both. More generally, we develop and apply methods for empirical estimation of dynamic mechanism design models and test the various models against each other using both data simulated from the models themselves and actual data on Thai households.

With few exceptions, the existing literature maintains a dichotomy, also embedded in the national accounts: households are consumers and suppliers of market inputs, whereas firms produce and hire labor and other factors. This gives rise, on the one hand, to a large literature which studies household consumption smoothing. There is voluminous work estimating the permanent income model, the full risk sharing model, buffer stock models (Zeldes, 1989; Deaton and Laroque, 1996) and, lately, models with private information (Phelan 1994; Ligon, 1998; Werning, 2001) or limited commitment (Ligon, Thomas and Worrall, 2002).

On the other hand, the consumer-firm dichotomy gives rise to an equally large literature on investment. For example, there is the adjustment costs approach of Abel and Blanchard (1983) and Bond and Meghir (1994) among many others. In industrial organization, Hopenhayn (1992) and Ericson and Pakes (1995) model the entry and exit of firms with Cobb-Douglas or CES production technologies where investment augments capital with a lag and output produced from capital, labor and other factors is subject to factor-neutral Markov technology shocks. Mostly, firms are modeled as risk neutral maximizers of expected discounted profits or of dividends to owners. There are also works attempting to explain stylized facts on firm growth, with higher mean growth and variance in growth for small firms, e.g. Cooley and Quadrini (2001), among others. The more recent works by Albuquerque and Hopenhayn (2004) and Clementi and Hopenhayn (2006) introduce either private information or limited commitment but maintain risk neutrality[1]. Here we set aside for the moment the issues of heterogeneity in technologies and firm growth and focus on a benchmark with financial constraints, investment and consumption data thinking of households as firms.

The literature that is closest to our paper, and complementary with what we are doing, features risk averse households as firms but largely *assumes* that certain markets or contracts are missing. For example, Cagetti and De Nardi (2006) follow Aiyagari (1994) in their study of inequality and assume that labor income is stochastic and uninsurable, while Angeletos and Calvet (2007) and Covas (2006) in their work on buffer stock motives and macro savings rates feature uninsured entrepreneurial risk. In the asset pricing vein, Heaton and Lucas (2000) model entrepreneurial investment as a portfolio choice problem, assuming exogenously incomplete markets in the tradition of Geanakoplos and Polemarchakis (1986) or Zame (1993). The methods of our paper might indicate how to build upon these papers, possibly with alternative assumptions on the financial underpinnings.

Indeed, this literature begs the question of how good an approximation are the various assumptions on financial regimes, different across the different papers. That is, what would be a reasonable assumption for the financial regime if that part too were taken to the data? We take this view below

[1] Applied general equilibrium models feature both consumption and investment in the same context, as Rossi-Hansberg and Wright (2007), but there the complete markets hypothesis justifies, within the model, a separation of the decisions of households from the decisions of firms. Alem and Townsend (2009) provide an explicit derivation of full risk sharing with equilibrium stochastic discount factors, rationalizing the apparent risk neutrality of households as firms making investment decisions.

to see how far we can get. For example, the adjustment costs investment literature may be picking up constraints implied by financing, not adjustment costs per se. The pecking order investment literature (Myers and Majuf, 1984) simply assumes that internally generated funds are least expensive, followed by debt, and finally equity, discussing wedges and distortions. Berger and Udell (2002) also have a long discussion in this spirit, of small vs. large firm finance. They point out that small firms seem to be informationally opaque[2] yet receive funds from family, friends, angels, or venture capitalists, leaving open the nature of the overall financial regime. The empirical work of Fazzari, Hubbard and Petersen (1988) picks up systematic distortions for small firms, but, again, the nature of the credit market imperfection is not modeled, leading to criticisms of their interpretation of cash flow sensitivity tests[3] (Kaplan and Zingales, 2000).

Our methods follow logically from Paulson, Townsend and Karaivanov (2006), (hereafter PTK), where we model, estimate, and test whether moral hazard or limited liability is the predominant financial obstacle causing the observed positive monotonic relationship between initial wealth and subsequent decision to enter into business. Buera and Shin (2007) extend this to endogenous savings decisions in a model with limited borrowing. Here, again, we abstract from occupational choice and focus much more on the dynamics. The recent work of Schmid (2008) is also an effort to estimate a dynamic model of financial constraints using regressions on model data, not maximum likelihood as here. Finally, Kinnan (2009) uses non-parametric methods to test inverse Euler equations or other implications of moral hazard, limited commitment, and unobserved-output financial regimes.

We naturally analyze the advantages of using a combination of data on consumption and the smoothing of income shocks with data on the smoothing of investment from cash flow fluctuations, in effect filling the gap in the dichotomy of the literature. In estimating both exogenously incomplete and endogenous information-constrained regimes we also break new ground. The only other similar efforts of which we are aware are Meh and Quadrini (2006), who compare and contrast numerically a bond economy to an economy in which unobserved diversion of capital creates an incentive constraint, and Attanasio and Pavoni (2008) who estimate and compare the extent of consumption smoothing in the permanent income model to that in a moral hazard model with hidden savings (see also Karaivanov, 2008).

In this paper we focus on whether, and in what circumstances, it is possible to distinguish financial regimes, depending on the data used. To that end, we first perform tests in which we have full control, that is, we know what the financial regime really is, using data generated from the model. Our paper is thus both a conceptual and methodological contribution. We show how all the financial regimes can be formulated as linear programming problems, often of large dimension, and how likelihood functions, naturally in the space of probabilities/lotteries or histograms, can be estimated. We allow for measurement error, the need to estimate the underlying distribution of unobserved state variables, and the use of data from transitions, before households reach steady state.

When using model-generated data we find that our ability to distinguish between the alternative financial regimes naturally depends on both the type of data used and the amount of measurement error. With low measurement error we are able to distinguish between almost all regime pairs. As expected, however, higher level of measurement error in the data reduces the power of our model comparison tests to the extent that some regimes cannot be distinguished from the data-generating baseline as well as from each other. For example, using investment/income data, or consumption/income data we cannot distinguish between the moral hazard and full information regimes when there is high

[2] Bitler, Moskowitz and Vissing-Jorgensen (2005) argue likewise that agency considerations play important role.

[3] Under the null of complete markets there should be no significant cash flow variable in investment decisions, but the criticism is that when the null is rejected, one cannot infer the degree of imperfection of financial markets from the magnitude of the cash flow coefficient. One needs to explicitly model the financial regime in order to make an inference, which is what we do in this paper.

measurement error. Using joint data on consumption, investment, and income markedly improves the ability to distinguish across the regimes when there is high measurement error. We also incorporate intertemporal data from the model through a panel which also significantly improves the ability to distinguish the regimes, relative to when using single cross-sections. The simulated data results are shown to be robust to various modifications of the baseline runs – no measurement error, different grid sizes, allowing for heterogeneity in productivity, and using data-generating parameters estimated from Thai data.

Additionally, we do take the next step and apply our methods to a featured emerging market economy, Thailand, to make the point that what we offer is a feasible, practical approach to real data when the researcher aims to provide insights on the source and nature of financial constraints. We chose Thailand for two main reasons. First, our data source (the Townsend Thai surveys) includes panel data on both consumption and investment and this is rare. We can thus see if the combination of consumption and investment data really helps in practice. Second, we also learn about potential next steps in modeling financial regimes. We know in particular, from other work with these data, that consumption smoothing is quite good, that is, it is sometimes difficult to reject full insurance, in the sense that the coefficient on idiosyncratic income, if significant, is small (Chiappori, Schulhofer-Wohl, Samphantharak, and Townsend, 2008). We also know that investment is sensitive to income, especially for the poor, but on the other hand this is to some extent overcome by family networks (Samphantharak and Townsend, 2009).

While we keep these data features in mind, we remain neutral in what we expect to find in terms of the best-fitting theoretical model. Hence, we test the full range of regimes (from autarky to full information) against the data. We are interested in how these same data look when viewed jointly though the lens of each of the various financial regimes modelled here. We also want to be assured that our methods which use grid approximations, measurement error, estimation of unobserved distribution of utility promises, and transition dynamics are, as a practical matter, applicable to actual data. This is our primary intent, to create an operational methodology for estimating and comparing across different dynamic models of financial regimes that can be taken to data from various sources.

We find that by and large out methods work with the Thai data and we obtain results consistent with those for model-simulated data. Using combined data on consumption, investment, and income, or using two-year panel data improves our ability to distinguish the regimes. In terms of the regime that fits the Thai data best, we echo previous work which finds that investment is not smooth and can be sensitive to cash flow fluctuations. Indeed, we find that investment and income data alone are most consistent with the borrowing and lending or savings only regimes, with ties depending on the specification. Results using combined consumption and investment data also lend support to the best fit of the savings and non-contingent borrowing regimes, and, in one instance, the moral hazard regime with unobserved investment.

We also echo previous work which finds that full risk sharing is rejected, but not by much, and indeed find that the moral hazard regime is not inconsistent with the income and consumption data alone, but often statistically tied with borrowing and lending or savings only, depending on the specification. We find some evidence that family networks move households toward less constrained regimes in regards to their consumption smoothing from income too. Stratifying the data by region – the richer, industrializing Central region vs. the poorer, predominantly agricultural Northeast shows evidence for regional differences in the best-fitting regime. Using consumption and income data alone reveals the mechanism design regimes as best fitting in the Northeast while borrowing and saving only fit best in the Central region while using combined consumption and investment data pins down the borrowing regime as best-fitting in the Central but cannot reject moral hazard (tied with borrowing and saving only) in the Northeast.

We also perform a range of additional runs that confirm the robustness of our results – imposing risk neutrality, imposing no measurement error, allowing for quadratic adjustment costs in investment, allowing for limited commitment in the moral hazard and full information regimes, different grid sizes, running on data cleaned from household fixed effects, and alternative assets and income definitions.

2 Theory

2.1 Basic Setting

Consider an economy of agents heterogeneous in their initial endowments (assets), k_0 of a single good that can be used for both consumption and investment. The agents live for T periods, where for the most part we take T to infinity. They can contract with a financial intermediary, entering into saving, debt, or insurance arrangements. Below we characterize the optimal dynamic financial contracts that arise between the agents and the intermediary in various information and credit access regimes.

Agents are risk averse and have time-separable preferences defined over consumption, c, and labor effort, z represented by $U(c, z)$ where $U_1 > 0$, $U_2 < 0$. They discount future utility using a discount factor $\beta \in (0, 1)$. We assume that c and z belong to the finite discrete sets (grids) C and Z respectively. The agents have access to a stochastic output technology, $P(q|z, k) : Q \times Z \times K \to [0, 1]$ which gives the probability of obtaining output/income, q from effort level, z and capital level[4], k. The sets Q and K are also finite and discrete[5]. In all financial regimes we study, output q is assumed to be observable and verifiable. However, one or both of the inputs, k and z may be unobservable to third parties, resulting in moral hazard and/or adverse selection problems. Capital, k depreciates at rate δ. Depending on the application we have in mind, the lowest k level ($k = 0$) could be interpreted as a 'worker' occupation (similar to PTK, 2006) or as firm exit but we do not do much with this in this paper.

The financial intermediary is risk neutral and has access to an outside credit market with opportunity cost of funds R. Using the linear programming approach of Prescott and Townsend (1984) and Phelan and Townsend (1991), we model financial contracts as probability distributions (lotteries) over assigned or implemented allocations of consumption, output, effort, and investment (see below for details). There are two possible interpretations. First, one can think of the intermediary (the principal) contracting with a single agent/firm at a time, in which case the contracts specify mixed strategies over allocations. Alternatively, one can think of a principal contracting with a continuum of agents, so that the optimal contract specifies the fraction of agents of given 'type' (state) that receive a particular deterministic allocation. It is further assumed that there are no aggregate shocks, there are no technological links between the agents, and the agents cannot collude. Finally, the intermediary and the agents are assumed to be able to fully commit to the ex-ante (constrained) optimal contract (later on we show how we can relax this assumption).

2.2 Information and Credit Access Regimes

We write down a dynamic linear programming problem determining the (constrained) optimal contract in each regime. To characterize the optimal contracts in the settings with incomplete information we invoke the revelation principle and study direct mechanisms in which the agents announce their types

[4]We can also easily incorporate heterogeneity in entrepreneurial ability across agents as in PTK (2006), for instance by adding a scaling factor in the production function. Note also that q, as defined, can be interpreted as income net of payments for any hired inputs other than z and k.

[5]This can be either a technological or computational assumption depending on the application.

truthfully[6].

The financial regimes we study can be classified into two groups. The first group are regimes with exogenously incomplete markets: *autarky (A)*, *savings only (S)*, and *borrowing and lending (B)*. In these regimes the feasible financial contracts can only take a specific, exogenously given form (no access to financial markets, a deposit/storage contract, or a non-contingent debt contract, respectively). In contrast, the second group of financial contracts we study are endogenously determined as the solutions to dynamic mechanism design problems, subject to information and incentive constraints. We look at two such endogenously incomplete markets regimes – *moral hazard (MH),* in which agents' effort is unobserved but capital and investment are observed, and *moral hazard with unobserved investment (UI)*, in which both effort and investment are unobservable. All regimes are compared to the *full information (FI)* benchmark (the complete markets first best).

2.2.1 Exogenously Incomplete Markets Regimes

Autarky

In autarky the agent has no access to financial intermediation or storage. The timeline is as follows. The agent starts the current period with initial capital k which he invests into production. At this time he also employs his effort z. At the end of the period output q is realized, the agent decides on the next period capital level, $k' \in K$, (we allow both downward and upward adjustments), and consumes $c \equiv (1-\delta)k + q - k'$. Capital, k is the single state variable in the recursive formulation of the agent's problem. The latter is relatively simple and can be solved by standard dynamic programming techniques. However, to be consistent with the solution method that we use for the mechanism design regimes, where non-linear techniques may be inapplicable due to non-convexities introduced by the incentive and truth-telling constraints (more on this below), we formulate the agent's problem in autarky (as well as all others) as a dynamic linear programming problem in the joint probabilities (lotteries) over all possible allocations (q, z, k') given k:

$$v(k) = \max_{\pi(q,z,k'|k)} \sum_{Q\times Z\times K} \pi(q, z, k'|k)[U((1-\delta)k + q - k', z) + \beta v(k')] \tag{1}$$

The maximization of the agent's value function[7], $v(k)$ in (1) is subject to a set of constraints on the choice variables, π. First, $\forall k \in K$, the probabilities $\pi(.)$ have to be Bayes-consistent with the technology-determined probability distribution over outputs, $P(q|z, k)$:

$$\sum_{K} \pi(\bar{q}, \bar{z}, k'|k) = P(\bar{q}|\bar{z}, k) \sum_{Q\times K} \pi(q, \bar{z}, k'|k) \text{ for all } (\bar{q}, \bar{z}) \in Q \times Z \tag{2}$$

Second, given that the $\pi(.)$'s are probabilities, they must satisfy $\pi(q, z, k'|k) \geq 0$ (non-negativity) $\forall (q, z, k') \in Q \times Z \times K$, and 'adding-up',

$$\sum_{Q\times Z\times K} \pi(q, z, k'|k) = 1 \tag{3}$$

The policy variables $\pi(q, z, k'|k)$ that solve the above maximization problem determine the agent's optimal effort and (state-contingent) investment in autarky for each k.

[6] The proofs that the optimal contracting problems can be written in a recursive form and that the revelation principle applies follow from Doepke and Townsend (2006) and hence are omitted.

[7] The variable k' takes values over the same set K as k, thus in (1) and later on, K under the summation sign refers to k' and not k and similarly for the set W below.

Saving Only / Borrowing

In this financial environment the agent is able to either only save (i.e., accumulate and run down a buffer stock), the *saving only (S)* regime, or borrow and save, the *borrowing (B)* regime, through a competitive financial intermediary. That is, the agent can use saving or debt in a risk free asset to smooth his consumption or investment, on top of what he could do under autarky. More specifically, if the agent borrows (saves) an amount b, next period he has to repay (collect) the amount Rb, independent of the state of the world. Involuntary default is ruled out by assuming that the principal refuses to lend to a borrower who is at risk of not being able to repay[8]. By shutting down all contingencies in the debt contracts we aim for better differentiation from the mechanism design regimes. Like all other variables, b is assumed to take values on the finite grid B. By convention, a negative value of b represents savings, i.e., in the S regime the upper bound of the grid B is zero, while in the B regime the upper bound is positive.

The timeline is as follows: the agent starts the current period with capital k and uses it in production together with effort z. At the end of the period, output q is realized, the agent repays Rb, and borrows or saves $b' \in B$. He also puts aside (invests in) next period's capital, k' and consumes $c \equiv (1-\delta)k + q + b' - Rb - k'$. The two 'assets', k and b are thus freely convertible into one another.

The problem of an agent with current capital stock k and debt/savings level b in the S or B regime can be written recursively as:

$$v(k,b) = \max_{\pi(q,z,k',b'|k,b)} \sum_{Q\times Z\times K\times B} \pi(q,z,k',b'|k,b)[U((1-\delta)k+q+b'-Rb-k',z)+\beta v(k',b')] \tag{4}$$

subject to the Bayes-consistency and adding-up constraints analogous to (2) and (3), and subject to $\pi(q,z,k',b'|k,b) \geq 0 \ \forall (q,z,k',b') \in Q\times Z\times K\times B$.

2.2.2 Mechanism Design Regimes

Full Information

With full information (FI) the principal observes and can contract upon agent's effort and investment. We write the corresponding dynamic principal-agent problem as an extension of Phelan and Townsend (1991) with capital accumulation. As is standard in such settings, to obtain a recursive formulation we use an additional state variable – *promised utility*, w, which belongs to the discrete set, W (Spear and Srivastava, 1987). The optimal full-information contract for an agent with current promised utility w and capital k thus consists of effort and capital levels $z, k' \in Z \times K$, next period's promised utility $w' \in W$, and a transfer τ belonging to the discrete set T. A positive value of τ denotes a transfer from the principal to the agent. The timing of events is the same as before, with the addition that transfers occur after output is observed.

Following Phelan and Townsend (1991), the set of promised utilities W has a lower bound, $w_{\min}$ which corresponds to assigning forever the lowest possible consumption, $c_{\min}$ (obtained from the lowest $\tau \in T$ and the highest $k' \in K$) and the highest possible effort, $z_{\max}$. The set's upper bound, $w_{\max}$ in turn corresponds to promising the highest possible consumption, $c_{\max}$ and the lowest possible effort forever:

$$w^{FI}_{\min} = \tfrac{U(c_{\min},z_{\max})}{1-\beta} \text{ and } w^{FI}_{\max} = \tfrac{U(c_{\max},z_{\min})}{1-\beta} \tag{5}$$

The principal's objective function, $V(w,k)$ when contracting with an agent at state (w,k) maximizes the expected value of output net of transfers plus the discounted value of future outputs and transfers.

[8] Computationally, this is achieved by assigning a very low utility value in such cases.

We write the mechanism design problem solved by the optimal contract as a linear program in the joint probabilities, $\pi(\tau, q, z, k', w'|w, k)$ over all possible allocations (τ, q, z, k', w'):

$$V(w,k) = \max_{\{\pi(\tau,q,z,k',w'|w,k)\}} \sum_{T\times Q\times Z\times K\times W} \pi(\tau,q,z,k',w'|w,k)[q-\tau+(1/R)V(w',k')] \tag{6}$$

The maximization in (6) is subject to the Bayes-consistency and adding-up constraints:

$$\sum_{T\times K\times W} \pi(\tau,\bar{q},\bar{z},k',w'|w,k) = P(\bar{q}|\bar{z},k) \sum_{T\times Q\times K\times W} \pi(\tau,q,\bar{z},k',w'|w,k) \text{ for all } (\bar{q},\bar{z}) \in Q\times Z, \tag{7}$$

$$\sum_{T\times Q\times Z\times K\times W} \pi(\tau,q,z,k',w'|w,k) = 1, \tag{8}$$

as well as non-negativity: $\pi(\tau, q, z, k', w'|w, k) \geq 0$ for all $(\tau, q, z, k', w') \in T \times Q \times Z \times K \times W$.

The optimal FI contract must also satisfy an additional, *promise keeping* constraint which reflects the principal's commitment ability and ensures that the agent's expected utility equals his current promise, w:

$$\sum_{T\times Q\times Z\times K\times W} \pi(\tau,q,z,k',w'|w,k)[U(\tau+(1-\delta)k-k',z)+\beta w'] = w \tag{9}$$

By varying the initial value of w one can trace the whole Pareto frontier for the principal and the agent. The optimal FI contract is the vector of allocation probabilities, $\pi^*(\tau, q, z, k', w'|w, k)$ that maximizes (6) subject to (7), (8) and (9).

Moral Hazard

In the moral hazard (MH) regime the principal can still observe and control the agent's capital and investment (k and k'), but he can no longer observe or verify the agent's effort, z. This results in a moral hazard problem, that is, the principal must induce effort by the agent. With capital observed and controlled, k can be also interpreted as endogenous collateral. The timing is the same as in the FI regime. However, the optimal MH contract $\pi(\tau, q, z, k', w'|w, k)$ must now satisfy an incentive-compatibility constraint[9] (ICC), in addition to (7)-(9). The ICC states that, given the agent's state (w, k), a recommended effort level $\bar{z}$, capital k', and transfer τ, the agent must not be able to achieve higher expected utility by deviating to an alternative effort level $\hat{z}$, i.e., $\forall(\bar{z}, \hat{z}) \in Z \times Z$:

$$\sum_{T\times Q\times W'\times K'} \pi(\tau,q,\bar{z},k',w'|w,k)[U(\tau+(1-\delta)k-k',\bar{z})+\beta w'] \geq$$

$$\geq \sum_{T\times Q\times W'\times K'} \pi(\tau,q,\bar{z},k',w'|w,k)\frac{P(q|\hat{z},k)}{P(q|\bar{z},k)}[U(\tau+(1-\delta)k-k',\hat{z})+\beta w'] \tag{10}$$

Apart from the extra ICC constraint (10), the MH regime differs from the FI regime in the set of feasible promised utilities, W. In particular, the lowest possible promise under moral hazard is no longer the value $w^{FI}_{\min}$ from (5). Indeed, if the agent is assigned minimum consumption forever, he would not supply effort above the minimum possible. Thus, the feasible range of promised utilities, W for the MH regime is bounded by:

$$w^{MH}_{\min} = \tfrac{U(c_{\min},z_{\min})}{1-\beta} \text{ and } w^{MH}_{\max} = \tfrac{U(c_{\max},z_{\min})}{1-\beta} \tag{11}$$

[9] For the details on the derivation of the ICC in the linear programming approach, see Prescott and Townsend (1984). The key term is the 'likelihood ratio', $\frac{P(q|\hat{z},k)}{P(q|\bar{z},k)}$ which reflects the fact that, if the agent deviates, he changes the probability distribution of output.

Intuitively, the principal cannot promise a slightly higher consumption in exchange for much higher effort such that agent's utility falls below $w_{\min}^{MH}$ since this is not incentive compatible. If the agent does not follow the principal's recommendations but deviates to $z_{\min}$ the worst punishment he can receive is $c_{\min}$ forever.

The constrained optimal contract in the moral hazard regime is the solution to the linear program defined by (6)–(10). The contract features incomplete consumption smoothing and intertemporal tie-ins, i.e., it is not a repetition of the optimal one-period contract (Townsend, 1982).

Moral Hazard with Unobserved Capital and Investment

Now suppose effort is still unobservable but, in addition, the principal also cannot observe the agent's current capital, k and the level planned for next period, k'. This adds a dynamic adverse selection problem (because of the unobserved state, k) to the moral hazard problem arising from the two unobserved actions, z and k'.

The agent sends a message about his capital level k to the principal who offers him a contract conditional on the message which consists of transfer τ, recommended effort z, investment k', and future promised utility. Because of the dynamic adverse selection problem in the state k, following Fernandes and Phelan (2000) and Doepke and Townsend (2006), instead of the scalar promise w from the MH regime, the proper state variable in the recursive representation is a promised utility *schedule*[10], $\mathbf{w} \equiv \{w(k_1), w(k_2), ..., w(k_{\#K})\} \in \mathbf{W}$, where k_1, k_2, etc. are the elements of the grid K. The #K-dimensional set $\mathbf{W}$ is endogenously determined (not all promise-assets combinations are feasible) and must be iterated upon together with the value and policy functions (Abreu, Pierce and Stacchetti, 1990).

The computational method we use to solve for the optimal contract in this unobserved investment (UI) regime requires separability in consumption and leisure, $U(c, z) = u(c) - d(z)$ (note, this was not needed for the MH, FI, or the exogenously incomplete regimes). The separability allows us to split each time period into two sub-periods and use dynamic programming within the time periods. This helps keep dimensionality in check, since the resulting sub-problems are of much lower dimension. The first sub-period includes the announcement of k by the agent, the principal's effort recommendation z, the agent's actual effort supply, and the realization of the output q. The second sub-period includes the transfer, the investment recommendation, and the agent's consumption and actual investment decisions. To tie the two sub-periods together, we introduce the extra variables, $\mathbf{w}_m$ that we call 'interim promised utility' – a representation of the agent's expected utility from the end of sub-period 1 (that is, from the middle of the period) onwards. The interim promised utility is a *schedule* (vector), $\mathbf{w}_m = \{w_m(k_1), w_m(k_2), ...\} \in \mathbf{W}_m$, similar to $\mathbf{w}$. Like $\mathbf{W}$, the set $\mathbf{W}_m$ is endogenously determined along the value function iteration.

The first sub-period problem for computing the optimal contract with an agent who has announced k and has been promised $\mathbf{w}$ is:

Program UI1

$$V(\mathbf{w}, k) = \max_{\{\pi(q,z,\mathbf{w}_m|\mathbf{w},k)\}} \sum_{Q\times Z\times \mathbf{W}_m} \pi(q, z, \mathbf{w}_m|\mathbf{w},k)[q + V_m(\mathbf{w}_m, k)] \tag{12}$$

The choice variables are the probabilities over allocations $(q, z, \mathbf{w}_m) \in Q \times Z \times \mathbf{W}_m$. The function $V_m(\mathbf{w}_m, k)$ is defined in the second sub-period problem (see Program UI2 below). The maximization in (12) is subject to the following constraints. First, the optimal contract must deliver the promised

[10] The reason why utility promises must, in general, depend on the state k, is the different incentives of agents entering next period with different capital levels (see Kocherlakota, 2004 for a detailed discussion).

utility on the equilibrium path, $w(k)$:

$$\sum_{Q \times Z \times \mathbf{W}_m} \pi(q, z, \mathbf{w}_m | \mathbf{w}, k)[-d(z) + w_m(k)] = w(k) \tag{13}$$

The utility from consumption and discounted future utility are incorporated in w_m. Second, as in the MH regime, the optimal contract must satisfy incentive compatibility in effort. That is, $\forall (\bar{z}, \hat{z}) \in Z \times Z$:

$$\sum_{Q \times \mathbf{W}_m} \pi(q, \bar{z}, \mathbf{w}_m | \mathbf{w}, k)[-d(\bar{z}) + w_m(k)] \geq \sum_{Q \times \mathbf{W}_m} \pi(q, \bar{z}, \mathbf{w}_m | \mathbf{w}, k) \frac{P(q|\hat{z}, k)}{P(q|\bar{z}, k)}[-d(\hat{z}) + w_m(k)] \tag{14}$$

Third, since the state k is private information, the agent needs incentives to reveal it truthfully. On top of that, the agent can presumably consider joint deviations in his announcement, k and his effort choice, z. To prevent such joint deviations, truth-telling must be ensured to hold regardless of whether the agent decides to follow the effort recommendation, z or considers a deviation to another effort level $\delta(z) \in Z$, where $\delta(z)$ denotes all possible mappings from recommended to actual effort, that is, from the set Z to itself. Such behavior is ruled out by imposing the following 'truth-telling' constraints, which must hold for all $\hat{k} \neq k$ and $\delta(z)$:

$$w(\hat{k}) \geq \sum_{Q \times Z \times \mathbf{W}_m} \pi(q, z, \mathbf{w}_m | \mathbf{w}, k) \frac{P(q|\delta(z), \hat{k})}{P(q|z, k)}[-d(\delta(z)) + w_m(\hat{k})] \tag{15}$$

In words, an agent who actually has $\hat{k}$ but considers announcing k triggering $\pi(.|\mathbf{w}, k)$ should find any such deviation unattractive. There are $(\#K - 1)\#Z^{\#Z}$ such constraints in total. Finally, the contract must satisfy the already familiar Bayes-consistency, adding-up, and non-negativity constraints for the probabilities $\pi(q, z, \mathbf{w}_m | \mathbf{w}, k)$.

To solve Program UI1, we first need to compute the principal's 'interim value function' $V_m(\mathbf{w}_m, k)$. The proper state variables are the schedule, $\mathbf{w}_m$ of interim utilities for each $k \in K$ and the agent's actual announcement k. Constraints will introduce truth-telling and obedience in the second-stage program. We need to ensure that, when deciding on k', the agent cannot obtain more than his interim utility, $w_m(k)$ for any announcement k.

Program UI2

$$V_m(\mathbf{w}_m, k) = \max_{\{\pi(\tau, k', \mathbf{w}' | \mathbf{w}_m, k)\}, \{v(\hat{k}, k', \tau | \mathbf{w}_m, k)\}} \sum_{T \times K \times \mathbf{W}} \pi(\tau, k', \mathbf{w}' | \mathbf{w}_m, k)[-\tau + (1/R)V(k', \mathbf{w}')] \tag{16}$$

Note that, in addition to the allocation lotteries, $\pi(\tau, k', \mathbf{w}' | \mathbf{w}_m, k)$ we introduce additional choice variables, $v(\hat{k}, k', \tau | \mathbf{w}_m, k)$ that we refer to as 'utility bounds' (see Prescott, 2003 for details). These bounds specify the maximum expected utility that an agent who is actually at $\hat{k}$ receiving transfer τ and an investment recommendation k' could obtain by reporting k and doing $\hat{k}'$. This translates into the constraint:

$$\sum_{\mathbf{W}} \pi(\tau, k', \mathbf{w}' | \mathbf{w}_m, k)[u(\tau + (1 - \delta)\hat{k} - \hat{k}') + \beta w'(\hat{k}')] \leq v(\hat{k}, k', \tau | \mathbf{w}_m, k) \tag{17}$$

which must hold for all possible combinations $\tau, k', \hat{k} \neq k$, and $\hat{k}' \neq k'$. To ensure truth-telling, the interim utility $w_m(\hat{k})$ that the agent obtains in the second sub-period by reporting k when the true state is $\hat{k}$, must satisfy, for all $k, \hat{k}$:

$$\sum_{T \times K} v(\hat{k}, k', \tau | \mathbf{w}_m, k) \leq w_m(\hat{k}) \tag{18}$$

The two sets of constraints, (17) and (18) rule out any joint deviations in the report k and the action k'. Finally, by definition, the interim utility must satisfy:

$$w_m(k) = \sum_{T \times K \times \mathbf{W}} \pi(\tau, k', \mathbf{w}'|\mathbf{w}_m, k)[u(\tau + (1-\delta)k - k') + \beta w'(k')] \tag{19}$$

and the probabilities $\pi(\tau, k', \mathbf{w}'|\mathbf{w}_m, k)$ must satisfy non-negativity and adding-up.

3 Computation

3.1 Numerical Solution Methods

We solve the dynamic programs for all financial regimes numerically[11]. Specifically, we use the linear programming (LP) methods developed by Prescott and Townsend (1984), Phelan and Townsend (1991) and PTK (2006). An alternative to the LP methodology in the literature is the 'first order approach' (Rogerson, 1985), used for instance by Abraham and Pavoni (2008), whereby the incentive constraints are replaced by their first order conditions[12]. A problem with that approach arises due to non-convexities introduced by the incentive and/or truth-telling constraints[13]. In contrast, the linear programming approach is extremely general and can be applied for *any* possible preference and technology specifications since, by construction, it convexifies the problem by allowing all possible lotteries over allocations. The only potential downside is that the LP method may suffer from the 'curse of dimensionality'. However, as shown above, by judicious formulation of the linear programs, this deficiency is minimized. The main reason for using discrete grids for all variables is not the dynamic programming part, which can be also solved without discretization (e.g., using splines), but our linear programming approach to the MH and UI regimes (necessitated by non-convexities) and our empirical application using the likelihood of the discretized joint distribution of the data.

To speed-up computation, we solve the dynamic programs for each regime using 'policy function iteration' (e.g., see Judd, 1998). That is, we start with an initial guess for the value function, obtain the optimal policy for that value function and iterate until convergence on the Bellman operator in policy space. At each iteration step we solve a linear program[14] in the policy variables π for each possible value of the state variables. In the unobserved investment (UI) regime the promised utilities set, $\mathbf{W}$ is endogenously determined and is solved for together with V. Using the incentive compatibility constraints, we restrict attention to non-decreasing promise vectors $\mathbf{w}(k)$. Specifically, we 'discretize' the set $\mathbf{W}$ by starting with a large set $\mathbf{W}_0$ consisting of linear functions $\mathbf{w}(k)$ with intercepts that take values from the grid $W = \{w_{\min}, w_2, ...w_{\max}\}$ defined in (11), and a discrete set of non-negative slopes. We initially iterate on the UI dynamic program using value function iteration, that is, we iterate over the promise set $\mathbf{W}$ together with the value function V, dropping all infeasible vectors $\mathbf{w}$ at each iteration and 'shrinking' $\mathbf{W}$ as a result (Abreu, Pierce and Stacchetti, 1990). Once we have successively eliminated all vectors in $\mathbf{W}$ for which the respective linear programs have no feasible solution, that is, once we have converged to the self-generating feasible promise set $\mathbf{W}^*$, we switch to (the much faster)

[11] Given our primarily empirical objectives, we chose general functional forms that preclude analytical tractability. We verify robustness by using different parameterizations and model specifications.

[12] The first order approach requires imposing (strong) monotonicity and/or convexity assumptions on the technology (Rogerson 1985; Jewitt, 1988) or, alternatively, as in Abraham and Pavoni (2008), employing a numerical verification procedure to check its validity for the particular problem at hand.

[13] We do find such non-convexities in our solutions for the MH and UI regimes and hence we cannot use the first order approach as it is not always valid in our set up.

[14] The coefficient matrices of the objective function and the constraints are created in Matlab while all linear programs are solved using the comercial LP solver CPLEX version 8.1. The computations were performed on a dual-core, 2.2 Ghz, 2GB RAM machine.

policy function iteration and continue iterating on the Bellman equation until convergence[15]. The same approach is used for the set $\mathbf{W}_m$.

3.2 Computing Joint Distributions of Model Variables

Our approach allows us to characterize and take to data the different model regimes using the optimal policy functions, $\pi^*(.)$ which solve the dynamic programs in section 2. We first construct the (Markov) state transition matrix for each regime. Formally, denote by $s \in S$ the current state – k in autarky, (b, k) in S/B, or (w, k) in the MH/FI/UI regimes. The transition probability of going from any current state s to any next-period state s' is computed using the optimal policy $\pi^*()$, integrating out all control variables. For example, for the MH regime we have:

$$\text{Prob}(w', k'|w, k) = \sum_{T \times Q \times Z} \pi^*(\tau, q, z, k', w'|w, k)$$

Putting these transition probabilities together for all state pairs, we form the state transition matrix $\mathbf{M}$ of dimension $\#S \times \#S$, (for example, for MH $\#S = \#K \times \#W$), with elements m_{ij}, $i, j = 1, ..., \#S$ corresponding to the transition probabilities of going from state i to state j in S.

The matrix $\mathbf{M}$ completely characterizes the dynamics of the model. Specifically, we can use $\mathbf{M}$ to compute the cross-sectional distribution over states at any time t, $\mathbf{D}_t(s) \equiv (d_t^1, .., d_t^{\#S})$, starting from an arbitrary given initial state distribution, $\mathbf{D}_0(s)$:

$$\mathbf{D}_t(s) = (\mathbf{M}')^t \mathbf{D}_0(s) \tag{20}$$

Setting $t = \infty$ gives the stationary state distribution if one exists. One can think of $D_0(s)$ as the population probability distribution (or, in the sample, frequency histogram) over the states s. In practice, in our empirical applications, some elements of the state s are unobservable to the researcher, for example, the state variable w in the MH and FI regimes. Thus, to initialize the model we use the empirical initial distribution of the observable state (k) and assume that the unobserved state is drawn from some known distribution, the parameters of which we estimate.

We further use the state probability distribution (20) in conjunction with the policy functions $\pi^*(.)$ to compute cross-sectional probability distributions $\mathbf{H}_t(x)$ for any model variable x (which could be k, k', z, τ, q, c, etc.), or any combination of these variables, at any time period. For example, in the MH regime, the time t joint cross-sectional distribution of next-period assets k' over the grid K with elements k'_i, $i = 1, .., \#K$ and current output q over the grid Q with elements q_h, $h = 1, .., \#Q$ is:

$$H_t(k'_i, q_h) \equiv \text{Prob}_t(k' = k'_i,\ q = q_h|\mathbf{D}_0) = \sum_{j=1..\#S} d_t^j \sum_{T \times Z \times W'} \pi_t^*(\tau, q = q_h, z, k' = k'_i, w'|s^j)$$

We also use the time-t distribution over states $\mathbf{D}_t(s)$ and the Markov matrix $\mathbf{M}$ to compute transition probabilities, $\mathbf{P}_t(x, x')$ for any model variable x, at any time period, t. The transition and the cross-sectional probabilities are easily combined to construct joint probability distributions encompassing several periods at a time as in a panel.

3.3 Functional Forms, Grids, and Baseline Parameters

Below are the functional forms we adopt for the empirical analysis. They are chosen and demonstrated below to be flexible enough to generate significant and statistically distinguishable differences across

[15]We also verified the results against proceeding with value function iteration all the way.

the financial regimes. Nevertheless, as argued earlier, our methods allow for any alternative or more general specifications of preferences and technology.

Agent preferences are of the CES form:[16]

$$U(c,z) = \frac{c^{1-\sigma}}{1-\sigma} - z^{\theta}$$

The production function, $P(q|z,k)$ representing the probability of obtaining output level, $q \in Q \equiv \{q_1, q_2, ..q_{\#Q}\}$, from effort $z \in Z$ and capital $k \in K$ is:

$$P(q = q_1|z,k) = 1 - \left(\frac{k^{\rho} + z^{\rho}}{2}\right)^{1/\rho}$$
$$P(q = q_i|z,k) = \frac{1}{\#Q - 1}\left(\frac{k^{\rho} + z^{\rho}}{2}\right)^{1/\rho} \quad \text{for } i = 2, .., \#Q$$

where q_1 is the lowest output. The probability of obtaining each output level is bounded away from zero[17]. The functional form for P encompasses a wide range of production technologies. With $\rho = 1$ we have a perfect substitutes technology; with $\rho \to 0$ the technology is Cobb-Douglas; and with $\rho \to -\infty$, the technology is Leontief[18].

To get an idea of the computational complexity of the dynamic contracting problems we solve, Table 1 shows the number of variables, constraints, and linear programs that need to be solved at each iteration for each regime for the grids we use in the empirical implementation. The number of linear programs is closely related to the grid size of the state variables while the total number of variables and constraints depends on the product of all grid dimensions. The biggest computational difficulties arise from increasing $\#K$ or $\#Z$ as this causes an exponential increase in the number of variables and/or constraints. This is why we keep these dimensions relatively low, whereas large $\#T$ is relatively 'cheap' computationally. In the UI regime the biggest computational difficulties arise from the huge number of linear programs to be computed over the two sub-stages.

The grids that we use in the estimation runs are defined in Table 2 and reflect the relative magnitudes and ranges of the variables in the Thai data. We can use (and did robustness runs with) much finer grids but the associated computational time cost is extremely high at the estimation stage because of the need to compute the linear programs and iterate at each parameter vector during the estimation. This is why we keep dimensions relatively low at present. Unfortunately, because of the extreme dimensionality and computational time requirements of the UI regime (see Table 1), we are currently unable to estimate it, with the exception of two runs with coarse grids[19].

Table 3 displays the baseline parameters used in the estimation runs with model-generated data that follow. These parameters are representative, from a set of many runs we did, and chosen to generate well-behaved interior solutions relative to the grids. In addition, we have done various robustness checks for other parameter values.

[16]Our LP numerical methodology does not require separable preferences. However, assuming separability is a common specification in the dynamic contracts literature and it also significantly simplifies the analysis of the UI regime.

[17]To have well-defined likelihood ratios and satisfy the full support condition, the probabilities we use in the computations are constrained between .01 and .99. That is, in computations, we actually use $\tilde{P} = \min\{.99, \max\{.01, P\}\}$.

[18]Note, that the functional form of P determines expected, and not actual output, thus our CES parameter ρ is not comparable to specifications and values from the macro literature.

[19]Currently, a single functional evaluation of the UI regime likelihood for our baseline grids takes about 45 minutes (as opposed to 7-9 sec in the MH regime) and over 1,500 such evaluations (47 days) are typically required to find the MLE parameters for a single estimation run. We are working on a parallel computing version of our estimation algorithm as well as optimizations based on the NPL approach (Aguirregabiria and Mira, 2002; Kasahara and Shimotsu, 2009).

4 Empirical Implementation

4.1 Maximum Likelihood Estimation

In PTK (2006) we estimated via maximum likelihood a one-period model of occupational choice with financial constraints. We used binary data on occupational choice and data on ex-ante wealth. Using the full-blown dynamic mechanism design models presented here, we significantly expand our previous approach and our goal is to fit cross-sections or panels of consumption, investment, productive assets and income[20]. We use both model-simulated and actual data in the estimation exercises that follow.

Specifically, let the data we have be $\{\hat{y}_{i,t}\}$ where $i = 1, ...n$ denotes sample units and $t = 0, ...T$ denotes time. For example, in this paper we use data from Thai households on assets, consumption and income, $\{\hat{k}_{i,t}, \hat{c}_{i,t}, \hat{q}_{i,t}\}$. We first use the initial period assets data $\{\hat{k}_{i,0}\}_{i=1}^n$ to form the initial asset distribution over the grid K (the observable state variable), $H_0(\hat{k}) \equiv \{h_0(k_1), ..., h_0(k_{\#K})\}$ which is an input needed to initialize the model. For the unobservable state variables, b and w we assume that they are drawn from a parametric initial distribution $\Phi(.)$ (e.g. normal, mixture of normals) the parameters of which will be estimated.

What we do next is take (a subset of) the data, for example the cross-sectional distribution $\{c_{i,t}, q_{i,t}\}_{i=1}^n$, and form the likelihood function for each model regime. To do so, we discretize the data in the form of frequencies over mutually exclusive 'cells' which are defined by the grid $C \times Q$. In general, suppose there are J such mutually exclusive grid cells over which we approximate / discretize the joint distribution of $\hat{y}_i$ in the data. Call these cells Y_j for $j = 1, ..., J$. For example, using grids of five points for c and five points for q we have $J = 25$ such mutually exclusive cells (a 5-by-5 cell matrix), the frequencies over which in the data approximate the joint bivariate distribution of $c_{i,t}$ and $q_{i,t}$.

Let the model we want to estimate generate the probability distribution function $f(y|\phi, H_0(\hat{k}))$ over the variables y (e.g., c and q) we are fitting, given the vector ϕ of the model structural parameters and distributional parameters for the unobserved states and given the initial distribution over the observable state, $H_0(\hat{k})$. We also allow for measurement error in the model variables with mean zero and standard deviation $\tilde{\gamma}_{me}$ – also a parameter included in ϕ (see below for more details). The conditional on $H_0(\hat{k})$ likelihood of the observed data, $\hat{y}$ (assuming they are i.i.d.) is:

$$L(\phi) = \prod_{i=1}^{n} f(\hat{y}_i|\phi, H_0(\hat{k}))$$

or, taking logs, we can re-write the above as:

$$\sum_{i=1}^{n} \ln Prob(y \in \hat{Y}_i|\phi, H_0(\hat{k}))$$

that is, log of the (joint) probability for the variables y (e.g., consumption and income) from the model to belong to cell $\hat{Y}_i$ (one of Y_1 through Y_J) to which data point $\hat{y}_i$ belongs, given the initial observable state distribution $H_0(\hat{k})$ from the data and given the model parameters, ϕ. The likelihood integrates over the unobserved state variable distribution and the measurement error distribution. This is done when $f(y|\phi, H_0(\hat{k}))$ is computed (see more on this below).

[20]In general, the MLE approach proposed here can be used on any data in the form of discretized (joint) distributions – that is any cross-sectional distibutions of model variables, panels, or transition probabilities. Almost any dataset can be put in this form by sorting the observations into appropriately chosen 'bins' and then using the bin frequencies in the estimation (see also Jappelli and Pistaferri, 2006 who perform a similar exercise for transition probabilities but using a minimum-χ^2 estimator).

Separating the above sum by cells and dividing by the sample size n we obtain the (normalized) log-likelihood criterion that we use to estimate each model:

$$\Lambda(\phi) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{J}\mathbf{1}(\hat{y}_i \in Y_j)\ln Prob(y \in Y_j|\phi, H_0(\hat{k})) \tag{21}$$

where $\mathbf{1}(\hat{y}_i \in Y_j)$ is the indicator function which equals one is the observation $\hat{y}_i$ belongs to cell Y_j and zero otherwise. Note that $\hat{y}_i$ could also have a temporal component too, for example in some runs below we use $\hat{y}_i = \{\hat{c}_{i,t}, \hat{q}_{i,t}, \hat{c}_{i,t+1}, \hat{q}_{i,t+1}\}$. The maximization in (21) is performed by an optimization algorithm robust to local maxima (we use pattern search and polytope). Standard errors are computed via bootstrapping, repeatedly drawing with replacement from the data.

Discretizing the joint variable distributions in the data into frequencies over mutually exclusive cells offers two important advantages. First, it allows us to write the likelihood function explicitly, fully employing the discrete nature of the LP approach. Second, being able to use maximum likelihood allows us to use a formal statistical test (Vuong, 1989), to compare across the competing model regimes. In principle, one could employ GMM or minimum-χ^2 estimation methods using arbitrary moments instead of our discretized joint distributions. Unfortunately, to our knowledge, no tractable and computationally feasible way of implementing a statistical test to compare across dynamic structural models like ours exists in this case[21].

4.2 Testing and Model Selection

We follow Vuong (1989) to construct and compute an asymptotic test statistic that we use to formally distinguish across the models of financial regimes using simulated or actual data. Vuong's test is based on the maximum likelihood method. Its most attractive feature is that it does not require that either of the compared models be correctly specified. Specifically, suppose the values of the estimation criterion function being minimized (i.e., minus the log-likelihood) for two non-nested[22] competing models are given by $L_n^1(\hat{\phi}^1)$ and $L_n^2(\hat{\phi}^2)$ where n is the common sample size and $\hat{\phi}^1$ and $\hat{\phi}^2$ are the parameter estimates for the two models. The pairwise nature of the test conveniently allows us to obtain a complete ranking by likelihood of all models we study. The null hypothesis, H_0 of the Vuong test, is that the two models are "asymptotically equivalent" relative to the true data generating process – that is, cannot be statistically distinguished from each other based on their 'distance' from the data.

Define the "difference in lack-of fit" statistic:

$$T_n = n^{-1/2}\frac{L_n^1(\hat{\phi}^1) - L_n^2(\hat{\phi}^2)}{\hat{\sigma}_n}$$

where $\hat{\sigma}_n$ is a consistent estimate of the asymptotic variance[23], σ_n of $L_n^1(\hat{\phi}^1) - L_n^2(\hat{\phi}^2)$ (the likelihood ratio). The main result is:

[21] Rivers and Vuong (2002) propose a general test that could be used in theory, however, it is computationally infeasible in our framework.

[22] For the functional forms and parameter space we consider and use in the estimations, the regimes we study are statistically non-nested. Formally, following Vuong (1989), we say that model A nests model B, if, for any possible allocation that can arise in model B, there exist parameter values such that this is an allocation in model A. The Vuong model comparison test can be also used for "overlapping" models, i.e. neither strictly nested nor non-nested, in which case the test statistic has a weighted sum of chi-squares distribution (see Vuong, 1989, p. 322).

[23] Such a consistent estimate is the sample analogue of the variance of the LR statistic (see Vuong, 1989, p. 314).

Proposition (Vuong, 1989):

Under certain regularity conditions (see Vuong, 1989, pp. 309-13), if the compared models are strictly non-nested, the test-statistic T_n is distributed $N(0, 1)$ under the null hypothesis.

4.3 Methodological Discussion

Initial conditions and transitions vs. steady states

The financial regimes we study naturally have implications for both the transitional dynamics and long-run distributions of the model variables such as consumption, assets, investment, etc. If one has reasons to assume that the actual data is from a steady state (stationary distribution), then they may not depend on the initial conditions or the transition path followed and so the regimes can be estimated by simply matching the simulated with the empirical cross-sectional stationary distributions. However, given our application to Thailand, an emerging, developing economy, we take the view that the actual data is more likely to correspond to a transition than to a steady state. Thus, estimating initial conditions is very important for us, as well as fitting the subsequent transition trajectory, which requires using intertemporal data. In addition, our computations show very slow dynamics for the model variables in the mechanism design regimes and also, theoretically, these regimes can have degenerate steady states (e.g. immiserization in MH) which are additional reasons to focus on transitions instead of steady states.

Specifically, in our models the initial conditions are the $t = 0$ values of the state variables (k for A, (k, b) for S/B, and (k, w) for the mechanism design regimes). As pointed out above, some of these states are unobserved by the econometrician – for example, the initial promise w. Thus, as we are interested in transitional dynamics, the initial distribution of w capturing unobserved heterogeneity across agents is imputed (and its parameters estimated) to initialize the MH/FI regimes. In contrast, we use the initial distribution of the observable state, k as an input to each model. More details follow in the estimation sections.

Identification

Because of the analytical complexity of our dynamic incomplete markets model, it is not possible to provide theoretical identification proofs while keeping the setting sufficiently general. In fact, we are aware (e.g., Honore and Tamer, 2006) that point identification may sometimes fail in complex structural models like ours. To address this issue we use the following algorithm – a form of 'numerical identification': Step 1 – take a baseline model regime parametrized by a vector of parameters, ϕ^{base}; Step 2 – generate simulated data from the baseline regime; Step 3 – estimate the baseline model using the data in Step 2 using maximum likelihood and obtain estimates, $\hat{\phi}^{base}$; and Step 4 – if the estimates from Step 3 are numerically close to the baseline ϕ^{base}, report success, otherwise report failure. In other words, before going to the actual data, we use data simulated from the model itself to verify that our estimation methodology performs as it should. We use this approach in all estimation runs in section 4 for various model specifications and data.

5 Estimation and Model Selection Using Simulated Data

To assess the performance of our empirical methods while keeping the environment under complete control we first estimate and test all regimes using simulated data from one of them, before moving on to actual data in the following section. Specifically, we adopt as a baseline the moral hazard regime with observed investment[24] (MH) and simulate data from it, which we then use to estimate and test across

[24]In the UI regime k and k' are unobserved to third parties in the theoretical model but are (ex-post) observed by the econometrician. The survey enumerators take the time and trouble to ask about these as best as possible. Of course,

all regimes, including MH to verify if we 'recover' the true regime and the data-generating parameters.

5.1 Generating Data From the Model

We use the grids from Table 2 and the baseline parameters from Table 3 to simulate data from the MH model regime[25]. To initialize the MH regime, we take an initial distribution over the states (k, w) which has an equal number of data points for each grid point in the capital grid K and is normally distributed in w, i.e., $w\tilde{\ }N(\mu_w, \gamma_w^2)$ for each $k \in K$. We set the mean μ_w to be equal to the average value in the promise grid, $\frac{w_{\max}+w_{\min}}{2}$, at the baseline parameters. We then draw n (the sample size) random numbers from $N(\mu_w, \gamma_w^2)$ (that is, we assign $n/\#K$ of these draws to each $k \in K$) and initialize the distribution over the state space[26]. Next, we compute the data-generating regime (MH) at the baseline parameters, ϕ^{base} (see Table 3) given the drawn initial distribution over states (k, w) as described above and use the LP solution π^* to generate the theoretical distributions, $f(.|\phi^{base}, H(k_0))$ of[27] $c, q,$and k' (including jointly), conditional on the observable state distribution $H(k_0)$. We use these conditional distributions from the baseline MH regime (already discretized given our LP method) as simulated 'data' to use in the estimation.

In addition, we allow for additive measurement error in consumption, c, assets, k and income, q. The measurement error, ε added to each variable belongs to the Normal distribution $N(0, \tilde{\gamma}_{me}^2)$. We use the Normal cdf to compute the joint distribution, $f(y|\phi^{base}, H_0(\hat{k}))$ (see (21)) with measurement error added, for example, the (c, q) joint distribution. The way we do it, basically, is compute analytically how much of the probability mass at each grid cell 'spills over' to nearby grid cells due to the measurement error. We perform all estimation and testing exercises in this section for two measurement error specifications: 'low measurement error', where we set $\tilde{\gamma}_{me}$ equal to 10% of the grid span of the respective variable, and 'high measurement error', with standard deviation $\tilde{\gamma}_{me} = 50\%$ of the grid span. That is, we set the standard deviation of the measurement error proportional to each variable's grid span: $\tilde{\gamma}_{me} = \gamma_{me}(grid\ span)$ where $\gamma_{me} = .1$ or $.5$ is the proportionality parameter[28].

The results displayed in this section are representative for many more runs we did, with various other parametrizations. We discuss some of these in the robustness section. In addition, we also re-do the estimation and model selection runs with simulated data reported in this section for the parameter estimates from the Thai data. While we feel we have done our best to verify the robustness of our findings here, this section is primarily intended as a 'proof-of-concept' run for our methods before applying them to the actual Thai data.

5.2 Baseline Results Using Simulated Data

The parameters, ϕ that we estimate are the distributional parameters for promises and measurement error (μ_w, γ_w and γ_{me}) and the three structural parameters of the model – the preference parameters, σ and θ and the technology parameter, ρ. The discount factor β, the outside cost of capital, R and the depreciation rate, δ are calibrated to standard values (see Table 3). For the S and B regimes, instead of the parameters μ_w and γ_w we estimate the mean, μ_b and standard deviation, γ_b of the distribution

---

incentives to hide or mis-report assets could still be present.

[25] In the computer code we re-write the MH and FI regimes' dynamic programs from section 2 in terms of consumption instead of transfers. The two formulations are equivalent.

[26] Our methods allow any other possible initializations (mixtures of normals or bivariate distributions), at the cost of additional parameters to be estimated and slower computation. In the Thai data application in the next section we use the actual initial discretized distribution of assets in the data.

[27] We also construct investment, i as $i \equiv k' - (1 - \delta)k$.

[28] By using relative rather than absolute level of measurement error, we keep its standard deviation commensurate across model variables with different grid ranges.

of b (assumed unobserved and drawn from a normal distribution, consistent with our treatment of w in the mechanism design regimes).

For each regime we follow the procedure described in section 4.1 – we first generate the initial state distribution, then generate simulated data, apply measurement error[29], and compute the discretized joint probability distribution of the variables of interest. We then form the likelihood function (21) and use a search-based global optimization routine[30] to solve for the estimates $\hat{\phi}^{MLE}$ maximizing the likelihood between the baseline data and the estimated regime. As mentioned above, one of the regimes we estimate is the data-generating regime (MH) itself, in order to verify whether we successfully recover the data-generating parameters (see the numerical identification discussion in section 4.3). Finally, we perform Vuong tests to establish whether we can distinguish statistically between the data-generating and the rest of the regimes, as well as between any counterfactual regime pairs (e.g., B and S).

5.2.1 Investment and Income

We first estimate and test the financial regimes based on their implications about assets, investment and income (cash flow). To that purpose we simulate a data sample of size $n = 1,000$ from the joint distribution of (k, i, q) in the baseline (MH) regime. Tables 4 and 5(A), respectively, display the parameter estimates and Vuong test results with these data. Table 4, using the low measurement error specification, shows that, when estimating the data-generating MH regime the baseline parameter values (last row in each section, in italics) used to generate the data for γ_{me}, θ and ρ are recovered relatively well but σ, μ_w and γ_w are a bit off. In terms of likelihoods, the MH regime naturally obtains the highest likelihood (as data-generating), followed very closely by the other mechanism design regime (FI), the B and S regimes, and finally autarky. With high measurement error (not reported in table 4 to save space) the likelihoods for all regimes are lower and some are very close so that several regime pairs are tied (table 5).

The parameter estimates in Table 4 differ across the estimated regimes, as the MLE procedure is trying to fit the common data, but the estimates are generally quite similar between the FI and MH regimes (apart from μ_w). This is not the case for the exogenously incomplete markets regimes (B, S, A) where to fit the data some of the parameters (e.g. ρ) take values far from the data-generating ones. The B, S and A regimes also seem to require higher measurement error to fit the data (especially the A regime), compared to the baseline value for γ_{me} (0.1). The bootstrap standard errors of the parameter estimates are in general relatively small.

Turning to the results of the pairwise Vuong tests (section 1 of Table 5), we find that in the low measurement error specification we are able to recover the data-generating regime (MH) as best-fitting and to distinguish between it and the B, S and A regimes almost perfectly (at the 1% significance level) but, with the (k, i, q) data we cannot distinguish between the MH and FI regimes statistically. We also distinguish across the regimes in all pairwise comparisons between counterfactual (non-MH) regimes with only one exception for B vs. S). That is, even if the researcher (incorrectly) believes that the data were, for example, generated from the FI regime, he/she can still distinguish it from the B, S and A regimes. In contrast, with high measurement error in the baseline simulated data, the distinction between the regimes is more blurred and, based on the Vuong test, we cannot discern statistically between the MH/FI, MH/B and FI/B regime pairs. This suggests that additional data is needed to distinguish these regimes from each other. In all cases, including high measurement error, all non-autarky regimes are statistically distinguishable at the 1% level from the autarky regime.

[29] Allowing for measurement error rules out zero probability events that would make the likelihood of some models infinite.

[30] We first perform an extensive grid search over the parameter space to rule out local extrema and then use the Matlab global optimization routines `patternsearch` and `fminsearch` to maximize the likelihood.

5.2.2 Consumption and Income

We next estimate and test whether we can distinguish between the regimes based on the degree of consumption smoothing they imply, as embedded in the consumption-income (c, q) joint distribution. The results are shown in Table 4 (second section) and Table 5, section B. As with the (k, i, q) data, the likelihood values are ordered MH, FI, B, S, and A from highest to lowest. Thus, the regime likelihood order, relative to the data-generating MH model remains robust and not affected by the type of data used. Using (c, q) data we recover the data-generating parameters better than with (k, i, q) data with the exception of ρ (compare the row for the MH regime estimates with the row for the baseline in Table 4). Again, the non data generating regimes seem to require more measurement error (especially A) to fit the data and the standard errors (with the exception of that for ρ in the FI regime) are low. The parameter estimates for the exogenously incomplete regimes in many instances differ significantly from the baseline ones as the MLE is adjusting them to fit the data best.

Regarding our ability to distinguish the regimes using (c, q) data, with low measurement error, the baseline (MH) regime is distinguished at the 1% significance level from all alternatives. This is not the case, however, for the counterfactual regimes where FI is tied with B and S. With high measurement error we are able to distinguish between the MH baseline and some alternative regimes (B, S) only at a lower confidence level, including a statistical tie with the FI regime. As in the results with (k, i, q) data, a number of other regime pairs are also tied at the 10% level with high measurement error. The autarky regime is again statistically distinguished from the others, including in the high measurement error specification.

5.2.3 Combined Data on Consumption, Investment, and Income

Theoretically it is known that with incomplete markets (which all our regimes except FI assume), the classical separation between consumption and production/investment decisions fails. A natural question is then whether using joint data on consumption and investment would enable us to distinguish the regimes even with substantial measurement error. More generally, more information should be present in the joint data on c, k, i and q than in consumption-income and investment-income data separately.

The parameter estimates with (c, q, i, k) data are reported in the third section of Table 4. The results are very similar to those with (c, q) and (k, i, q) data described above. The data-generating regime is 'recovered' as best-fitting and also the parameter estimates are very close to the baseline parameters (compare with the (k, i, q) case in particular). The regime order in likelihood remains MH, FI, B, S, A; the parameters are precisely estimated and the exogenously incomplete regimes require higher measurement error to fit the data (in the range .22-.72 compared to the .1 baseline).

Section C of Table 5 reports the estimated likelihoods and Vuong test results using data on the joint distribution of c, q, k and i. The ability to distinguish the data generating regime (MH) from all alternatives is nearly perfect (at the 1% level) with both low and high measurement error. That is, using the joint data causes a significant improvement (especially in the high measurement error case) in the ability to distinguish between the data-generating (MH) regime against each alternative, compared to when using consumption or investment data separately – compare the number of ties in sections A and B to that in section C of Table 5. The ability to distinguish between counterfactual regimes (i.e., those different from MH) also improves significantly especially relative to when using (c, q) data (the number of ties falls from five to three overall). Only two cases remain (FI/B and FI/S with high measurement error) in which if the researcher guesses the data-generating regime incorrectly he/she would be unable to distinguish it from some other regime. Overall, even substantial measurement error (50% of the entire range of values that c, k and k' may take) does not impede our ability to distinguish the true regime once joint consumption, investment and cash flow data are used.

5.2.4 Intertemporal Data – Panel

We also estimate and test the financial regimes using their implications about the dynamics of the consumption and income joint distribution. Specifically, we use simulated data on the joint distribution of consumption and income, (c, q) in two different periods: $t = 0$ and 1, or $t = 0$ and 50, as in a panel dataset. Section D of Table 5 reports the Vuong test statistics. We use the high measurement error specification only and investigate whether using intertemporal data improves our ability to distinguish across our dynamic models of financial markets compared to when using a single cross-section (at $t = 1$) as in Table 5(B).

Compared to part B, part D of Table 5 demonstrates that using intertemporal data significantly improves our ability to distinguish the regimes – the number of ties diminishes from four (plus two marginal comparisons) – see row 2 in part B – to zero or one ties, depending on the panel time span (rows 1 and 2 in part D). Also, the Vuong test statistics are larger in most cases showing the regimes are distinguished better. The improvement in ability to discern the regimes using intertemporal data is comparable and even slightly better compared to when joint data on consumption, investment and cash flow was used (compare sections C and D). The time period between the panel periods has negligible effect on the results.

5.3 Robustness Runs

We report a number of additional estimation runs to study the robustness of our results.

5.3.1 Using Simulated Data from the Borrowing Regime

First, instead of generating the baseline data from the MH regime, we generated it from the borrowing regime (B). Table 6(A) presents the Vuong test statistics using simulated data on consumption, assets, investment, and income (c, k, i, q) generated from the B regime at the baseline parameters from Table 3. Comparing the results to those in Table 5C, we see from the test-statistics that, with data generated by the B regime, the likelihood order naturally changes, with B now producing the highest likelihood, followed by MH, B, FI and A. Once again, the autarky regime is furthest away from the data-generating regime and it is distinguished from all alternatives at the 1% level. With low measurement error we can distinguish across all regime pairs at the 5% significance level. As before, larger measurement error reduces our ability to distinguish some of the regimes – the MH/B and MH/S regime comparisons produce ties.

5.3.2 Additional Robustness Runs

Table 6 part B contains the results from five additional robustness runs using (c, q, i, k) data generated from the MH regime. Row 1 analyzes the effect of generating the data without measurement error[31] (we set $\gamma_{me} = 0$). The Vuong test statistics for the comparisons between the MH and the rest of the regimes go up relative to Table 5C but some counterfactual regime pairs are indistinguishable. The autarky regime remains statistically distinguishable from all others.

Rows 2 and 3 of Table 6B study the effect of varying the simulated sample size using the high measurement error specification. We find that reducing the sample size, n from 1,000 to 200 significantly reduces the power of the Vuong test and, as a result, we cannot distinguish between any of the MH, FI, B and S regimes, only autarky stands out. In contrast, increasing n to 5,000 achieves very similar results to our $n = 1,000$ baseline run with (c, q, i, k) data.

[31] If a regime implies zero probability in some joint distribution frequency cell, we assign a large negative number (instead of minus infinity) for $\ln m(\phi)$ in (21) for that cell.

The fourth row of Table 6 part B checks the sensitivity of our results to grid dimensionality. Reducing the size of all grid to three points (from five) does not affect the likelihood values or the Vuong statistics significantly relative to those in Table 5(B) which is reassuring for the robustness of our findings. The last row in Table 6B allows for additional heterogeneity in the model, through allowing for 'productivity' differences across agents. Specifically, we draw ten productivity values from a uniform distribution on [0.75,1.25] and compute the MH regime multiplying the grid Q by each productivity factor, to capture 'skill' heterogeneity. We draw the simulated data from these heterogeneous joint distributions, ending up with a joint (c, q, i, k) distribution that corresponds to that of a mixture of households with different productivities. We then estimate all regimes as if those differences do not exist (i.e., as if we mistakenly treat the data as generated without such differences). Line 5 in Table 6B shows that allowing for this additional source of unobserved heterogeneity (and mis-specification) in the model does not affect the robustness of our results. We still recover MH as the best-fitting regime, distinguished at the 1% level from all others. The counterfactual regime likelihood order is also preserved relative to the baseline runs.

We also report the Vuong test statistics when we generate data from the MH regime at the estimated parameters from the 1999 Thai data (see next section). Table 6C shows that the best fitting regimes are exactly the same as in our low measurement error specification – the MH regime is recovered as best-fitting (tied with FI in the (k, i, q) case). There are some differences in terms of ties among the counterfactual regimes. Note that here we cannot directly compare the number of ties across the runs with different data as they are computed at different parameter vectors (see table 8 below).

Finally, in Table 6 part D we perform a run with the moral hazard with unobserved investment (UI) regime for our coarse, three-point grid specification with the (c, q, i, k) data simulated from the MH regime (read together with line 4 in part B). As explained above (see footnote 19) the UI regime is extremely computationally heavy to estimate so we only compute this single run as proof of concept. The results in the table show that, as expected, the data-generating MH regime achieves highest likelihood followed by FI, UI, B, S and A.

6 Application to Thai Data

6.1 Data

In this section we apply our structural estimation methodology to actual data from a developing country. The data we use come from the Townsend Thai Monthly Survey (Townsend, Paulson and Lee, 1997). The survey began in August 1998 with a comprehensive baseline questionnaire on an extensive set of topics, followed by interviews roughly every month. Initially consumption data were gathered weekly, then bi-weekly. The data we use here begins in January 1999 so that technique and questionnaire adjustments were essentially done. We use a panel of 531 households observed for seven consecutive years, 1999 to 2005. The data were gathered from 16 villages in four provinces, two in the relatively wealthy and industrializing Central region near Bangkok, and two in the relatively poor, semi-arid Northeast. All variables were added up to produce annual numbers.

Consumption expenditures, c include owner-produced consumption (rice, fish, etc.). Income, q is measured on an accrual basis (see Samphantharak and Townsend, 2009) though, at an annual frequency, this is close to cash flow from operations. Business assets, k include business and farm equipment, but exclude livestock, and household assets such as durable goods (we do not attempt to distinguish farm from nonfarm enterprise, though the distinction between assets and durable goods is sometimes not so obvious as well). We perform a robustness check with respect to the asset definition – see section 6.3. Assets other than land are depreciated. All data are in nominal terms but inflation was low over this period. The variables are not converted to per-capita terms, i.e., household size is not brought into consideration. Brief summary statistics for the Thai data are displayed in Table 7. We construct a measure of investment using the assets in two consecutive years as: $i_t \equiv k_{t+1} - (1-\delta)k_t$ for each household.

We convert the data into model units rather than Thai currency by dividing all currency values by the 90-th percentile of the assets distribution (1,742,557 Thai baht).The first step in our empirical strategy is to put the data in the frequency form we need to perform MLE estimation. The normalized asset values are placed on a five-point grid[32], $K \equiv [0, .02, .1, .4, 1]$. The unequal spacing of the grid reflects the skewness of the asset distribution in the data with numerous small and few large values. Normalized consumption, c is placed on an evenly-spaced five-point grid on the interval $[.1, .9]$, while normalized income, q is fitted on a five-point uniformly spaced grid on the interval $[.05, 2]$. These grids imply an upper bound of .9 for the B grid in the B regime to ensure no default. The chosen grid bounds reflect the actual magnitudes of c and q in the data relative to those of assets k.

Having discretized the Thai data over these grids, we construct cross-sectional joint distributions of model variables, e.g., (c, q), (k, i, q), (k, i, q, c) for any data year. We then use the procedures described in section 4.1 to generate the initial state distribution[33] assuming normal distributions for the unobservable states w and b with means and variances to be estimated. As in the runs with simulated data, measurement error with standard deviation proportional to grid span is allowed for all data variables.

6.2 Estimation Results

We compute each regime and, using the (normalized and discretized) Thai data, maximize the likelihood function between the model-generated joint distributions and their counterparts in the data, as in (21). As before, we estimate the structural parameters σ (risk aversion), θ (effort curvature), and ρ (technology curvature), together with the distributional parameters, μ_w, γ_w (μ_b, γ_b for B and S), and the standard error parameter, γ_{me}. For robustness purposes we estimate both using the first two ('99-00) or the final two ('04-05) years of the Thai data. Parameter estimates and bootstrap standard errors are in Table 8 (reported for '99-00 only) while the Vuong regime comparison test results are in Table 9 (for both '99-00 and '04-05).

As with the simulated data, the estimates differ across the regimes (see Table 8) as the MLE optimization adjusts the parameters for each model regime to attain best fit with the common data. In addition, holding the best-fitting regime (e.g., B) constant, if we compare across the sections in Table 8 using different data variables, we see that the parameter estimates inclusive of bootstrap standard errors are sensitive to the data used, something which does not happen with model-generated data.

6.2.1 Business Assets, Investment, and Income Data

We start by estimating and testing the implications of the different regimes about assets, investment and income using the joint cross-sectional distribution of (k, i, q) in the Thai data. When estimated from k, i, q 1999-00 cross-sectional data, the financial regimes rank in decreasing order of likelihood as: S, B, A, FI, MH (Table 8, last column), with the Vuong test unable to reject the hypothesis that the B and S regimes are equally close to the data (Table 9, section 1). With the '04-05 data the B, S and MH regimes are tied for best fit (row 1.2 in Table 9). As in the baseline runs with simulated data, the autarky regime is distinguished at the 1% significance level in all pairwise comparisons, but here

[32] We use a standard histogram function based on distance to the closest gridpoint (Matlab's command `hist`).

[33] We use the actual distribution of k in the first year of the data, 1999 to initialize each model. The frequencies at the five grid point on K in the 1999 Thai data are respectively 28.8%, 18.1%, 17.9%, 18.1% and 17.1%.

it does not come last in terms of fit. The MH, FI and A regimes which obtain the lowest likelihoods also feature the highest estimated level of measurement error (γ_{me}), again similar to our results with simulated data. The likely explanation is that, to compensate for the bad fit, the MLE procedure is raising the level of measurement error.

6.2.2 Business Assets, Investment, Income, and Consumption Data

Next, we evaluate the gains from using combined data on assets, income, and consumption as opposed to using (k, i, q) data only. The regimes' likelihood ranking (see the 'winning' regime in each pairwise comparison) remains the same with the B and S regimes coming on top, followed by A, MH and FI (section 2 of Table 9). Adding the consumption data to the set of variables, the joint distribution of which we estimate, we observe a significant improvement in our ability to distinguish the regimes relative to using (k, i, q) data only – the number of ties falls from four to two and there are no marginal significances. This confirms our previous results from the runs with simulated data.

6.2.3 Consumption and Income Data

We also test whether we can distinguish between the regimes based solely on the degree of consumption smoothing they imply relative to the Thai data using the (c, q) joint cross-sectional distribution (Table 9, section 3). We observe significantly more regime ties – seven ties and one marginally significant comparison over all years considered, compared to when using combined investment and consumption data (two ties). This once again confirms our finding that using combined investment and consumption data significantly improves our ability to distinguish the incomplete market regimes. Unlike the simulated data in section 4, the Thai consumption and income data alone seem to be unable to pin down precisely the best fitting regime.

The regime likelihoods using the joint (c, q) data, here in this sub-section, also rank differently compared to those using (k, i, q) or (c, k, i, q) data in the above sub-sections. The autarky regime has the worst likelihood and is always rejected against the others, but now the moral hazard regime achieves slightly higher likelihood than the exogenously incomplete (B and S) regimes with the 1999 data. However, the likelihood difference is not statistically significant – MH is tied with B in 1999 and with B and S in 2005. Thus, we cannot reject the hypothesis that the B (and S for 2005) regimes are as close to the data as the MH regime in these cases.

6.2.4 Model Dynamics: Panel Data

We also estimate and test across the alternative financial regimes using panel data, targeting differences in the models' dynamics. Specifically, we use data from the joint distribution of consumption and income in two different time periods as in a panel[34] (section 4 in Table 9). The regimes' likelihood ordering is consistent with that in the (c, q) single cross-section results (section 3 of Table 9), although the MH regime is now rejected, just like when we use the combined data. The ability to distinguish across the regimes with the 1999/00 (c, q) panel is much better than in the single 1999 (c, q) cross section (zero vs. three ties respectively). However, our ability to distinguish the regimes worsens with the gap between years included in the panel, with the autarky regime coming close and tied with B and S (this is the only instance in which this happens, among all baseline runs in Table 9) when the 1999/05 (c, q) panel is used.

[34]Unfortunately, using panel data of (k, i, q) or (c, q, i, k) is infeasible with our Thai data due to the high number of joint distribution cells required to be fitted (up to over 15,000) with only 531 observations available.

6.3 Robustness Runs

We also performed a battery of additional estimation runs (Table 10) to check the robustness of our results with the Thai data and shed more light onto the regime ordering patterns with consumption vs. investment data. Unless stated otherwise the table uses 1999-00 data.

A re-estimation imposing risk neutrality (that is, fixing $\sigma = 0$ instead of estimating σ) naturally hurts the fit of all regimes with the Thai data, but especially the MH and FI regimes. As a result, the B and S regimes fit best using the 1999 Thai (c, q) data (row 1.1 in Table 10). This shows that, as argued in the introduction, allowing for risk aversion explicitly can be important in identifying the underlying financial market imperfections in the data. Otherwise, imposing risk neutrality pins down the borrowing regime as the best fitting using the 1999 (k, i, q) and (c, q, i, k) cross-sections.

Estimating without allowing for measurement error (fixing its standard deviation, γ_{me} to zero) also reduces the regimes' fit (especially for autarky) and preserves the MH and B regimes' best fit with the Thai data on consumption and income while B emerges as the single best fitting regime with (k, i, q) data.

Another robustness check focuses on a sub-sample of households ($n = 391$) who are related by blood or marriage, as in a kinship network, using cross-sectional (c, q) and (c, k, i, q) data and (c, q) panel data (section 3 of Table 10). Compared to the whole sample results and likelihoods (Table 9), using this networked sub-sample data allows us to pin better the best fitting regime as moral hazard with the consumption and income cross-sectional data. This presents evidence that family networks help in consumption smoothing in the cross-section. This result is not very robust, however, as the 1999/00 (c, q) panel and the joint (c, q, i, k) data indicate that the B regime remains best-fitting (as in the whole sample) within the networked sub-sample when 'richer' data are used.

Section 4 of Table 10 re-estimates the regimes and compares across them when we allow for quadratic adjustment costs in investment. In particular, our full information regime with adjustment costs corresponds to the standard adjustment costs model in the literature (Bond and Meghir, 1994 among many others). In this specification, the (k, i, q) data alone is insufficient to discern across the B, S, MH and FI regimes (only autarky is rejected). That is, the B and S regimes are tied with the 'pure adjustment costs' regime from the literature (our FI) as well as with the moral hazard plus adjustment costs regime – presumably we need more data to distinguish the regimes, but our conclusions from before survive. Using combined consumption and investment data, however, recovers S and B with adjustment costs as the best-fitting regimes (although MH with adjustment costs is tied with the second-best regime B too).

Next, in section 5 of Table 10 we explore whether there are regional differences in the best fitting regime. In the Central (richer and fast industrializing) region the B regime (tied with S with (c, q) data alone) is revealed as the one characterizing best the joint data distribution. However, in the North-East (poorer and mostly agricultural) region, we cannot reject the moral hazard (MH) regime as best-fitting (it achieves the highest likelihood), including when using joint consumption, investment and income data.

In section 6 of Table 10 we perform several additional robustness runs. A re-estimation with another data sub-sample using individuals related via personal loans or gifts (Kinnan and Townsend, 2009) and (c, q) data puts the MH regime on top in terms of likelihood but statistically tied with the B and FI regimes. Again, some evidence of the role of such networks in better consumption smoothing is present. A run in which we used (c, q, i, k) data cleaned from household fixed effects imposing the average family structure and time and regional effects (as in Kaboski and Townsend, 2009) pins down the borrowing regime as best fitting, consistent with the baseline runs. A run with coarser grids (three points each, as in the simulated data section) produces the same results as the benchmark.

We also check robustness with respect to our definitions of assets and income in the Thai data.

We re-estimate using the (c, q, i, k) data including all household assets and livestock in the definition of k and exclude households who have only labor income. The sample size drops to 297 but our main findings about the best fitting regime from Table 9 (row 2.1) do not change, although there are now two regime ties because of the lower n. Finally, we also tried imposing a lower bound on promises in the MH and FI regimes equal to the agent's autarky value at $k = 0$, as in a limited commitment model (Ligon, Thomas and Worrall, 2002). This worsens the fit of the MH regime relative to the data – the ability to smooth consumption and investment is diminished with limited commitment. In other words, the degree of consumption smoothing in the Thai data is more consistent with our original moral hazard regime without the limited commitment constraint. The B and S regimes remain best-fitting as a result.

Due to extreme computational time requirements we are unable to estimate the UI regime for all runs in tables 9 and 10. Nevertheless, to show that our methods work in principle, in Table 10 section 7 we did a single run for the moral hazard with unobserved investment (UI) regime using the coarse, three-point grid specification and 1999-00 (c, q, i, k) Thai data (read together with line 3 in section 6). We find that for this run the UI regime achieves the highest likelihood among the six regimes we compute. This result shows evidence for a complex financial structure in the Thai data that is more constrained than pure moral hazard. Still, the UI regime is tied in this run with the S regime and barely edges B at the 10% confidence level, so our overall conclusions from before stand.

6.4 Discussion

The financial regimes we study postulate endogenous constraints on the ability of firms to adjust assets or, in other words, endogenize the degree of persistence of assets/capital k. For example, the FI regime stipulates that an agent, facing no financial constraints, would immediately adjust to the optimal capital level, k' no matter what the initial k is. Such adjustment is however subject to the incentive compatibility constraints in the MH regime and subject to even more stringent borrowing constraints (e.g., zero borrowing under savings only and autarky) in the exogenously incomplete markets regimes. A salient feature of the Thai data is that capital is very persistent and investment events are infrequent (Samphantharak and Townsend, 2009). This is also depicted on Figure 1 which plots the joint distribution of k and k' in the 1999-00 Thai data[35]. The persistence in capital favors the B (and often S) regimes overall. It is also the reason why in our robustness runs with quadratic adjustment costs the likelihood of the MH and FI regimes with the (k, i, q) and (c, q, i, k) data improves notably. On the other hand, the autarky regime is rejected in virtually all runs, as it apparently predicts excessive persistence and inability to smooth investment relative to the data.

Our models also imply endogenous, theoretical restrictions on the ability to insure consumption against income shocks, with the moral hazard model predicting more insurance (consistent with the 1999 Thai consumption-income data) than the exogenously incomplete markets regimes. However, the good fit of the MH regime with the Thai consumption and income data is not robust to some of the alternative specifications discussed in the robustness section. Overall, we find that the consumption data alone do not provide conclusive evidence on the nature of the financial regime. Indeed, our results using combined data on consumption, income and investment where the best-fitting regimes are the same as when using investment and income data alone, and where our ability to distinguish regimes is better than using consumption data alone, suggest that the type of financial constraints represented by the borrowing regime (often tied with the saving only) are the leading factor in shaping the overall patterns in the Thai data. On the other hand, for subsamples such as family networks and the Northeast region there is some evidence of moral hazard.

[35]The picture looks qualitatively the same for all years.

Finally, the results in this section can be put in perspective relative to our previous findings in PTK (2006) where we estimated a one-period model of binary occupational choice between starting a business and subsistence farming. In that paper we found moral hazard (rather than limited liability) to be the predominant source of financial constraints for rural Thai entrepreneurs, but the borrowing and saving only regimes we study here were not tested. As in this paper, PTK found evidence for differences in the best fitting regime in some specifications, e.g., when stratified by region. On the other hand, Karaivanov (2008) finds that, in an occupational choice setting similar to PTK, one cannot distinguish statistically between a model of moral hazard vs. a model of borrowing with default similar to what we find in the (c, q) cross-section specifications in this paper.

7 Conclusions

We formulate and solve numerically a wide range of models of dynamic financial constraints with exogenous or endogenous contract structure that allow for moral hazard and unobservable capital and investment. We characterize the optimal allocations implied by the regimes from both cross-sectional and intertemporal perspectives. We develop methods based on mechanism design theory and linear programming and used them to structurally estimate, compare, and statistically test between the different financial regimes. The compared regimes differ significantly with respect to their implications for investment and consumption smoothing in the cross-section and transitions. Combined consumption and investment data were found particularly useful in pinning down the financial regime generating the data. Our methods can handle unobserved heterogeneity, grid approximations, transitional dynamics, and reasonable measurement error.

One important finding is that in our baseline runs using combined consumption and investment data we can readily distinguish exogenously incomplete financial regimes from endogenously incomplete ones, where the latter are solutions to mechanism design problems with unobserved actions and state variables. As the literature we surveyed in the introduction typically takes one route or the other, we believe this ability to distinguish will prove helpful in future research and the applications of others. We are also able to distinguish within these regime groups, though this depends on measurement error, the variables in the available data set, and whether or not we have more than a single cross-section of data. Of course, we do not claim that we have covered all possible models of financial contracts, only six common prototypes. Obvious inclusions for future work are models with observed effort but unobserved ability or productivity, unobserved output (costly state verification), or limited commitment.

We are still somewhat limited on the computational side, though we are encouraged with recent advances we have been making. We had difficulty estimating the moral hazard regime with unobserved capital and investment. In an on-going collaboration with computer scientists, we have been exploring the use of parallel processing to speed up our codes and allow more complexity. What we have done thus far is, for want of better terminology, brute force. There would be further gains from more streamlined programs and more efficient search, i.e., where to refine the grids, when to use non-linear or mixed methods, the use of nested pseudo-likelihood methods, and so on.

We have also established that our methods work on actual data from villages in Thailand. We echo previous work which finds that full risk sharing is rejected, but not by much, and indeed find that the moral hazard regime is not inconsistent with the income and consumption data. However, this result is not robust to some alternative specifications as discussed in the robustness section and we find that, overall, consumption and income data alone do not provide conclusive evidence on the nature of the financial regime. In terms of investment, we confirm previous work which finds that investment is not smooth and may be sensitive to cash flow and, indeed, find that the borrowing and saving regimes seem to characterize best the investment and income data, as well as the combined consumption, income and investment data.

We do recover a more sophisticated contract theoretic regime (moral hazard constrained credit) if we restrict attention to family networks and 1999 data, confirming related work by Chiappori, Samphantarak and Townsend (2009) and Kinnan and Townsend (2009). Still, this result is not robust overall, which we suspect could be due to the infrequent nature of investment in the Thai data and the relatively large size of investment compared to capital when investment takes place. On the other hand, the financial regimes we study postulate endogenous constraints on the ability of firms to adjust assets, embedding the degree of assets persistence. The feature of the Thai data that capital is persistent thus favors the B (or S) regimes where assets adjustment is subject to more stringent constraints than in MH or FI. Evidently we have learned something from our approach, beginning to distinguish, in a sense, capital adjustment costs from financial constraints.

Natural future steps include allowing for distinctions across different technologies (fish/shrimp, livestock, business, etc.) and aggregate shocks (shrimp disease, rainfall, etc.). We would also like to return to the issue of entrepreneurial talent, as in our earlier work (PTK, 2006) and allow for heterogeneity in project returns in the actual data. Related work (Pawasutipaisit and Townsend, 2008) shows that ROA varies considerably across households and is persistent. On the other hand, such data summaries have trouble finding consistent patterns with respect to finance, suggesting the data be viewed through the lens of revised models.

We have our eyes on other economies as well, in part because we get more entry and exit from business in other countries, and in part because we need large sample sizes for our methods to work. Unfortunately, we do not typically find both consumption and investment data, which is why we chose the Thai data to begin with. Work in progress (Karaivanov, Ruano, Saurina and Townsend, 2009) with non-financial firms data from Spain shows evidence that the number of firms' bank relationships matters for whether they exhibit excess cash flow sensitivity of investment. We use the methods described in the current paper to evaluate which of four financial regimes (autarky, non-contingent debt, moral hazard and complete markets) best characterizes the degree of financial constraints for unbanked, single-banked and multiple-banked firms. Our methods allow in principle for transitions across financial regimes which is another extension we plan.

References

[1] Abraham, A. and N. Pavoni (2008), "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending: A Recursive Formulation", *Review of Economic Dynamics* (forthcoming).

[2] Abel, A. and O. Blanchard (1983), "An Intertemporal Model of Savings and Investment", *Econometrica*, 51(3), pp. 675-92.

[3] Abreu, D., D. Pearce and E. Stacchetti (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", *Econometrica* 58(5), pp. 1041-63.

[4] Aiyagari, S.R. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics*, 109(3), pp. 659-84

[5] Albuquerque, R. and H. Hopenhayn (2004), "Optimal Lending Contracts and Firm Dynamics," *Review of Economic Studies*, 71(2), pp. 285-315.

[6] Alem, M. and R. Townsend (2009), "An Evaluation of Safety Nets and Financial Institutions in Crisis and Growth", working paper, MIT.

[7] Angeletos, G. and L. Calvet (2007), "Idiosyncratic Production Risk, Growth, and the Business Cycle", working paper, MIT.

[8] Attanasio, O. and N. Pavoni (2008), "Risk Sharing in Private Information Models with Asset Accumulation", working paper, UCL.

[9] Berger, A. and G. Udell (2002), "Small Business Credit Availability and Relationship Lending: The Importance of Bank Organisational Structure", *Economic Journal*, 112, pp. F32-F53.

[10] Bitler, M., T. Moskowitz and A. Vissing-Jorgensen (2005), "Testing Agency Theory with Entrepreneur Effort and Wealth", *Journal of Finance*, 60(2), pp. 539-76.

[11] Bond, S. and C. Meghir (1994), "Dynamic investment models and the firms' financial policy", *Review of Economic Studies*, 61(2), pp. 197-222.

[12] Buera, F. and Y. Shin (2007), "Financial Frictions and the Persistence of History: A Quantitative Exploration", working paper, Northwestern.

[13] Cagetti, M. and M. De Nardi (2006), "Entrepreneurship, Frictions, and Wealth", *Journal of Political Economy*, 114(5), pp. 835-70

[14] Chiappori, P., S. Schulhofer-Wohl, K. Samphantharak, and R. Townsend (2008), "Heterogeneity and Risk Sharing in Thai Villages", working paper, University of Chicago.

[15] Clementi, G. and H. Hopenhayn (2006), "A Theory of Financing Constraints and Firm Dynamics", *Quarterly Journal of Economics*, 121(1), pp. 229-65.

[16] Cooley, T. and V. Quadrini (2001), "Financial Markets and Firm Dynamics", *American Economic Review*, 91(5), pp. 1286-1310.

[17] Covas, F. (2006), "Uninsured idiosyncratic production risk with borrowing constraints", *Journal of Economic Dynamics and Control*, 30(11), pp. 2167-90.

[18] Deaton, A. and G. Laroque (1996), "Competitive Storage and Commodity Price Dynamics", *Journal of Political Economy*, 104(5), pp. 896-923.

[19] Doepke, M. and R. Townsend (2006), "Dynamic Mechanism Design with Hidden Income and Hidden Actions", *Journal of Economic Theory,* 126(1), pp. 235-85.

[20] Ericson, R. and A. Pakes, (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work", *Review of Economic Studies*, 62(1), pp. 53-82.

[21] Fazzari, S., R. Hubbard and B. Petersen (1988), "Financing Constraints and Corporate Investment", *Brooking Papers on Economic Activity*, 1, pp. 141-206.

[22] Fernandes, A. and C. Phelan (2000), "A Recursive Formulation for Repeated Agency with History Dependence", *Journal of Economic Theory*, 91, pp. 223-247.

[23] Geanakoplos, J. and H. Polemarchakis (1986), "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When Markets are Incomplete", in: W.P. Heller, R.M. Ross, D.A. Starrett (eds.): *Uncertainty Information and Communication. Essays in Honor of Kenneth Arrow Vol. 3*, Cambridge.

[24] Heaton, J. and D. Lucas, (2000), "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk", *Journal of Finance*, 55(3), pp. 1163-98.

[25] Honore, B. and E. Tamer (2006), "Bounds on Parameters in Panel Dynamic Discrete Choice Models", *Econometrica*, 74(3), pp. 611-30.

[26] Hopenhayn, H. (1992), "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(2), pp. 1127-50.

[27] Jappelli, T. and L. Pistaferri (2006), "Intertemporal Choice and Consumption Mobility", *Journal of the European Economic Association*, 4(1), pp. 75-115.

[28] Jewitt, I. (1988), "Justifying the First-Order Approach to Principal-Agent Problems,", *Econometrica*, 56(5), pp. 1177-90.

[29] Judd, K. (1998), *Numerical Methods in Economics*, MIT Press, Cambridge, MA.

[30] Kaplan and Zingales (2000), "Investment and Cash Flow Sensitivities are not Valid Measures of Financing Constraints", *Quarterly Journal of Economics*, 115(2), pp. 707-12.

[31] Karaivanov, A. (2008), "Financial Constraints and Occupational Choice in Thai Villages", working paper.

[32] Karaivanov, A., S. Ruano, J. Saurina and R. Townsend (2009), "Formal Financial Access for Older Firms; Does it Matter for Firms' Investment?", working paper.

[33] Kinnan, C. and R. Townsend (2009), "Kinship Networks, Financial Access and Consumption Smoothing", working paper, MIT

[34] Kocherlakota, N. (2004), "Figuring Out the Impact of Hidden Savings on Optimal Unemployment Insurance", *Review of Economic Dynamics* 7(3), pp. 541-554.

[35] Ligon, E. (1998), "Risk Sharing and Information in Village Economies", *Review of Economic Studies* 66(4), pp. 847-64.

[36] Ligon, E., J. Thomas and T. Worrall, (2002) "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies", *Review of Economic Studies*, 69(1), pp. 209-44.

[37] Meh, C. and V. Quadrini (2006), "Endogenous Incomplete Markets With Investment Risks", *Journal of Economic Dynamics and Control*, 30(11), pp. 2143-65.

[38] Myers, S. and N. Majuf (1984), "Corporate financing and investment decisions when firms have information that investors do not have", *Journal of Financial Economics*, 13(2), pp. 41-56.

[39] Paulson, A., R. Townsend and A. Karaivanov (2006), "Distinguishing Limited Liability from Moral Hazard in a Model of Entrepreneurship", *Journal of Political Economy*, 114(1), pp. 100-144.

[40] Pawasutipaisit, A. and R. Townsend (2008), "Enterprise and the Growth of Wealth", working paper, MIT.

[41] Phelan, C. (1994), "Incentives, Insurance and the Variability of Consumption and Leisure", *Journal of Economic Dynamics and Control*, 18, pp. 581-99.

[42] Phelan, C. and R. Townsend (1991), "Computing Multi-Period, Information-Constrained Equilibria", *Review of Economic Studies*, 58, pp. 853-881.

[43] Prescott, E.C. and R. Townsend (1984), "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica*, 52, pp. 21–45.

[44] Prescott, E. S. (2003), "Communication in Models with Private Information: Theory and Computation", *Geneva Papers on Risk and Insurance Theory* 28(2), pp. 105-30.

[45] Rivers, D. and Q. Vuong, (2002), "Model Selection Tests for Nonlinear Dynamic Models", *The Econometrics Journal*, 5, pp. 1-39.

[46] Rogerson, W. (1985), "The First Order Approach to Principal-Agent Problems", *Econometrica* 53, pp. 1357-68.

[47] Rossi-Hansberg, E. and M. Wright (2007), "Establishment Size Dynamics in the Aggregate Economy", *American Economic Review*, 97(5), pp. 1639-66.

[48] Samphantharak, K. and R. Townsend (2009), *Households as Corporate Firms: An Analysis of Household Finance Using Integrated Household Surveys and Corporate Financial Accounting*, Cambridge University Press.

[49] Schmid, L. (2008), "A Quantitative Dynamic Agency Model of Financing Constraints", working paper, Duke University.

[50] Spear, S. and S. Srivastava (1987), "On Repeated Moral Hazard with Discounting", *Review of Economic Studies,* 53, pp. 599-617.

[51] Townsend, R. (1982), "Optimal Multi-period Contracts and the Gain from Enduring Relationships Under Private Information", *Journal of Political Economy* 61, pp. 166-86.

[52] Townsend, R., A. Paulson and T. Lee (1997), *Townsend Thai Project Surveys*, http://cier.uchicago.edu

[53] Vuong, Q. (1989), "Likelihood Ratio Tests for Model Selection and Non-nested Hypotheses." *Econometrica*, 57(2), pp. 307-333.

[54] Werning, I. (2001), "Repeated Moral Hazard with Unmonitored Wealth: A Recursive First Order Approach", working paper, MIT.

[55] Zame, W. (1993), "Efficiency and the Role of Default when Security Markets are Incomplete", *American Economic Review*, 83, pp. 1142-64.

[56] Zeldes, S. (1989), "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence", *Quarterly Journal of Economics*, 104(2), pp. 275-98.

Table 1 - Problem Dimensionality

Number of:	linear programs	variables	constraints
Autarky	5	75	16
Saving / Borrowing	25	625	16
Full Information	25	3,125	17
Moral Hazard	25	19,375	23
Unobserved investment, stage 1	250	1,650	122
Unobserved investment, stage 2	550	11,625	2,506
Unobserved investment, total	137,500	n.a.	n.a.

Note: This table assumes the following grid sizes that we use in the estimation: #Q=5, #K=5, #Z=3, #B=5, #T=31 #W=5; and #**W**=50 and #**W**m=110 for moral hazard with unobserved investment

Table 2 - Variable Grids

Variable	grid size, #	grid range
income/cash flow, Q	5	{.05, 2}
assets, K	5	[0, 1]
effort, Z	3	[.01, 1]
savings/debt, B	6 (5 for S regime)	S: [-2, 0], B: [-2, .9]
transfers/consumption	31	[10^{-3}, 1.2]
promised utility, W	5	[w_{min}, w_{max}]

Table 3 - Baseline Parameters

Parameter	Value(s)
depreciation rate, δ	0.05
agent's discount factor, β	0.95
principal's discount factor, 1/R	0.95
risk aversion, σ	0.5
effort curvature, θ	2
technology parameter, ρ	0

Table 4 - Parameter Estimates using Simulated Data from the MH model regime

Simulated (k,i,q) data

Model	γ_{me}	σ	θ	ρ	$\mu_{w/b}$	$\gamma_{w/b}$	**LL Value**[2]
Moral Hazard (base) - MH *	0.1015 *(.008)*	0.1944 *(.114)*	1.9749 *(.345)*	0.0771 *(.012)*	0.2380 *(.158)*	0.1985 *(.064)*	**-2.8138**
Full Information - FI *	0.1012 *(.007)*	0.2182 *(.079)*	1.9902 *(.308)*	0.0773 *(.052)*	0.8000 *(.002)*	0.1271 *(.015)*	**-2.8140**
Borrowing & Lending - B	0.1031 *(.008)*	0.7008 *(.031)*	8.0625 *(.000)*	-50.000 *(.000)*	0.5148 *(.068)*	0.1626 *(.081)*	**-2.8817**
Saving Only - S	0.1163 *(.010)*	2.1151 *(.020)*	5.0975 *(.333)*	-50.000 *(.000)*	0.2062 *(.050)*	0.3880 *(.027)*	**-2.8921**
Autarky - A	1.0000 *(.021)*	0.0000 *(.000)*	1.2000 *(.150)*	3.0129 *(.009)*	n.a.	n.a.	**-3.7253**
baseline values	*0.1000*	*0.5000*	*2.0000*	*0.0000*	*0.5000*	*0.3500*	

Simulated (c,q) data

Model	γ_{me}	σ	θ	ρ	$\mu_{w/b}$	$\gamma_{w/b}$	**LL Value**[2]
Moral Hazard (base) - MH *	0.1024 *(.007)*	0.5003 *(.036)*	1.9788 *(.239)*	0.2329 *(.390)*	0.4994 *(.015)*	0.3915 *(.024)*	**-2.4696**
Full Information - FI	0.1021 *(.006)*	0.4814 *(.019)*	0.0350 *(.022)*	-50.000 *(4.26)*	0.5436 *(.012)*	0.3210 *(.020)*	**-2.5273**
Borrowing & Lending - B	0.1168 *(.003)*	4.9151 *(.030)*	0.8590 *(.000)*	-50.000 *(.000)*	0.5701 *(.051)*	0.9968 *(.043)*	**-2.5274**
Saving Only - S	0.1219 *(.007)*	0.0000 *(.000)*	1.5782 *(.000)*	-4.6791 *(.001)*	0.2378 *(.052)*	0.9999 *(.000)*	**-2.5456**
Autarky - A	0.5901 *(.018)*	0.0000 *(.000)*	1.2000 *(.000)*	-4.8495 *(.000)*	n.a.	n.a.	**-2.7222**
baseline values	*0.1000*	*0.5000*	*2.0000*	*0.0000*	*0.5000*	*0.3500*	

Simulated (c,q,i,k) data

Model	γ_{me}	σ	θ	ρ	$\mu_{w/b}$	$\gamma_{w/b}$	**LL Value**[2]
Moral Hazard (base) - MH *	0.1001 *(.004)*	0.5013 *(.022)*	1.9854 *(.139)*	0.0773 *(.052)*	0.5105 *(.013)*	0.3724 *(.015)*	**-4.2163**
Full Information - FI	0.1370 *(.005)*	0.1629 *(.007)*	2.6233 *(.020)*	0.0584 *(.017)*	0.6047 *(.010)*	0.2404 *(.009)*	**-4.5799**
Borrowing & Lending - B	0.2240 *(.019)*	0.0000 *(.000)*	1.7975 *(.000)*	-2.9849 *(.005)*	0.3231 *(.017)*	0.3390 *(.015)*	**-4.7664**
Saving Only - S	0.2753 *(.013)*	2.2115 *(.000)*	4.9279 *(.159)*	-50.000 *(.000)*	0.4016 *(.015)*	0.2798 *(.014)*	**-4.7676**
Autarky - A	0.7209 *(.016)*	0.0000 *(.000)*	1.2000 *(.000)*	3.0129 *(.000)*	n.a.	n.a.	**-5.3922**
baseline values	*0.1000*	*0.5000*	*2.0000*	*0.0000*	*0.5000*	*0.3500*	

Notes: 1. Bootstrap standard errors in the parentheses; 2. Normalized (divided by n) log-likelihood values; 3. $\mu_{w/b}$ and $\gamma_{w/b}$ reported relative to grid span

* denotes the best fitting regime (including tied)

All estimates reported above are for the low measurement error specification.

Table 5 - Model Comparisons[1,2] Using Simulated Data from the MH Model Regime - Vuong Test Z-Statistics

Comparison	MH v FI	MH v B	MH v S	MH v A	FI v B	FI v S	FI v A	B v S	B v A	S v A	Best Fit
A. (k,i,q) simulated data											
1. low measurement error	0.68(tie)	5.50***(MH)	4.44***(MH)	33.9***(MH)	5.48***(FI)	4.42***(FI)	34.0***(FI)	0.78(tie)	27.1***(B)	25.1***(S)	MH,FI
2. high measurement error	0.66(tie)	-0.40(tie)	1.77*(MH)	8.83***(MH)	-0.42(tie)	1.76*(FI)	8.83***(FI)	2.14**(B)	8.95***(B)	8.70***(S)	MH,FI,B
B. (c,q) data simulated											
1. low measurement error	5.70***(MH)	5.59***(MH)	6.10***(MH)	12.7***(MH)	0.01(tie)	1.08(tie)	9.28***(FI)	1.83*(B)	11.1***(B)	9.27***(S)	MH
2. high measurement error	0.88(tie)	1.66*(MH)	1.80*(MH)	4.75***(MH)	0.95(tie)	0.89(tie)	4.10***(FI)	0.04(tie)	4.63***(B)	4.45***(S)	MH,FI
C. (c,q,i,k) simulated data											
1. low measurement error	7.60***(MH)	15.4***(MH)	17.7***(MH)	37.9***(MH)	4.05***(FI)	4.63***(FI)	14.4***(FI)	0.05(tie)	15.0***(B)	17.4***(S)	MH
2. high measurement error	3.13***(MH)	2.87***(MH)	3.27***(MH)	10.2***(MH)	-0.61(tie)	-0.03(tie)	8.37***(FI)	1.66*(B)	9.73***(B)	10.1***(S)	MH
D. (c,q) panel simulated data											
1. high m.e.; t = 0, 1	1.71*(MH)	6.08***(MH)	14.8***(MH)	18.2***(MH)	4.62***(FI)	13.4***(FI)	17.1***(FI)	12.0***(B)	15.5***(B)	15.1***(S)	MH
2. high m.e.; t = 0, 50	1.59(tie)	5.11***(MH)	16.8***(MH)	16.7***(MH)	3.43***(FI)	15.2***(FI)	15.6***(FI)	17.1***(B)	15.0***(B)	2.20**(S)	MH

NOTES:

1. The data-generating regime is MH; n = 1,000
2. *** = 1%, ** = 5%, * = 10%, "tie" >10% two-sided test significance level - the better fitting regime is in the parentheses.
3. "high m.e." - high measurement error; "low m.e." - low measurement error

Table 6 - Model Regime Comparisons[1,2] Using Simulated Data - Robustness

Comparison	MH v FI	MH v B	MH v S	MH v A	FI v B	FI v S	FI v A	B v S	B v A	S v A	Best Fit
A. Using (c,k,i,q) data from the borrowing (B) regime											
1. low measurement error	13.6***(MH)	-13.3***(B)	2.44**(MH)	21.0***(MH)	-20.4***(B)	-9.66***(S)	-10.7***(FI)	12.9***(B)	32.1***(B)	22.8***(S)	B
2. high measurement error	3.73***(MH)	-0.61(tie)	1.39(tie)	8.89***(MH)	-3.76***(B)	-2.99***(S)	6.33***(FI)	2.78***(B)	9.51***(B)	8.86***(S)	B
B. Additional robustness runs with simulated (c,q,i,k) data from the moral hazard (MH) model regime											
1. without meas. error	12.2***(MH)	29.0***(MH)	16.6***(MH)	57.4***(MH)	0.88(tie)	0.33(tie)	10.5***(FI)	-0.65(tie)	18.7***(B)	13.3***(S)	MH
2. n = 200; high m.e.	1.44(tie)	1.03(tie)	1.61(tie)	5.49***(MH)	-0.40(tie)	0.09(tie)	4.57***(FI)	0.98(tie)	5.35***(B)	5.34***(S)	MH,FI,B,S
3. n = 5,000; high m.e.	6.90***(MH)	5.74***(MH)	5.69***(MH)	22.9***(MH)	-1.55(tie)	-1.08(tie)	18.7***(FI)	0.89(tie)	22.2***(B)	22.8***(S)	MH
4. coarser grids; low m.e.	6.55***(MH)	14.1***(MH)	14.4***(MH)	42.4***(MH)	8.27***(FI)	8.08***(FI)	36.8***(FI)	2.36**(B)	29.7***(B)	26.8***(S)	MH
5 productivity heterogeneity	7.53***(MH)	13.8***(MH)	18.1***(MH)	35.4***(MH)	3.38***(FI)	5.77***(FI)	12.5***(FI)	4.38***(B)	13.8***(B)	15.1***(S)	MH
C. Simulated data from the moral hazard (MH) regime at the Thai data estimated parameters											
1. (c,q,i,k) data	17.2***(MH)	8.94***(MH)	8.71***(MH)	8.05***(MH)	-11.7***(B)	-12.1***(S)	-12.5***(A)	-0.16(tie)	-1.14(tie)	-1.28(tie)	MH
2. (k,i,q) data	-0.78(tie)	6.17***(MH)	6.12***(MH)	6.56***(MH)	6.17***(FI)	6.12***(FI)	6.56***(FI)	1.02(tie)	1.74*(B)	0.89(tie)	MH,FI
3. (c,q) data	3.69***(MH)	6.33***(MH)	6.65***(MH)	9.42***(MH)	2.20**(FI)	3.84***(FI)	7.95***(FI)	3.50***(B)	9.02***(B)	10.5***(S)	MH
D. Run with the unobserved investment (UI) regime with simulated (c,q,i,k) data from the moral hazard (MH) model regime											
	UI v MH	**UI v FI**	**UI v B**	**UI v S**	**UI v A**						
coarser grids; low m.e.	-7.59***(MH)	-2.89***(FI)	5.31***(UI)	5.73***(UI)	38.8***(UI)						MH

NOTES:

1. The data-generating financial regime is moral hazard with observed investment (MH); the sample size is n = 1,000
2. *** = 1%, ** = 5%, * = 10%, "tie" >10% two-sided test significance level - the better fitting regime is in the parentheses.
3. "m.e." - measurement error

Table 7 - Thai Data Summary Statistics[1,2]

Variable	mean	median	std. dev.	min	max
Assets, k	80,298	13,688	312,010	0	8,083,800
Net Income from production, q	128,700	65,016	240,630	-1,473,100	5,211,800
Consumption expenditure, c	64,172	47,868	53,284	4,210	610,860

1. Sample size is 531 households observed over 7 years (1999-2005). Units are Thai baht.
2. The summary statistics are computed for the pooled data.

Table 8 - Parameter Estimates Using 1999-00 Thai Data

Investment and income, (k,i,q) data

Model	γ_{me}	σ	θ	ρ	$\mu_{w/b}$	$\gamma_{w/b}$	LL Value[2]
Moral Hazard (base) - MH	0.4497 *(.027)*	1.8943 *(.015)*	2.8991 *(.003)*	-4.0854 *(.000)*	0.0001 *(.000)*	1.0000 *(.000)*	**-4.2259**
Full Information - FI	0.4845 *(.007)*	2.3348 *(.079)*	0.0052 *(.308)*	-4.0854 *(.052)*	0.9400 *(.002)*	0.0684 *(.015)*	**-4.2126**
Borrowing & Lending - B *	0.3890 *(.020)*	1.9970 *(.000)*	5.2958 *(.000)*	-50.000 *(.000)*	1.0000 *(.000)*	0.0000 *(.000)*	**-4.0872**
Saving Only - S *	0.3638 *(.017)*	9.7346 *(.000)*	2.0000 *(.000)*	-50.000 *(.000)*	0.9999 *(.000)*	0.2164 *(.000)*	**-4.0870**
Autarky - A	0.4310 *(.016)*	9.1306 *(.000)*	1.2000 *(.000)*	-50.000 *(.000)*	n.a.	n.a.	**-4.1236**

Consumption and income, (c,q) data

Model	γ_{me}	σ	θ	ρ	$\mu_{w/b}$	$\gamma_{w/b}$	LL Value[2]
Moral Hazard (base) - MH *	0.1366 *(.009)*	0.5148 *(.000)*	0.1086 *(.000)*	0.0235 *(.057)*	0.4258 *(.011)*	0.1014 *(.013)*	**-2.6031**
Full Information - FI	0.1993 *(.009)*	0.5326 *(.017)*	0.8912 *(.067)*	-5.1380 *(2.09)*	0.5044 *(.025)*	0.0619 *(.024)*	**-2.6625**
Borrowing & Lending - B *	0.1442 *(.011)*	1.1542 *(.080)*	7.6853 *(.000)*	0.3880 *(.035)*	0.8750 *(.070)*	0.0000 *(.012)*	**-2.6198**
Saving Only - S	0.1519 *(.009)*	8.6168 *(.247)*	8.8090 *(.000)*	-0.2559 *(.755)*	0.5082 *(.088)*	0.6934 *(.092)*	**-2.6334**
Autarky - A	0.2574 *(.011)*	0.0000 *(.000)*	0.6150 *(.333)*	0.6947 *(.175)*	n.a.	n.a.	**-2.7616**

Combined consumption, investment, and income, (c,q,i,k) data

Model	γ_{me}	σ	θ	ρ	$\mu_{w/b}$	$\gamma_{w/b}$	LL Value[2]
Moral Hazard (base) - MH	0.3733 *(.017)*	0.2223 *(.033)*	0.0636 *(.035)*	-4.0854 *(.001)*	0.0001 *(.008)*	0.6108 *(.269)*	**-5.7982**
Full Information - FI	0.5701 *(.080)*	0.2223 *(.017)*	0.0213 *(.003)*	7.0052 *(.000)*	0.2666 *(.033)*	0.0281 *(.012)*	**-6.1113**
Borrowing & Lending - B *	0.3350 *(.011)*	1.6952 *(.342)*	242.14 *(.000)*	-50.000 *(.000)*	1.0000 *(.000)*	0.0000 *(.004)*	**-5.5018**
Saving Only - S *	0.3196 *(.014)*	9.6861 *(.383)*	8.7020 *(.000)*	-34.457 *(.000)*	0.9061 *(.027)*	0.1893 *(.020)*	**-5.5166**
Autarky - A	0.3870 *(.012)*	8.8584 *(.000)*	0.9682 *(.000)*	-50.000 *(.000)*	n.a.	n.a.	**-5.6028**

Notes: 1. Bootstrap standard errors in the parentheses; 2. Normalized (divided by n) log-likelihood values; 3. $\mu_{w/b}$ and $\gamma_{w/b}$ reported relative to grid span

* denotes the best fitting regime (including tied)

Table 9 - Model Regime Comparisons[1,2,3] Using Thai Data - Baseline; Vuong Test Z-Statistics

Comparison	MH v FI	MH v B	MH v S	MH v A	FI v B	FI v S	FI v A	B v S	B v A	S v A	Best Fit
1. Using (k,i,q) data											
1.1. years: 99-00	-1.81*(FI)	-6.80***(B)	-6.70***(S)	-4.79***(A)	-5.34***(B)	-5.54***(S)	-3.51***(A)	-0.02(tie)	3.36***(B)	2.00*(S)	B,S
1.2. years: 04-05	7.30***(MH)	1.59(tie)	-0.01(tie)	3.63***(MH)	-2.15**(B)	-4.20***(S)	-0.08(tie)	-3.24***(S)	9.28***(B)	5.83***(S)	S,MH,B
2. Using (c,q,i,k) data											
2.1. years: 99-00	6.20***(MH)	-10.7***(B)	-9.97***(S)	-8.18***(A)	-11.7***(B)	-11.6***(S)	-9.45***(A)	0.94(tie)	6.71***(B)	3.50***(S)	B,S
2.2. years: 04-05	4.31***(MH)	-9.83***(B)	-9.35***(S)	-4.90***(A)	-10.0***(B)	-11.2***(S)	-6.40***(A)	1.48(tie)	10.9***(B)	5.79***(S)	B,S
3. Using (c,q) data											
3.1. year: 99	4.79***(MH)	1.23(tie)	1.82*(MH)	6.26***(MH)	-2.31**(B)	-1.34(tie)	3.41***(FI)	1.16(tie)	7.22***(B)	8.20***(S)	MH,B
3.2. year: 05	3.92***(MH)	-0.56(tie)	0.15(tie)	2.77***(MH)	-3.96***(B)	-4.33***(S)	-0.60(tie)	0.75(tie)	4.58***(B)	3.11***(S)	B,MH,S
4. Two-Year Panel											
4.1. (c,q), yrs: 99 and 00	14.2***(MH)	-10.1***(B)	-4.45***(S)	-2.24**(A)	-20.5***(B)	-17.7***(S)	-13.0***(A)	2.08**(B)	5.24***(B)	4.21***(S)	B
4.2. (c,q), yrs: 99 and 05	14.1***(MH)	-11.2***(B)	-6.87***(S)	-7.25***(A)	-17.3***(B)	-14.7***(S)	-16.5****(A)	0.02(tie)	0.06(tie)	0.05(tie)	B,S,A

NOTES:

1. *** = 1%, ** = 5%, * = 10% two-sided significance level, the better fitting regime is in the parentheses
2. Z-statistics cutoffs: 2.575 = *** 1.96 = ** 1.645 = * "tie"
3. Investment, *i* is constructed from the firm assets data as $i = k' - (1 - \delta)k$ with $\delta = .05$

Table 10 - Model Regime Comparisons[1,2,3] Using 1999-00 Thai Data - Robustness; Vuong Test Z-Statistics

Comparison	MH v FI	MH v B	MH v S	MH v A	FI v B	FI v S	FI v A	B v S	B v A	S v A	Best Fit
1. Risk neutrality											
1.1 (c,q) data	3.57***(MH)	-10.2***(B)	-10.2***(S)	-8.98***(A)	-10.2***(B)	-10.8***(S)	-8.29***(A)	1.63(tie)	5.17***(B)	5.18***(S)	B,S
1.2 (k,i,q) data	-1.85*(FI)	-6.55***(B)	-2.06***(A)	-5.21***(A)	-5.07***(B)	-1.20(tie)	-3.01***(A)	6.51***(B)	6.56***(B)	-1.57(tie)	B
1.3 (c,q,i,k) data	14.4***(MH)	-9.46***(B)	-6.19***(S)	-12.1***(A)	-15.9***(B)	-15.7***(S)	-17.7***(A)	5.18***(B)	3.53***(B)	-0.79(tie)	B
2. No measurement error											
2.1 (c,q) data	4.92***(MH)	0.47(tie)	1.84**(MH)	9.28***(MH)	-2.10**(B)	-0.39(tie)	8.62***(FI)	1.33(tie)	9.56***(B)	9.79***(S)	MH,B
2.2 (k,i,q) data	0.81(tie)	-2.78***(B)	0.71(tie)	12.5***(MH)	-2.84***(B)	0.65(tie)	12.6***(FI)	3.61***(B)	11.4***(B)	9.64***(S)	B
2.3 (c,q,i,k) data	-2.93***(FI)	0.14(tie)	7.86***(MH)	11.2***(MH)	2.52**(FI)	9.63***(FI)	11.8***(FI)	8.37***(B)	11.8***(B)	4.33***(S)	MH,B
3. Networks sub-sample											
3.1. (c,q) data	2.13**(MH)	1.92*(MH)	1.97**(MH)	5.99***(MH)	-0.28(tie)	-0.05(tie)	3.66***(FI)	0.34(tie)	6.53***(B)	7.33***(S)	MH
3.2 (c,q,i,k) data	11.9***(MH)	-6.96***(B)	-7.29***(S)	-4.27***(A)	-14.9***(B)	-13.8***(S)	-13.7***(A)	0.22(tie)	5.56***(B)	3.27***(S)	B,S
3.3 (c,q) panel 99 and 00	14.9***(MH)	-8.12***(B)	-2.92***(S)	-1.24(tie)	-19.9***(B)	-16.0***(S)	-13.5***(A)	3.45***(B)	6.18***(B)	3.38***(S)	B
4. Investment adjustment costs											
(k,i,q) data	0.23(tie)	-0.42(tie)	-1.39(tie)	3.01***(MH)	-0.49(tie)	-1.29(tie)	1.32(tie)	-4.31***(S)	4.71***(B)	4.68***(S)	B,S,MH,FI
(c,q,i,k) data	15.1***(MH)	-1.60(tie)	-1.87*(S)	2.22**(MH)	-15.5***(B)	-15.7***(S)	-14.4***(A)	-0.62(tie)	5.18***(B)	2.87***(S)	S,B
5. Stratified by Region											
Central, (c,q,i,k) data	5.28***(MH)	-9.98***(B)	-5.07***(S)	-0.96(tie)	-9.16***(B)	-8.91***(S)	-6.24***(A)	1.81*(B)	10.8***(B)	5.27***(S)	B
Central, (c,q) data	-0.25(tie)	-2.18**(B)	-1.82*(S)	0.90(tie)	-1.86*(B)	-1.81*(S)	1.32(tie)	0.08(tie)	4.79***(B)	4.76***(S)	B,S
North-East, (c,q,i,k) data	12.8***(MH)	0.38(tie)	-0.59(tie)	4.12***(MH)	-12.7***(B)	-13.7***(S)	-10.1***(A)	-1.30(tie)	3.60***(B)	3.48***(S)	MH,B,S
North-East, (c,q) data	1.62(tie)	1.85*(MH)	3.76***(MH)	5.78***(MH)	-0.29(tie)	1.99**(FI)	4.49***(FI)	3.52***(B)	5.52***(B)	6.15***(S)	MH,FI
6. Other runs											
networks v.2; (c,q) data	0.88(tie)	0.23(tie)	2.39**(MH)	3.89***(MH)	-0.34(tie)	1.85*(FI)	3.33***(FI)	2.91***(B)	5.78***(B)	3.79***(S)	MH,B,FI
FE removed (c,q,i,k)	5.77***(MH)	-12.3***(B)	-10.9***(S)	1.74*(MH)	-15.9***(B)	-15.7***(S)	-4.82***(A)	3.78***(B)	15.4***(B)	13.6***(S)	B
coarser grids (c,q,i,k)	12.7***(MH)	-3.00***(B)	-3.21***(S)	-2.08**(A)	-12.4***(B)	-16.2***(S)	-12.5***(A)	0.02(tie)	2.71***(B)	1.88*(S)	B,S
alternative k definition, (c,q,i,k)	1.90*(MH)	-4.59***(B)	-1.51(tie)	2.72***(MH)	-4.56***(B)	-2.10**(S)	1.02(tie)	3.08***(B)	7.48***(B)	4.88***(S)	B
limited commitment, (c,q,i,k)	-0.21(tie)	-11.7***(B)	-11.6***(S)	-9.45***(A)	-14.4***(B)	-15.0***(S)	-11.5***(A)	0.94(tie)	6.71***(B)	3.50***(S)	B,S
7. Run with the unobserved investment (UI) regime											
	UI v MH	**UI v FI**	**UI v B**	**UI v S**	**UI v A**						
coarser grids (c,q,i,k)	4.09***(UI)	12.7***(UI)	1.66*(UI)	0.99(tie)	3.20***(UI)						UI,B,S

NOTES:

1. *** = 1%, ** = 5%, * = 10% two-sided significance level, the better fitting regime is in the parentheses

positive sign of the Vuong statistic indicates that the first regime in the pair being compared has higher likelihood; negative sign indicates the opposite

2. Z-statistics cutoffs: 2.575 = *** 1.96 = ** 1.645 = * "tie"

3. Investment, i is constructed from the household assets data as $i = k' - (1 - \delta)k$ with $\delta = .05$

Figure 1: Joint Distribution of k and k' in the Thai Data