

The effects of financial development in the short and long run

Scott Fulford*

This draft: May 2011
First draft: April 2008

Abstract

Although many view financial access as a means of reducing poverty or increasing growth, empirical studies have produced contradictory results. One problem is that most studies cover only a short time frame and do not consider dynamic effects. I show that introducing credit creates a boom in consumption and reduces poverty initially, but eventually reduces mean consumption because credit substitutes for precautionary wealth. Using new consistent consumption data, my empirical findings show that increased access to bank branches in rural India increased consumption initially and reduced poverty, but consumption later fell and poverty rose.

JEL classification: O16, D91

*Boston College Department of Economics, 140 Commonwealth Ave, Chestnut Hill, MA 02467; email: scott.fulford@bc.edu. I am grateful to Angus Deaton, Chris Paxson, Sam Schulhofer-Wohl, Donald Cox, Fabio Schiantarelli, Susanto Basu, and numerous seminar participants for helpful thoughts and comments.

1 Introduction

There are two views of financial development. One holds that financial development is a crucial contributor to growth. This view drove India's efforts to extend bank branches to rural areas in the 1970s and 1980s, the more recent expansion of microcredit to hundreds of millions of households across the world, and the 2006 Nobel Peace Prize to Muhammad Yunus and the Grameen Bank for their work extending financial services to the poor. Yet where credit has long been available a darker and more pessimistic view of the consequences of better financial access also exists. Concerns about over-indebtedness, debt spirals, and farmer suicide mark this view, leading to the worry that credit reinforces poverty rather than alleviating it.

The empirical evidence for whether access to financial services helps the poor is contradictory, despite a strong correlation between financial development and growth across countries.¹ While Burgess and Pande (2005) find that the large expansion of branch banks into rural India in the 1970s and 1980s significantly reduced poverty, Kochar (2005) and Panagariya (2008) disagree. Microcredit has been the subject of a similar debate, with some studies finding benefits for the poor (Khandker, 2005; Pitt and Khandker, 1998), others questioning the evidence (Morduch, 1998; Roodman and Morduch, 2009), and recent experimental and quasi-experimental studies finding mixed results (Banerjee et al., 2009; Kaboski and Townsend, 2009). Yet the conflicting evidence has not stopped practitioners from making strong claims about the positive impact of microfinance.

What is missing from most studies is a coherent treatment of dynamic effects. I show in a general model of intertemporal consumption that the observable effects of gaining access to credit vary in a substantial and fundamental way over time. In the short run credit has large positive effects, but in the long run access to credit decreases consumption and can increase poverty. Access to credit in the short run allows people to consume more, which explains why it sometimes looks like such an effective way to help the poor. When people save for precautionary reasons, however, credit substitutes for keeping wealth in the long run. Lower wealth means lower income, as farmers mortgage their land, for example, or sell their livestock. So any estimate of the effect of credit will change depending on when it is measured: initially it appears to have very positive effects, but later on credit may look like it has harmed a community, and both views need to be considered in context. One reason for the contradictory empirical results is that since most studies do not collect the evidence necessary to separate out the long term from the short term, the measured effects should vary from study to study even when looking at the same program in the same country.

Allowing for changes over time requires consistent estimates of consumption and poverty over a long enough time period to convincingly capture dynamics, a strong requirement. Perhaps the

¹See Levine (2005) for a critical review of the theoretical issues and empirical questions on the effects of financial development on growth.

largest expansion in access to credit took place during the branch building boom in India during the 1970s and 1980s, and I use this expansion to examine whether there are important changes over time. Although the Indian National Sample Surveys cover a long period of time, they are not entirely consistent with each other. The most well known problems are from the 1999-2000 large round (Deaton and Kozel, 2005; Tarozzi, 2007), but other rounds are inconsistent as well. I take a direct approach and reconstruct consistent consumption and poverty estimates from the item level consumption for each household. Using the model to guide the empirics, I find that rural areas in which the branches per capita increased had a boom in consumption and reduced poverty initially. But consumption later fell and poverty rose. In the the long term, the total effect of the changes is that poverty increased slightly while consumption was unchanged or increased slightly, leaving open the possibility of benefits to production. The dynamic effects of credit matter both theoretically and empirically.

Why do the effects of credit vary fundamentally over time? The intuition comes from thinking about the ultimate consumer, rather than the investments he or she makes.² Households keep wealth for many reasons, but an important one is to protect the household from bad events. Indeed, for a long enough lived consumer that is the only reason to keep wealth—at least the only reason that does not lead to unlimited accumulation (Schechtman and Escudero, 1977). In the short term credit allows more consumption, but in the long term households adjust their wealth holdings. Access to credit means they no longer need to keep as much wealth to receive the same level of insurance, which eventually reduces wealth and income.³ These effects come not because modern-day Shylocks are taking advantage of the poor, but through optimal choices that households make over time.

There is a growing empirical literature on the role of financial development in consumption smoothing by the poor.⁴ The theoretical model in this paper clarifies several areas of ambiguity. In particular, credit can be welfare improving, even if it does not increase income, but measuring

²This paper complements a large literature which considers the effects of credit on production. The overlapping generation poverty-trap models (Banerjee, 2001; Banerjee and Newman, 1993; Galor and Zeira, 1993) typically allow for non-convexity in some form of investment such as in education or health, but limit the consumption decision. Similarly, Giné and Townsend (2004) fit a model of occupational choice to Thai data. Bencivenga and Smith (1991) and Acemoglu and Zilibotti (1997) emphasize the effects of financial intermediation on diversification in overlapping generations models. Townsend and Ueda (2006) build an infinite horizon intertemporal choice model based on Greenwood and Jovanovic (1990) which they fit to Thai data. In their model financial access allows for diversification and a better rate of return, but not credit.

³The mechanism in this paper is distinct from Jappelli and Pagano (1994) who show in an overlapping generation model with no uncertainty that when there are positive externalities from capital accumulation, forced savings by the young may increase income and improve welfare if the externality is large enough. That a reduction in the long term is possible was suggested, but not proved, by Aiyagari (1994). There is also a literature which examines changes in uncertainty, which has a similar effect to changes in credit (Krusell et al., 2009).

⁴Menon (2006) finds that microcredit aids in seasonal smoothing in Bangladesh over 18 months. Gertler, Levine, and Moretti (2009) find that microfinance helps families smooth health shocks in Indonesia. That loans, even when ostensibly for productive purposes, may be used for consumption, is highlighted by Dichter (2007).

improvements in welfare may be difficult since the observable benefits are likely to vary over time. Credit may have negative long-term consequences on consumption and poverty as it leads “booms” in the short term and “busts” in the long term. Finally, while credit aids consumption smoothing initially, in the long term credit will not necessarily help households smooth, since while credit adds to wealth used as self-insurance in the short term, it substitutes for it in the long term. Both views of financial development are valid, but understanding financial development requires seeing both the good and the bad, and understanding that they come together.

2 Modeling changes in financial access in a buffer-stock economy

While much of the literature which examines the distributional aspects of incomplete financial markets focuses on the investment decision, this paper focuses on the dynamics of consumption in a community which gains access to credit. I model a community of long-lived households who face uninsurable income shocks and must decide how much to consume today and how much to save for tomorrow.⁵ Households would prefer to consume more now, but save to create a buffer of wealth which allows them to self-insure against bad shocks in the coming periods.

What happens when the residents of such a community suddenly gain access to credit from an outside source—and what can an economist who can observe consumption and investment, but not utility, learn from those changes? On gaining access to credit everyone in the community, even those who do not borrow today, immediately consumes more. Credit looks great for this community based on the observables; poverty has fallen and everyone is consuming more than before. Yet credit is the promise to pay in the future, and from its peak right after the change, I show that mean consumption in the community falls continuously, and in the long term, mean consumption is *lower* than before. Measured in the long term, credit appears to have harmed the community since consumption and income have fallen. Credit, at least within the model, does improve welfare, but not as much as the initial rise in consumption would suggest. Of course, improving welfare depends in part on having a constant discount rate; dividing up generations, rather than examining total family utility, may alter the positive welfare result, as might other ways in which there is a present bias for consumption (Laibson, 1997).

The key to these results comes from the self-insurance nature of savings. With credit, the command over resources of everyone in the community suddenly expands, making them effectively wealthier. This extra wealth allows them to consume more, since they no longer need to delay consumption as much to maintain a buffer. Yet in the long term, members of the community do not

⁵This “income fluctuations” problem has received much attention. Schechtman and Escudero (1977) provide an early treatment; Clarida (1987) extends the results. Deaton (1991) considers the “buffer-stock” nature of this model, with results that are extended by Rabault (2002). Very similar models are considered by Carroll, Hall, and Zeldes (1992), and Carroll (2001).

need to maintain as much wealth for self-insurance. As the wealth of the community falls, so does its income since the wealth paid a return. Then in the new ergodic distribution, consumption must fall since the income of the community has fallen. These results hold generally, with no additional conditions on utility or income beyond those required to establish buffer-stock behavior.

2.1 The model

The consumer has the standard infinite horizon, intertemporally additive, expected utility maximization problem with exponential discounting:

$$\max_{\{c_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right], \quad (1)$$

where c_t is consumption in period t subject to an accumulation equation for disposable wealth w :

$$w_{t+1} = (1 + r)(w_t - c_t) + \rho h + y_{t+1}, \quad (2)$$

and a borrowing constraint that must be met every period:

$$c_t \leq w_t + \alpha B; \quad (3)$$

where: (1) the instantaneous utility function $u(\cdot)$ belongs to the set U of increasing, strictly concave, and differentiable functions that satisfy $u'(0) = \infty$; (2) the discount factor is $\beta \in (0, 1)$; (3) labor income y_{t+1} is i.i.d. and has distribution G with density g , whose support has a lower bound y_l , and $E_t[y_{t+1}] < \infty$; (4) the initial wealth $w_0 \geq y_l$ is given; (5) the rate of return on liquid wealth is $r > 0$; (6) the rate of return on illiquid wealth is $\rho > 0$; (7) the illiquid wealth is $h > 0$, which can potentially be converted into liquid wealth; (8) the degree of illiquidity is given by $\alpha \in [0, 1]$; and (9) the maximum that can be borrowed and paid back almost surely from the illiquid asset is $B = (\rho/r)h > 0$. I consider variations in some of these assumptions later, particularly the restrictions on the income process and rate of return. Income does not need to be i.i.d., and with some additional restrictions can be non-stationary.

The illiquid asset h and the returns r and ρ determine the amount that a bank or another institution might be willing to lend against the asset. A bank that pays r on its deposits must earn at least r on its loans. A consumer who borrows B must be able to pay back rB every period. Then the largest value of B for which h can serve as collateral is $B = (\rho/r)h$. If $r = \rho$, this simplifies to $B = h$, which implies that since the bank gets the same return as the individual, selling the asset is equivalent to taking a loan on it.

The parameter α then determines how much can be borrowed. With $\alpha = 0$, the illiquid wealth

can only be saved: it provides income ρh in every period, but is illiquid in the sense that the principal cannot be consumed now. The illiquid wealth could be human capital, for example, which raises the expected labor earnings by ρh each period. Alternatively it could be land or property which provides rental or crop income, but is difficult or impossible to sell.⁶ When $\alpha = 1$, the consumer may borrow up to $B = (\rho/r)h$, the most that can be paid back from the income from the asset. If $y_t = 0$ then this limit is the “natural debt limit” (Aiyagari, 1994), otherwise it represents the natural limit on borrowing from a particular asset.⁷

The model can be rewritten to one in which no borrowing is allowed, and is thus equivalent to the standard intertemporal consumption model with additive income fluctuations.⁸ The reason to employ two parameters α and h to form the debt limit is because changes in α and h have different effects that a change in a single parameter B would confound. If either α or h is zero, of course, no borrowing is possible; there are two ways not to be able to borrow: have no future income to borrow against so that no one will lend to you ($h = 0$), or live in an area with incomplete markets so that no one is able to lend to you ($\alpha = 0$). Increases in h have both lending effects and income effects since h affects both the constraint and the income earned each period. To examine changes in financial access, however, it is changes in α that are important.

When does a stationary distribution of wealth and consumption exist? For a given α and h this problem is very well studied, particularly by Schechtman and Escudero (1977), Deaton (1991), and most recently by Rabault (2002) who provides the most general treatment. If people are sufficiently impatient so that they generally prefer consuming now, given the price, to later ($\beta(1 + r) < 1$); if they are not so risk averse that they will keep saving indefinitely as insurance against bad shocks ($\lim_{c \rightarrow \infty} u''(c)/u'(c) = 0$); and if the worst thing that can happen does not result in infinite marginal utility ($y_t > 0$), then a stationary distribution of consumption and wealth exists, and the borrowing constraint will bind from time to time for everyone.⁹

⁶Such restrictions may occur because markets do not exist, or because there is a strong social stigma to selling family land as was the case in at least some parts of rural India (Srinivas, 1976). Alternatively, if there is a large specific human capital element in the returns on the asset, there may be little or no value from selling it, but it can still be used as collateral if credit markets exist. See Kiyotaki and Moore (1997) for a model in which assets have specific human capital and are worth much less to anyone else.

⁷One way to think of h is as a certain income stream pulled from the stochastic income process y_{t+1} . Viewed this way, the underlying asset—say land—may change in value, which will be captured in y_{t+1} . h is the portion of that asset which can act as collateral since it holds its value in all states.

⁸See Rabault (2002) or Aiyagari (1994) for examples of the change of variables to convert a fixed borrowing limit into an equivalent problem with no borrowing limit.

⁹See the mathematical appendix for additional detail on this result, which is a special case of a more general result from Rabault (2002) who allows for $y_t \geq 0$, and so for marginal utility to be potentially unbounded. In the more general case, however, borrowing constraints will not necessarily bind, although they will come arbitrarily close to binding almost surely with some additional restrictions. There are a number of variations on this result in the literature. Clarida (1987) assumes that the income process has a zero lower bound, but that marginal utility is bounded: $u'(0) > \infty$. Lemma 4 in the appendix establishes that it is possible to rewrite the current model into the model of Clarida (1987), and so the results obtained there are applicable here. The case when there is a positive probability of

A buffer-stock community, as used here, is a mass of buffer-stock consumers which is small and open to the rest of the world so that prices are given and capital flows need not balance. At any point in time in a buffer-stock community, there are some relatively wealthy families, and some who have consumed all of their assets and enter the next period with only their labor income. Over time, although the identities of the rich and poor change, the distribution does not. Buffer-stock models seem to be a reasonable way to describe family dynasties in village economies which have not changed much for long periods of time. Srinivas (1976, pp. 107-118), for example, describes the “rise and fall of lineages” in a village in the Indian state of Karnataka, in a way that suggests that the model description of a stationary distribution with dynasties moving around within it is a good one. With extended families living together, moreover, the age and life-cycle considerations of individuals members become less important, leaving only the long term consequences of household consumption, exactly the situation the model describes. This extended family approach is supported by the data used in the empirical section: when I break the National Sample Surveys into age cohorts, I find almost no spread of inequality with the age of the household head as one might expect if households savings are dependent on head age (Deaton and Paxson, 1994). Other work points to the importance of assets used as buffers for consumption smoothing in developing countries.¹⁰ Moreover, a growing literature suggests that even in developed countries precautionary savings are important and necessary to understand consumption and savings patterns.¹¹

2.2 Changes in the liquidity constraint

What happens when the ability to borrow in a buffer-stock community increases? Mean consumption initially increases, but from then steadily declines to a new steady state which is below the previous one, so that in the long term communities with greater lending have lower mean consumption. This result depends on the self-insurance nature of savings: Buffer-stock consumers have wealth largely as a precaution against future poor income draws. Allowing them to borrow

zero earnings has been studied extensively by Carroll, Hall, and Zeldes (1992), Carroll (1997), and Carroll (2001). Schechtman and Escudero (1977) demonstrate the need for some restrictions on risk aversion, and prove a somewhat more restrictive result.

¹⁰Deaton (1997, 389-90) suggests that the panel data on households in the ICRISAT villages—rural villages in the Indian states of Maharashtra and Karnataka—are most consistent with the autarkic buffer-stock model described here. Rosenzweig and Wolpin (1993) point to the use of bullocks in India as asset buffers, although their dual use as lumpy productive assets complicates the analysis. Kazianga and Udry (2006) similarly examine live-stock holdings during a drought in Burkina Faso.

¹¹Gourinchas and Parker (2002) find that even for individuals in the United States, the buffer-stock model is a good description of consumption habits until life-cycle concerns come to dominate sometime after age 40, while Hubbard, Skinner, and Zeldes (1995) find that precautionary motives are a necessary addition to life-cycle models in explaining US wealth distributions. Caballero (1991) suggests that precautionary savings represent a substantial portion of US wealth. Gross and Souleles (2002) present evidence on credit card debt and liquidity constraints in the United States that supports this model and the predictions in the next section. Fuchs-Schündeln (2008) finds precautionary motive is necessary for understanding German consumption and savings following reunification.

against more of their future income has a similar effect to increasing their wealth in the current period, and makes them immediately better off (proposition 1). Since they are now able to borrow, in addition to using their own savings, consumers are now very well insured. But consumers do not want so much insurance, and so they immediately start consuming more (proposition 2). Since they are consuming more, the wealth of the community starts to fall. This wealth paid a positive return, and so when the new stationary distribution is reached, the mean income after paying interest on debts of the entire community is lower. In the stationary distribution, this implies that consumption is lower (proposition 3). In the long term, loosening liquidity constraints reduces mean consumption. All proofs are in the appendix.

Reducing liquidity constraints cannot lead to steady growth (see corollary 2 in the appendix) but still improves welfare. Proposition 1 shows that for an individual with any given level of wealth, loosening liquidity constraints improves utility. Since we are increasing the choice set, such a change cannot reduce utility since the same consumption path is still available, but showing that it improves utility requires the consumers to actually be liquidity constrained at some point, a more stringent requirement.

Proposition 1. *Under assumptions 1, 2, and 3, let $V(w_0, \alpha) = E_0 [\sum_{t=0}^{\infty} \beta^t u(c_t)]$ be the value function of following the optimal policy $c(w_t, \alpha)$. Then $V(\cdot, \cdot)$ is strictly increasing in its second argument.*

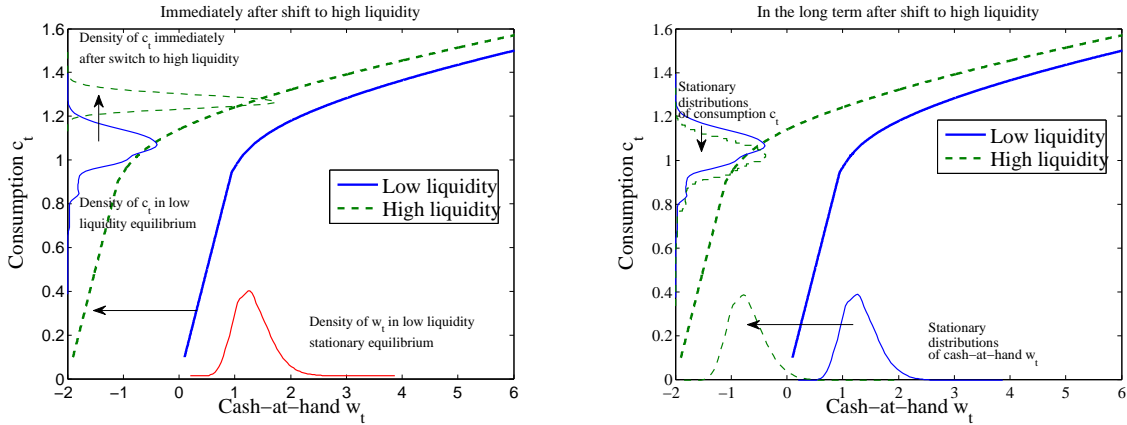
A more subtle point is that relaxing liquidity constraints (raising α) will increase consumption for every wealth level, whether or not the individual is actually constrained at that wealth level.

Proposition 2. *For a given level of wealth for which both consumption functions are defined ($w \geq y_l - \alpha^0 h$), increasing liquidity raises consumption: for $\alpha^1 > \alpha^0$, $c(w, \alpha^1) > c(w, \alpha^0)$.*

Corollary 1. *For a given distribution of wealth $F(w)$, define the distribution of consumption as $J(c, \alpha) = F(c^{-1}(c, \alpha))$ where $c^{-1}(c, \alpha)$ is the inverse of the consumption function $c(w, \alpha)$ which gives the wealth associated with every level of consumption. Then relaxing liquidity constraints, so that $\alpha^1 > \alpha^0$ implies the following: (1) The distribution of consumption with higher liquidity stochastically dominates consumption with lower liquidity; (2) Mean consumption strictly increases; (3) Poverty, measured as the headcount ratio for some poverty line z , is non-increasing: $J(z, \alpha^0) \geq J(z, \alpha^1)$.*

Relaxing liquidity constraints improves utility and increases consumption immediately for everyone, and causes poverty to fall (assuming poverty is not already zero). This change is illustrated in figure 1 which shows the consumption policies for low liquidity ($\alpha = 0$) and high liquidity ($\alpha = 1$), and the shift up in the distribution of consumption on the left axis. The consumption function for high liquidity is everywhere higher than for low liquidity. The distribution of wealth is unchanged, however, and so the entire distribution of consumption shifts up.

Figure 1: Consumption policy and distributions



Notes: The stationary distribution of cash-at-hand in the left panel is the simulated distribution for the low liquidity policy, which is used to obtain both distributions of consumption on the vertical axis. Each distribution is smoothed using a normal kernel density with bandwidth 0.08 for cash-at-hand and 0.02 for consumption. The consumption distributions are the (kernel) distributions of consumption immediately before and immediately after a change from low to high liquidity. In the right panel the stationary distributions of cash-at-hand correspond to the distributions for consumption on the vertical axis. The community is composed of one million individuals with identical CRRA preferences with $\gamma = 2$, $\beta = .95$. The income process is a nine point Gaussian quadrature approximation of a log normal with location parameters $\sigma = 0.2$ and $\mu = -\sigma^2/2$. The rate of return is $r = 0.02$. Each individual has illiquid wealth of 2 which goes from totally illiquid to totally liquid. The simulations use function approximation and convergence routines from Miranda and Fackler (2002).

What happens in the long term, however, is not obvious. Higher consumption now means less accumulated wealth in the future. In the steady state with the relaxed liquidity constraint, consumption at every wealth level is higher, but if the distribution of wealth is low enough, mean consumption will be lower. The following proposition establishes that the fall in wealth is sufficient to reduce consumption.

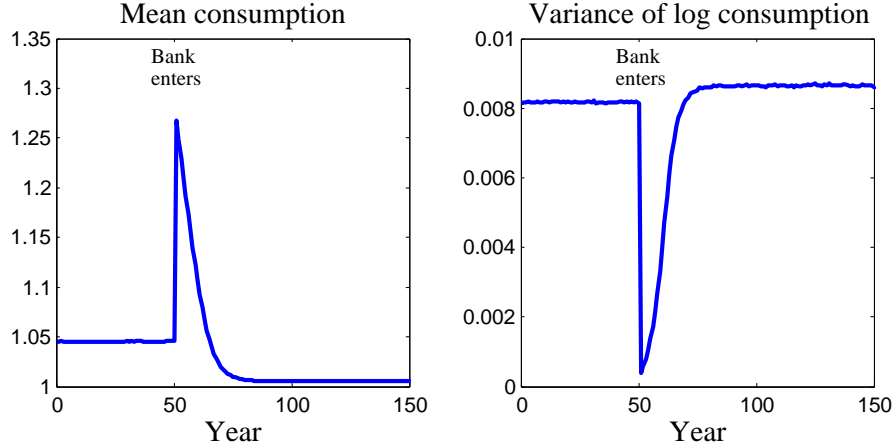
Proposition 3. *Mean consumption is lower in the stationary distribution with looser liquidity constraints: Let $F(w, \alpha)$ be the stationary distribution of wealth with liquidity parameter α . For $\alpha^1 > \alpha^0$, $\int c(w, \alpha^1) dF(w, \alpha^1) < \int c(w, \alpha^0) dF(w, \alpha^0)$.*

Then combining propositions 2 and 3, consumption increases immediately after liquidity constraints are relaxed, but then falls steadily to a new mean that is lower than before. The path is shown in the left panel of figure 2. That mean consumption is non-increasing follows directly from the (weak) stochastic dominance of wealth at every period used to show proposition 3.

2.2.1 Lending and inequality

While the evolution in mean consumption is well defined, what happens to inequality after a change in liquidity constraints is less clear. Figure 1 suggests some characteristics that may be general,

Figure 2: Consumption and inequality of a community with changing liquidity constraints



Notes: Wealth worth twice the mean income is converted from illiquid to liquid in year 50. See the notes for figure 1 for parameter values.

however. When marginal consumption goes to a constant as is true for CRRA utility, then a shift left of the consumption function decreases consumption inequality, since the wealth distribution is producing a consumption distribution in a region of the consumption function with less curvature. As the wealth distribution starts to shift down to the new steady state, the curvature of the consumption function over the support of the wealth distribution increases, and so does consumption inequality. The path is illustrated in the right panel of figure 2. In the simulations presented, consumption inequality is slightly higher in the stationary distribution, driven by an increase in the curvature of the consumption function.¹²

In the model described, inequality comes solely from differing shocks to individuals who are the same in all other respects, but allowing for other sources of heterogeneity is straightforward. While the path of adjustment will not be the exactly the same for subgroups with different preferences or income process, the general shape will be the same as long as the income process and preferences fit the model assumptions. The population path would then be a weighted average of subgroup paths. Yet under some circumstances the paths of different subgroups with different mean incomes will be exactly the same up to some multiplicative constant all along the path of adjustment.

Proposition 4. *Let all individuals in all subpopulations have the same discount factor, and CRRA preferences: $u(c) = c^{1-\gamma}/(1-\gamma)$. The income process in each subpopulation are multiplicatively shifted so that $G^j(\mu^j y) = G^i(\mu^i y)$ for $\forall y$ and each subpopulation i and j , where $\mu^j > 0$ is*

¹² By rewriting both functions in terms of “effective wealth” $z_t = w_t + \alpha h$, the high liquidity function would lie below the consumption function for low liquidity above the point where the high liquidity consumption function transitions from consuming all available wealth to saving some for the next period. Then the high liquidity function is more concave around this point, which seems to drive the higher consumption inequality in the stationary distribution.

the mean income of each subpopulation. Further, the illiquid wealth is shifted in the same way: $h^j/\mu^j = h^i/\mu^i$ for each i and j . Then in the stationary distribution for any liquidity constraint, and the transition from one stationary distribution to a new one after a change in the liquidity constraint which is the same across all subpopulations, the distribution of consumption and wealth will be multiplicatively related as well: $J_t^j(\mu^j c) = J_t^i(\mu^i c)$ for $\forall c$ and each subpopulation i and j , and each t starting with a stationary distribution.

Very usefully, the proposition implies that when taking logs, differences in income produce the same consumption and savings behavior up to an additive constant. If $\ln c_t^j$ is the mean log consumption for subpopulation j at time t , then $\ln c_t^j - \ln \mu^j = \ln c_t^i - \ln \mu^i$ for all subpopulations i and j and time t starting from a stationary distribution. In other words, income differences do not affect identification, since by including fixed effects the effect of credit on communities with different levels of income will be *exactly* the same in logs.

2.2.2 Some extensions

This section considers several extensions to the basic lending model. The model links the liquid rate of return and the borrowing rate, and requires that both be positive. Equating the two returns can be justified by an appeal to a competitive market for loans and savings, but this justification is strained by the model motivation of opening a new bank branch in a previously unbanked area. Allowing different rates of return for borrowing and savings, say r_b and r_s with $r_b > r_s$ introduces a kink in the asset accumulation equation. A kink complicates results since if r_b is high enough, the stationary distribution may not include any borrowing. Indeed, if $r_b = \infty$ then this is equivalent to $\alpha = 0$. With no one liquidity constrained, changing α is irrelevant, since the consumers are constrained not by the ability to borrow but their own willingness to pay the high costs of borrowing. Differentiating borrowing and lending costs does not overturn the results, but adds a caveat: if no one is liquidity constrained in the stationary distribution because borrowing costs are too high, then loosening liquidity constraints has no effect. If consumers are liquidity constrained from time to time, then they benefit when the liquidity constraints are relaxed, and will consume more immediately and less later on. Linking the returns is clearly a simplification, but it allows the direct use of the powerful results on the standard buffer-stock model in proposition 6 in the appendix.

For the poor in a developing country the assumption of positive returns on wealth may not be a good one. If we allow negative returns so that $r_s < 0 < r_b$ the analysis becomes even more complicated. Liquid assets now have value only as a buffer since the consumer has to pay to keep them, but has to pay at least as much to borrow. Borrowing allows the consumer to carry less of the costly positive wealth as a buffer, which increases income. The downside is that the consumer

has to borrow from time to time, which decreases income, and so mean community income may increase or decrease in the long run. In the short term, there will still be an increase in consumption followed by a steady decline, as in figure 2. While liquid assets with negative returns may be a reasonable assumption for agricultural communities without access to banking, when a bank enters a community and offers lending, it will generally also take deposits and pay a positive return. The introduction of banking will then have the dual effect of allowing borrowing, and raising the rate of return which I consider briefly in the next section.

It is straightforward to allow for more general income processes with serial correlation, although the stationarity of assets and consumption has not been shown generally.¹³ The intuition for the path of mean consumption remains the same. Even with some serial correlation, everyone will have enough bad shocks to exhaust their resources eventually. Then allowing borrowing is beneficial even for those who have had good shocks in the past, and so are likely to have high earnings in the future, just as it is for those who have acquired large amounts of wealth and so will continue to have high income in the future. Everyone consumes more now, and so less later since they reduce their stock of income producing assets.

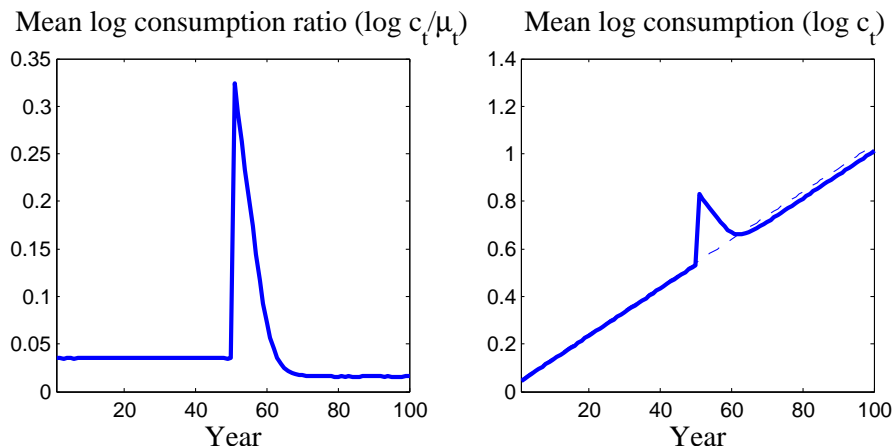
Developing countries may be growing countries and so the assumption of a stationary income process may not be reasonable. The following proposition shows that under some circumstances the ratio of consumption to the underlying growing income process follows the same path described without growth.

Proposition 5. *Let preferences be CRRA: $u(c) = c^{1-\gamma}/(1-\gamma)$. Define $y_t = \mu_t \epsilon_t$ where ϵ_t has the same restrictions as the income process used in the basic model: it is i.i.d. with finite mean and positive lower bound. Allow the illiquid assets to scale with μ_t as well so that $h_t = \mu_t h$ are the illiquid assets at time t . Let $\mu_{t+1}/\mu_t = (1+g)$ with $g \geq 0$ non-stochastic and $\mu_0 > 0$. Then if $\beta(1+r)(1+g)^{-\gamma} < 1$ and $r > g$, for a population of such consumers all of whom have the same μ_t , the ratio $\hat{c}_t = c_t/\mu_t$ has a stationary distribution and follows the path described in propositions 2 and 3.*

Under the assumptions in the proposition, the long-term effect of lending on log consumption is a downward shift in the path. Figure 3 illustrates this shift, showing the path of the mean of $\ln \hat{c}_t$ and the path of the mean $\ln c_t$ for a community with a 1% growth rate. Otherwise preferences and the stationary portion of income are the same as in figure 2. The stationary distribution of the log consumption ratio is lower after lending in the first panel, and the path of consumption is lower in the second panel, as marked by the dashed line which shows the path without credit. Note that the change after lending is a level shift in the path of log consumption in the long run, not a change

¹³See Deaton (1991) and Rabault (2002), where the major cost is added notation since it is no longer possible to reduce the state space to “cash-at-hand.”

Figure 3: Consumption in a community with growth and changing liquidity constraints



Notes: The top panel shows the mean log ratio of consumption to the non-stationary portion of income: $\ln(c_t/\mu_t)$. The bottom panel shows the mean of consumption. Wealth worth twice the mean income is converted from illiquid to liquid in year 50. The community is composed of 10,000 individuals with identical CRRA preferences with $\gamma = 2$, and $\beta = 0.95$. The i.i.d. part of the income process is a nine point Gaussian quadrature approximation of a log normal with location parameters $\sigma = 0.2$ and $\mu = -\sigma^2/2$. Initial income $\mu_0 = 1$. The rate of return is $r = 0.02$ and growth is $g = 0.01$. The simulations use approximation and convergence routines from Miranda and Fackler (2002).

in the growth rate. This specification suggests, moreover, the circumstances under which it makes sense to remove the trend in empirical specifications, and still recover the underlying effects. In particular if $\ln c_t$ is the mean log consumption at time t and $\hat{c}_t = c_t/\mu_t$ the ratio of consumption to mean income, then $\ln c_t = \ln \hat{c}_t + \ln \mu_t$. Moreover, $\ln \hat{c}_t$ is constant in the stationary distribution, increases immediately when liquidity constraints are loosened, and falls to a new steady state below the original, just like consumption for stationary distributions.

Finally, suppose different populations have different growth rates and different starting incomes. For example, some states may be better governed than others and so grow faster. Then it is possible to combine propositions 4 and 5 to allow for not just different levels of income but different growth rates of income as well. The paths in logs of such communities are not exactly the same but will follow the same increase then decrease to a new lower steady state, and so the problem is the same as differences in preferences considered earlier.¹⁴ An important consequence of all of these proofs is that banks may enter endogenously seeking profits in fast growing communities or forced by the government into poor communities, but the endogeneity does not affect identification since the dynamic effects in logs with trends and fixed effects are the same.

¹⁴Knowing the growth rates, it is possible to be even more precise, since the consumer's problem for communities with different growth rates can be rewritten to be exactly the same in ratios up to a discount rate of $\beta(1+g)^{1-\gamma}$. Then in logs faster growing communities want to borrow more, and move more to the present, since marginal utility is decreasing more in the future.

2.3 Changes in the rate of return

While the previous section examines what happens when the liquidity constraint changes, this one briefly examines changes in consumption and wealth when the rate of return r increases. One possible role for banks is to offer better returns on savings to people in rural areas. Even very low but positive returns in a bank account may be better than other available means of saving for the future which may even have negative returns: rice and grains saved for the next year can rot or be consumed by pests. Banks may offer better returns, through, for example, better access to projects, ability to differentiate between good and bad projects, or increasing returns to scale available to large pools of capital. For small changes, the effects of the increased return are ambiguous in both the short and long term, because the competing income and substitution effects of changing a price, in this case r , are present intertemporally as well. Some people may choose to bring consumption forward since they now have higher income at any wealth level, others may choose to delay consumption since the cost of consuming now is higher. If the increase in the rate of return is large enough to convert buffer-stock consumers into accumulators of wealth, then consumption and wealth will grow without bound. With a high enough rate of return, the long-term consequences provide a striking contrast to the stationary distribution that prevailed before. The cross-sectional consumption path will have an immediate decline or rise, depending on whether the substitution or income effect is dominant, followed by ever increasing levels of consumption which eventually surpass any bound.

If the discount factor or the rate of return differs across or within communities, moreover, inequality will increase. Suppose bank branches open in some communities, but not others, putting the communities with new bank branches on an accumulation path. The newly patient communities will start to diverge and inequality will increase. It is cross-community inequality that drives the increase as some communities become increasingly wealthy and consume more. A similar observation holds if some groups are more patient than others. When a bank comes in the more patient groups may start accumulating, leading to a long term rise in inequality. Of course, if the very rich migrate to the cities, or refuse to participate in surveys, it will be very difficult to measure such changes.¹⁵

What happens if both credit and rate of return change at the same time? Theory has little guidance since in general the results are ambiguous. Big enough changes in the rate of return and strong enough income or substitution preferences can swamp a small change in credit. Changes in the rate of return will either enhance the short term effect of credit if better returns bring consumption forward, or dampen them if the better returns cause increased savings. Since credit has a large initial effect, and a smaller long term effect, it may dominate in the short term, but that will clearly

¹⁵The working paper version of this paper has an extended analysis of changes in inequality and consumption after an increase in the rate of return.

depend on the magnitudes of the changes. One reading of the empirical results in this paper, which suggest a large increase in consumption in the short term declining to a small one in the long term, is that they are consistent with such a combined effect.

3 Simulated paths

The theory gives some clear predictions about the path of mean consumption, and some suggestions as to the evolution of inequality for a given community with increased liquidity or higher rates of return. One key result is that the relationship between changes in borrowing constraints and consumption depends crucially on how, and particularly when, the relationship is measured. Moreover, the relationship will be obscured by different communities gaining access at different times, by shocks at a regional or community level that affect everyone, and by underlying diversity. To examine how these results will affect empirical strategies and what approaches to real world data are likely to correctly capture the dynamics, I produce a simulated economy with many communities gaining access to banking and regions with common regional shocks, and restrict the analysis to aggregates at a regional level. This approach creates a panel of aggregates over time as in many empirical studies (Burgess and Pande, 2005; Khandker, 2005; King and Levine, 1993). In the next section, I perform the same analysis on regional level aggregates from India. The preferences, income process, and changes in credit underlying the simulations remain largely educated guesses, however, relying on small surveys from one part of India, or estimates from the developed world, and so I place little weight on the actual magnitudes. The point of the simulations is to show that with reasonable parameters the effects described by the theory are important, and that estimates of a single “effect” of banking depends sensitively on the timing of measurement.

The simulated economy consists of 50 regions (or countries) each divided into 100 communities with 1000 households in each—giving an underlying population of five million households. The underlying preferences are the same as those used to produce figure 2: CRRA preferences with $\gamma = 2$, and $\beta = 0.95$, which are similar to those used or estimated in recent studies of the United States (Carroll, Dynan, and Krane, 2003; Gourinchas and Parker, 2002; Storesletten, Telmer, and Yaron, 2004). The (real) rate of return is 2%. The income process has three multiplicative components which in the cross-section have a mean of one: a permanent household component that is constant over time for a household but varies across the population; a stochastic household component that is independent of all other households and independent over time; and a stochastic regional component that is the same for all households within a region, but independent of other regions and over time. The cross-sectional variance of the income for the community has a coefficient of variation of 0.53, somewhat lower than 0.90, the mean of the three ICRISAT villages

reported in Townsend (1994), which is likely to contain some measurement error.¹⁶

Communities gain access to banking services, which convert these illiquid assets to liquid assets, at different times. At the end of year 1970, after running for 50 years from an initial draw of income, the first communities gain access to banking. To create regions with different paths, the regions are divided into two blocks. In the first block, in each year from 1970 to 1979, 75 communities per year gain access to banking, randomized across regions. For the next 10 years, 100 communities gain access year. In the second block, 50 communities gain access for the first decade and 75 a year for the next decade. By year 1990, three fifths of all communities have access. No more communities gain access after year 1990. This path is chosen to mimic the path of rural branch expansion in India starting in 1970 and ending in 1990 and shown in figure 4 in which there was a rapid buildup of access to banking which ends after approximately 20 years. Since the choice of which communities get banks is random in the simulated economy, there will also be randomization in the proportion of communities with banks at the end of the buildup, which provides additional regional variation. Gaining a branch is outside of the control of individuals within a community, and they do not include the possibility of getting a branch in their decisions until a branch actually appears.

The underlying “true” path for a single community on gaining access to a bank is similar to that shown in figure 2.¹⁷ With the assumed discount factor of 0.95, consumption only falls below pre-banking steady state after approximately 15 years, and takes as much as 30 years to approach the new steady state. The length of this adjustment, and size of the initial increase in consumption, is sensitive to the assumed discount factor and rate of return, with higher impatience giving a larger initial increase and more rapid adjustment. Given the speculative nature of the underlying preferences and income process, the conclusion I draw from this exercise is that with reasonable

¹⁶Labor income in each period t for household j in region s takes the form $y_{jt} = y_j^u y_{jt}^i y_{st}^r$, where y_j^u is fixed for the household over time, but varies across the population; y_{jt}^i is a household idiosyncratic shock; and y_{st}^r is a regional shock that is shared by all households in the region. Since there are no general equilibrium effects, the household only cares about the combined income y_{jt} . The components y_{jt}^i and y_{st}^r are i.i.d. over time, and each component is independent of the others, and of other households and regions. The distribution of each subcomponent of y_{jt} is a five point Gaussian quadrature approximate of a lognormal, with variance parameters s_u^2 , s_i^2 and s_r^2 and location parameters which give each subcomponent at mean of 1. The overall distribution is thus approximately a lognormal with mean one, and variance $e^{s_u^2 + s_i^2 + s_r^2} - 1$. The variance parameters of the subcomponents are chosen so that the idiosyncratic shock is the most important $s_i^2 = 0.6s^2$, followed by the underlying permanent component $s_u^2 = 0.3s^2$, and the regional component $s_r^2 = 0.1s^2$.

¹⁷The difference is that figure 2 does not have a permanent component of income or regional shocks, but the underlying preferences are otherwise the same. Mean log consumption before a bank enters is 0.01, which comes from a population mean labor income of one plus the return on illiquid assets and saved wealth. On gaining a bank, mean log consumption immediately increases by 0.17, and thereafter declines. The decline is around 0.013 each year for the first five years, with smaller declines after that. The total effect of gaining a bank is -0.038, which is the sum of each yearly effect, or the difference between the mean steady state before and after. These effects are sizable effects: since they are measured in log consumption, the initial change is an increase in consumption of approximately 17%, but that depends on the size of illiquid wealth which can now be borrowed against.

guesses the effects of credit can be large, and adjustment can be lengthy.

If one were to take these data without knowledge of the underlying process, and try to calculate the adjustment path, what approaches would succeed? Regressions that do not include lags are substantially biased depending on what part of the data one considers. Table 1 shows the results of regressing mean log consumption in the region on the proportion of communities with bank branches in the region over different periods and including different lags. Below each coefficient is a standard error in parentheses from the regression, and in brackets the interval in which the coefficients from 95% of 200 additional draws of “worlds” of 50 regions fall. The sum of the coefficients, with a 95% coverage interval from the 200 world draws are on the last two rows. A regression that included only the current proportion of communities with branches in a region gives a dramatically different picture of the effect of banks depending on what subsample is used. Banks first enter communities at the end of year 1970, and a regression which observes consumption in years 1971 through 1990 will find that banks increase consumption in at least 95% of world draws. A shift of ten years, to a subsample of years 1981 through 2000, changes the picture entirely, estimating that banks cause a large decrease in consumption. Thus including only a contemporaneous banking variable in a limited sample will not provide an estimate of any parameter of interest in general, and the coefficient on this variable will also be unstable, producing different results over different time periods.


In contrast, including lags out to 15 years catches most of the important dynamics for either period. The lags miss some information, but since much of the adjustment takes place within 15 years, they capture most of the important changes—both subsamples have a large positive followed by negatives. Neither gets the sum of lags quite right, and whether the initial effect is overestimated or underestimated depends on the sample, but both are closer than with only the contemporaneous banking variable, and they are reasonably stable across subsamples. As shown in the last column, however, if consumption can only be observed in certain years, the sample may not be large enough, either in time periods or in cross section, to separate out the effects of new bank branches from the underlying regional volatility. The last column restricts the observations of consumption to years in which one of the Indian National Sample Surveys used in the next section occurred: 1983, 1987, 1993, 1999, 2000, 2001, 2003, and 2005. The regression limited to this subsample captures much of the dynamics, but it is sensitive to the simulation assumptions, and the estimates are biased. These results suggest that we should not expect too much from such a limited subsample, even one spread over 23 years.

Table 1: Regressions from simulated data

Dependent Variable: Log Consumption x 100					
	Consumption Sample				
	1970-1990	1980-2000	1970-1990	1980-2000	NSS subsample
Proportion of communities with a bank					
1 lag	5.970*** (0.597) [4.0,9.6]	-7.951*** (1.112) [-11.1,-1.0]	23.68*** (4.582) [3.5,27.1]	22.47*** (4.989) [8.2,29.4]	17.19 (14.58) [-4.5,42.9]
3 lags			-15.04** (6.236) [-14.1,8.3]	-4.746 (7.020) [-14.1,9.5]	-6.733 (19.70) [-43.8,38.2]
5 lags			1.558 (5.242) [-18.1,6.2]	-8.556* (5.123) [-17.1,3.6]	2.974 (14.81) [-38.5,26.9]
10 lags			-7.733* (4.194) [-19.3,7.9]	-3.046 (3.791) [-18.9,2.4]	-8.201 (5.630) [-18.2,8.2]
15 lags			-4.858 (4.560) [-17.3,9.2]	-8.434** (3.504) [-16.0,5.1]	-7.614** (3.189) [-14.8,1.2]
Constant	2.370*** (0.204) [0.7,3.1]	7.605*** (0.592) [4.4,9.7]	1.207*** (0.319) [-0.5,2.5]	0.334 (0.761) [-2.1,3.0]	0.581 (2.357) [-3.7,5.2]
Observations	1000	1000	1000	1000	400
Regions	50	50	50	50	50
R-squared	0.095	0.051	0.121	0.231	0.312
Fixed Effects	Yes	Yes	Yes	Yes	Yes
Sum of lags	5.970	-7.951	-2.393	-2.311	-2.382
Mean all draws	6.786 [4.0,9.6]	-6.173 [-11.1,-1.0]	-0.254 [-10.7,9.4]	-1.157 [-6.4,4.5]	-1.337 [-9.9,6.3]

Notes: From simulated data. The NSS subsample is with consumption observed only in years 1983, 1987, 1993, 1999, 2000, 2001, 2003, and 2005; the proportion of communities with banks is observed in every year. Communities start to gain access to banking services in 1970. It is assumed that the entire history of whether a community has access is known but consumption is limited to the subsample. In brackets are bounds within which fall 95% of the coefficients from the regressions of 200 additional draws of "worlds" of 50 regions. The brackets are the 5th and 195th coefficient from the regressions.

Figure 4: Number of bank branches in India



Notes: See separate data appendix for sources and details for construction.

4 The dynamic effects of bank branches in India

4.1 The data

India nationalized its major banks in 1969, and proceeded to institute several social banking policies that encouraged the placement of bank branches in rural unbanked areas. There were several different policies, with somewhat different emphases, but the net result was an explosion of rural bank branches (Shah, Rao, and Shankar, 2007). Figure 4 shows the total number of rural and urban bank branches by year using data from the Reserve Bank of India (RBI), India's central bank.¹⁸ With the liberalization following the currency crisis in 1990-1991, there was much less emphasis on putting new branches in rural areas, but there was a general policy to keep branches from closing which continues today (Leeladhar, 2008). These new bank branches both expanded credit and increased the bank deposits in rural areas, a process which continued even after the branch expansion slowed (Shah, Rao, and Shankar, 2007).¹⁹

The Indian National Sample Surveys (NSS) are large nationally representative surveys that examine major aspects of Indian life. Major rounds of the survey include around 120,000 households

¹⁸See the separate data appendix for more detail on the construction of this series.

¹⁹The model has predictions about the flow of funds into and out of rural areas, as well as the path of consumption. After liquidity constraints are loosened, lending from outside will tend to increase in order to fund the higher consumption for those who were previously liquidity constrained. In steady state, these debts must be repaid, and so funds will tend to flow out of the community more than before. Credit has increased more quickly than deposits in rural areas during the initial expansion, but the credit to deposit ratio later stabilized (see Shah, Rao, and Shankar (2007, table 1) based on RBI banking statistics). Bank deposits have been increasing the entire time.

and 600,000 individuals and are conducted approximately every five years: in 1983, 1987-88, 1993-94, 1999-2000, and 2004-05. Smaller surveys are conducted between the major rounds, of which the only ones used here were conducted in 2000-01, 2001-02, and 2003. All of the surveys used take place evenly throughout a year of sampling, and so avoid seasonal effects.²⁰ Each NSS includes a measure of per capita consumption for each household which is widely used to compare states and calculate poverty measures. Although the rounds are internally comparable, there are sometimes significant changes in survey structure between rounds which makes comparing across rounds perilous. The debate over poverty from the 55th round in 1999-2000 is the most well known example (see Deaton and Kozel (2005)), but there were significant changes to other surveys as well. A separate data appendix details other changes in the surveys and describes how to recalculate consumption in a consistent manner.²¹

The NSS divides India into 77 regions which correspond approximately to agro-climatic zones and break the survey up into an urban and rural subsamples. I use only the rural data since these areas had a clear change in access to banking over the period. While the number of banks in urban areas certainly increased, it is not clear that this increase represents a change in access. The regions subdivide populous states, and coincide with the smaller states, and so are much more comparable in terms of size than states. Regions are also the smallest geographical identifier that can be examined across all rounds. They have maintained consistency even as several new states were formed in the late 1990s. I exclude some regions because of very small sample sizes or missing data leaving a panel of 68 regions each survey year, with an average sample size across the surveys of 770 households in rural areas per region.²²

As figure 4 makes clear, the number of bank branches in rural areas increased substantially over this period.²³ The theory suggests that any relaxation of borrowing constraints will produce

²⁰The smaller surveys used here are the 56th round in 2000-01 with 57,000 households, the 57th round in 2001-02 with 63,000 households, and the 59th round in 2003 with 42,000 households. Some of the other smaller surveys take place over only six months: the 60th from January-June 2004, the 58th from July-December 2002, and the 54th from January to June 1998. 51 through 54 are not comparable with the other surveys for other reasons as well, as is described in the separate appendix.

²¹To deal with the changes in the surveys, the appendix recalculates consumption based on individual item level data to make them comparable across rounds. Recalculating is possible for many of the surveys since there is a common portion of recall periods used through all of the surveys, although not necessarily used in the calculation of consumer expenditure by the NSSO, and in official estimates of poverty. The 55th round is still incompletely comparable, because of the problem of dual recall periods, but its exclusion does not have a large effect on the results.

²²The following regions are excluded and the reason for exclusion is given in parentheses: Nagaland (only urban in 38 and 43); Lakshadweep (small); Dadra and Nagar Haveli (small); Hills region of Assam (small); Manipur Plains (only has urban in 1983); Arunachal Pradesh (only has urban in 1983), Pondicherry (mostly urban); Delhi (mostly urban); Chandigarh (small).

²³The increase may create concern that the series is non-stationary, and so a regression with banks per capita and consumption which is also generally increasing may create spurious results (Granger and Newbold, 1974). Since rural banking in India is interesting exactly because there was a large, and explained, increase in the number of banks per capita during the period, it is not clear that this increase creates a problem of spurious regression coefficients. Not surprisingly, it is not possible to reject that the region average rural banks per capita has a unit root (using the Phillips-

the same general effect; all that is required is that the proportion of illiquid assets which can be used as collateral (α) increase. So a bank opening in a nearby town or village will have a similar, although possibly attenuated, effect as one opening locally as long as that bank is willing to lend in the area. The effect does not need to be direct: self-help groups and ROSCAs may be able to extend their activities by using a nearby bank, as may local money lenders.

4.2 Estimates

The model suggests that areas in which new bank branches open should have an immediate spike in consumption, which should then decrease to a new steady state below the previous one. The model suggests that the effects of gaining access to credit will change over time, and the estimation strategy must take the change into account. The estimating equation is then of the form:

$$\ln c_{rt} = \theta_r + \theta_t + A(L)b_{rt} + \epsilon_{rt}, \quad (4)$$

where $A(L)$ is a polynomial lag operator, b_{rt} is a measure of banking access in region r at time t , $\ln c_{rt}$ is mean log consumption, either total or non-durable, and θ_r and θ_t are region and time fixed effects. Propositions 4 and 5 show that including regional effects and time effects with log consumption will exactly deal with differences in earnings by region and growth over time under the assumption of CRRA preferences and some restrictions on the income process. Time effects also deal with any remaining problems with the National Sample Surveys that are purely mean shifting in logs, leaving the shape of the distribution unchanged. If the effect of banks is to increase the rate of return, however, or provide investment funds to credit constrained entrepreneurs, then including time effects may absorb these changes, and so I also report results without them. The model predicts that the first coefficients of $A(L)$ will be positive, with the rest negative, and the sum of coefficients should be less than zero.

Observations of consumption are limited to the years in which a comparable survey occurred, but the banking variables are available yearly, and so any lag structure can be used. Survey measurement error is not an issue since the variable measured with error appears on the left, and so will not necessarily lead to attenuation bias unless consumption itself is measured with bias. I measure access to banking as the number of branches per 1000 people at the end of the year, and so I use a one period lag to give the branches open when consumption is measured.

Perron test). This result is entirely a problem with a limited sample: the series has a natural upper and lower bound (no banks per person and one bank per person) and so cannot have a unit root. By allowing a trend and trend break between 1970 and 1990, however, it is possible to reject that the residuals are non-stationary. Simulating results as in Granger and Newbold (1974) which regresses two random walks on each other including a number of lags shows no tendency to produce a large number of results with significant first and last lag coefficients of opposite sign, suggesting that this form of spurious regression is not a problem.

Table 2: Effects of the branch expansion in India on consumption

	Log total consumption			Log non-durable consumption			Var(log total consumption)		
	[A]	[B]	[C]	[D]	[E]	[F]	[G]	[H]	[I]
Branches per 1000									
1 lag	3.559*** (0.702)	9.061*** (2.411)	3.288 (2.499)	4.763*** (0.918)	10.73*** (3.156)	6.011* (3.121)	-0.456 (0.319)	-2.917*** (1.099)	-1.727* (1.047)
3 lags		-4.924 (3.497)	3.749 (3.460)		-4.099 (4.576)	3.191 (4.321)		2.599 (1.593)	0.419 (1.449)
5 lags		0.184 (2.493)	-1.031 (2.530)		-1.363 (3.263)	-2.644 (3.160)		0.870 (1.136)	1.244 (1.060)
10 lags		0.998 (1.475)	1.994 (1.574)		2.299 (1.931)	3.426* (1.966)		-2.300*** (0.672)	-2.459*** (0.659)
15 lags		-2.994*** (0.964)	-0.565 (0.886)		-4.346*** (1.261)	-2.136* (1.106)		1.460*** (0.439)	1.540*** (0.371)
Time effects	Yes 475	Yes 475	No 475	Yes 475	Yes 475	No 475	Yes 475	Yes 475	No 475
Observations									
Regions	68	68	68	68	68	68	68	68	68
R-squared	0.560	0.581	0.485	0.405	0.432	0.364	0.104	0.139	0.106
Sum of all lags		2.325	7.436		3.217	7.848		-0.288	-0.982
p-value sum = 0		0.002	0.000		0.001	0.000		0.399	0.000
Sum L3+L5+L10+L15		-6.736	4.148		-7.509	1.837		2.629	0.745
p-value sum = 0		0.006	0.067		0.020	0.516		0.019	0.432

Notes: Standard errors in parentheses. All regressions include region effects. The banking variable is available yearly. The dependent variables are per capita household consumption from the major NSS consumer expenditure rounds in 1983, 1987-88, 1993-94, 1999-2000, and 2004-05, and the minor rounds in 2000-01, 2001-02, and 2003, consumption is updated for changing prices based on Deaton (2003) and the CPI for rural laborers). The p-value of the sum = 0 is the value from the F-stat of the test with the null that the sum of all lags equals 0. The variance of log consumption is measured within region, so the changing variance does not automatically mean there is heteroskedasticity across regions or over time.

Table 3: Effects of branch expansion on consumption—variation in the lag structure

Dependent variable: log total per capita consumption	First Lag [A]	Second Lag [B]	Sum of all lags [C]	p-value ([C]=0)	Sum of lags after first [D]=[C]-[A]	p-value ([D]=0)	Simulations	
							MSE [E]	95% Interval [F]
Lags of rural branches per 1000								
L1	3.56 (0.702)		3.56	0.000		0.000	9156	[3596,16247]
L1 L3 L5 L10 L15	9.06 (2.411)	-4.92 (3.497)	2.33	0.002	-6.736	0.006	3255	[1002,7528]
L1 L3 L5 L9 L13	8.36 (2.392)	-5.08 (3.513)	2.44	0.001	-5.911	0.016	3614	[1201,8816]
L1 L2 L4 L8 L12	31.59 (5.216)	-32.39 (5.985)	2.31	0.002	-29.284	0.000	3860	[1321,8586]
L1 L2 L3 L5 L10 L15	32.89 (5.288)	-34.99 (6.962)	1.81	0.014	-31.089	0.000	3679	[1135,8475]
L1 L3	7.71 (2.272)	-4.42 (2.302)	3.29	0.000	-4.417	0.056	4550	[1806,9270]
L1 L3 L5	7.49 (2.392)	-3.62 (3.525)	3.26	0.000	-4.225	0.078	6163	[2850,11879]
L1 L3 L5 L7	8.02 (2.381)	-5.84 (3.591)	3.22	0.000	-4.804	0.045	4178	[1797,8390]
L1 L3 L5 L7 L9	7.84 (2.400)	-5.63 (3.610)	3.07	0.000	-4.778	0.046	4061	[1406,8494]
L1 L3 L5 L7 L9 L11	8.43 (2.413)	-6.05 (3.606)	2.97	0.000	-5.455	0.024	4215	[1322,8829]
L1 L3 L5 ... L13	9.06 (2.428)	-6.23 (3.596)	2.58	0.001	-6.476	0.009	4256	[1363,9055]
L1 L3 L5 ... L15	9.40 (2.436)	-6.17 (3.591)	2.56	0.001	-6.847	0.006	4100	[1320,8908]
L1 L3 L5 ... L17	9.80 (2.467)	-6.41 (3.598)	2.49	0.002	-7.312	0.004	4190	[1366,9964]

Notes: All regression include region and time fixed effects. The lags are past values of the banking variable which is available yearly, so L7 means that a seven year lag of the banking variable is included. The dependent variable is the full per capita household consumption from the major NSS consumer expenditure rounds in 1983, 1987-88, 1993-94, 1999-2000, and 2004-05, and the minor rounds in 2000-01, 2001-02, and 2003. The p-value of the sum $\neq 0$ is the value from the F-stat of the test with the null that the sum of all lags equals 0. The p-value that the first lag \neq sum of lags is the value from the F-stat of the test with the null that the first lag equals the sum of all lags. Columns [E] and [F] are from simulations based on that lag structure for the limited sample of NSS consumption years. [E] is a measure of model fit based on simulations assuming the truth of the model; smaller is better. [E] is the mean over 200 “worlds” of 50 regions of the sum of squared deviations from the true lag parameters for regressions with the given model structure run on a limited sample of available consumption dates, as described in the text. The 95% coverage interval gives the bounds within which 95% of the worlds’ squared errors fit.

Table 2 shows the results of a parsimonious lag model for both non-durable and total (durable and non-durable) per capita household consumption. For each measure of consumption, the first column includes only one lag, the second column includes lags of banking extending out to 15 years, and the third the same lags, but without time effects. Including only one lag suggests that banks are associated on average with increased consumption by either measure. Allowing the effects to vary over time, however, shows that the dynamic effects predicted by the model are important. The initial effect is large and positive while later lags are negative, and so the total effect is smaller than the initial effect. The total sum of lags, however, is generally positive and significantly different from zero, which is different from the long-term prediction of the model. One explanation for this result may be that not enough time has passed to achieve convergence. Alternatively, banks may improve productive capacity through credit, which may raise the long-term income in the area.

Picking a lag structure is difficult in these circumstances, as was illustrated by the simulation regressions, and due to the small number of observed time periods, even though they cover a span of 23 years, standard time series methods for picking lag structure are not useful here. To deal with the difficulty of model selection, table 3 shows a number of lag variations, which suggest that there is some robustness in the path of consumption following an increase in the number of branches. All of the regressions include time effects. For each row the lags included are shown, with the first and second coefficient as well as the sum, and the sum of lags after the first. The p-values of Wald tests on the hypotheses that the sum is zero, and that the sum of lags after the first is zero, are shown in the last two columns. While the coefficients are clearly sensitive to which lags are included, the general picture of a positive and large initial effect, followed by negatives, is still supported. Moreover, the inclusion of extra lags at two year intervals shows that with more lags the initial effect increases, while the sum of later lags becomes more negative and significant, and the total effect declines. Adding more lags allows more curvature in the path to be estimated, and helps separate the long-term effect from the initial effect. The results become more supportive of the model as additional lags are added, which suggests that the path described by the theory is a good description of the data.

The simulations from the previous section can also help in model selection, even if they are sensitive underlying model parameters. I define a loss-function as the sum of squared deviations from the true path that each lag model implies. For each lag model and world (draw of 50 regions), I estimate the coefficients on the lags. These coefficients give a path of adjustment following a change in liquidity constraints where I assume that when there are gaps in the lags, the path is constant. I calculate the squared difference of the true path in each period from the estimated path, and end the comparison after 30 periods.²⁴ The last two columns of table 3 show the mean

²⁴Let the path following the adjustment be the coefficients a_i where $i = 1$ is the change immediately following the

squared deviation from 200 world draws, as well as the 95% coverage interval for a selection of different lag models using only the limited NSS subsample. Models with more lags around the greater curvature at the beginning of the path, and at least one longer lag do the best. The statistic also tends to value model parsimony, perhaps because in the limited sample imprecisely estimated lag coefficients can cause large errors. The statistic tends to favor the model (L1 L3 L5 L10 L15) which I use as the base specification, although other variations are close, and the models selected are a small subset of possible models.

The effects, whether initial or long-term, are meaningfully large. The mean number of branches per 1000 people in rural areas in 1981 was 0.052, while it was 0.088 in 1991, an increase of 69% over 10 years during the height of the branch expansion period.²⁵ The increase in branches per 1000 is approximately 0.0036 per year. Using the values for total consumption in column two of table 2, an increase of 0.0036 in the number of branches per 1000 people is associated with an immediate increase of 3.3% in consumption relative to trend, and a long-term increase of 0.84% which is statistically different from zero. The number of branches seems to do a better job of explaining changes in total consumption than consumption excluding durables. Since durables are the most likely to be bought on credit and so having banks in the area may facilitate such purchases, this result seems reasonable. The results using the two measures of consumption are very similar, however.

The model also suggests that after the first coefficient, later coefficients should be getting smaller as the community approaches its long term mean. That prediction does not appear to hold in the estimates, which move around and are often not generally precisely estimated individually. One reason may be that I am increasing the length between lags, so each lag is capturing more changes, even though the changes are smaller; the last lag, for example, captures all remaining changes in the sample.

The model has much less clear predictions about the course of inequality, although the simulations for relaxing the liquidity constraint suggest a large initial drop in inequality, followed by increases to a steady state which may be slightly higher than before. Alternatively, if some communities start accumulating after an increase in the rate of return, inequality should rise. The last columns of table 2 shows estimates of equation 4 for inequality measured as the variance of log

change in access, and $\sum_{i=1}^k a_i$ is the cumulative effect (I calculate this path to much higher precision than the reported significant figures). For each model with the coefficients on lagged proportion of banks given by $L_1, L_3 \dots$, define the implied path $v_i = \sum_{j=1}^i L_j$ where the L_j in a model which has not estimated that coefficient is zero. Then the sum of squared deviations $D^2 = \sum_{i=1}^3 0(a_i - v_i)^2$.

²⁵The mean number of branches is the unweighted regional average, which is used because it corresponds to the unweighted regressions. This translates into approximately 11,300 people per branch in 1991. To give perspective there were approximately 3300 people per branch in the US in 2003 (Critchfield et al., 2004, using a population of 290 million).

Table 4: Effects of the branch expansion in India on poverty

	Poverty		
	[A]	[B]	[C]
Branches per 1000			
1 lag	-0.122 (0.375)	-3.618*** (1.257)	1.962 (1.509)
3 lags		4.541** (1.823)	-4.151** (2.088)
5 lags		-2.101 (1.300)	-0.967 (1.527)
10 lags		-0.217 (0.769)	-1.078 (0.950)
15 lags		2.106*** (0.502)	0.324 (0.535)
Time effects	Yes	Yes	No
Observations	475	475	475
Regions	68	68	68
R-squared	0.674	0.703	0.511
Sum of all lags		0.711	-3.911
p-value sum = 0		0.069	0.000
Sum L3+L5+L10+L15		4.329	-5.873
p-value sum = 0		0.001	0.000

Notes: Standard errors are in parentheses. All regression include regional fixed effects. The banking variable is available yearly. Poverty is the headcount ratio (on a 0-1 scale), calculated using the consistent consumption data (see appendix; consumption is updated for changing prices based on Deaton (2003) and the CPI for rural laborers). The p-value of the sum = 0 is the value from the F-stat of the test with the null that the sum of all lags equals 0. The p-value that L3+L5+L10 +L15 = 0 is from the F-stat with the null that the sum of lags after the first equals 0.

per capita consumption including durables.²⁶ The long-term effect of an increase in branches per 1000 people appears to be a decline in inequality whether with or without time effects. The result is sensitive to variations in the lag structure, and to changes in the measure of consumption and access to banking. Since the variance of log consumption is quite sensitive to outliers, it is difficult to draw any strong conclusions from these results, but there is no strong evidence for inequality increasing from the presence of additional bank branches.

Banking may affect other outcomes as well. Poverty, measured as the headcount ratio, is a description of the lower tail of the consumption distribution. It is reasonable to expect that it will follow the opposite path of mean consumption: decreasing substantially when liquidity constraints are loosened (see corollary 1), then increasing to a new steady state above the previous level. Table 4 shows the results of estimating equation 4 with poverty instead of consumption.²⁷ The results do

²⁶While other inequality measures are possible, the variance of log consumption provides a close link with the log consumption measure used. Note that the inequality measured is within a region, so the changing variance does not automatically mean there is heteroskedasticity across regions or over over time.

²⁷Poverty is measured as the headcount ratio based on an the all-India rural poverty line for 1987 of 115.7 (per

mirror those for consumption: new bank branches seem to decrease poverty immediately, but later on this effect is diminished. These results are thus supportive of both the findings in Burgess and Pande (2005) and in Kochar (2005). Extending financial access can have multiple effects.²⁸

To check the robustness of the results, table 5 shows several variations of the consumption regressions. The regressions in the first four columns cluster the standard errors at the region level to allow for the variance of the error term to differ across regions. Using the sandwich estimator tends to reduce the significance of results, but leaves the basic interpretation unchanged. Including time effects may be taking much of the temporal variation from the data, however, so columns C and D replace the time effects with a time trend. The results are similar to the base specification, with an initial positive and later negatives.

One concern is that banks were not placed randomly, and so there may be selection effects. Including fixed effects, which may be correlated with other variables, deals with this endogeneity at its most important level. As shown by proposition 4, in logs bank branches placed in high income and low income places have the same effect, up to a constant. It is possible, however, that the placement is also related to future growth, and that this selection changed after the nationalization. Separate regional time trends would control for this problem as shown in 5, but the data are not quite rich enough to separate out time trends for every region, fixed effects, and dynamic effects of banking. While this form of endogeneity will potentially affect the long term, it should not affect the basic conclusion that the short term effect is much larger than the long term effect, only the estimate of the magnitude of the difference. If banks are placed in regions which will grow slower, the estimate of the long term will be biased downwards; if banks are placed in regions which grow faster than the long-term effect will tend to be biased upwards. One test for whether this bias is an issue is to allow regions with more branches before nationalization to grow at a different rate—basically allowing for a fast growing and slow growing time trend. Column D of table 5 allows an interaction with the time trend and the branches per capita in 1970. It does not seem that after including lags, those regions with more branches had significantly different growth.

While durables are less well measured than non-durables, since their use and definition changed
capita per month in Rupees). It is updated for changing prices based on Deaton (2003) and the CPI for rural laborers.

²⁸To understand whether the increase in consumption is coming from education investments I also examine the relationship between education and banking since it is sometimes asserted that credit constraints are a reason for low educational attainment. I examine two measures of education: the average number of years of schooling in the entire population, and the proportion of women age 15-30 who are literate. Both have seen very large increases since 1983, but female education, starting from a very low level, has increased substantially more than for men. Using the same lag model as for consumption, while there does seem to be some relationship, it is small. The same 0.0036 increase in branches per 1000 people used to examine consumption has a long-term effect of raising education by less than a tenth of a year using the sum of lags from the regression without time effects, and an even smaller amount with the time effects. A 0.0036 increase in branches per 1000 increases the proportion of women 15-30 who are literate by a percentage point. While not entirely negligible, this effect is only a small fraction of the increases in education that have occurred in India over the same period.

Table 5: Robustness checks

	Log total rural cons.		Log total rural cons.		Log durable rural cons.		Log total urban cons.	
	[A]	[B]	[C]	[D]	[E]	[F]	[G]	[H]
Branches per 1000	(rural branches)		(rural branches)		(rural branches)		(urban branches)	
1 lag	3.537 (2.363)	8.831** (4.126)	7.488** (3.453)	8.058** (3.306)	1.484 (1.127)	10.55** (3.754)	1.836*** (0.601)	0.122 (2.409)
3 lags		-5.137 (10.93)	0.709 (7.197)	0.452 (7.183)		-17.24*** (5.572)		0.350 (3.760)
5 lags		0.727 (5.492)	-3.340 (3.509)	-3.338 (3.598)		11.85*** (3.964)		2.784 (2.901)
10 lags		0.354 (1.215)	2.063 (1.316)	1.594 (1.265)		-5.470** (2.286)		-4.582** (1.998)
15 lags		-2.355* (1.384)	-3.492*** (1.273)	-2.856** (1.200)		0.846 (1.495)		2.859* (1.457)
Year			0.0184*** (0.00267)	0.0182*** (0.00270)				
Branches per 1000 in 1970 x Year			-0.0763 (0.0847)					
Time effects	Yes	Yes	No	No	Yes	Yes	Yes	Yes
Observations	543	543	543	543	543	543	515	515
Regions	68	68	68	68	68	68	65	65
s.e. clustered at region	Yes	Yes	Yes	Yes	No	No	No	No
R-squared	0.558	0.574	0.541	0.542	0.613	0.626	0.601	0.607
Sum of all lags		2.420093	3.424	3.909		0.541		1.533
p-value sum = 0		0.1835016	0.000	0.000		0.655		0.039
Sum L3+L5+L10+L15		-6.4113588	-4.017	-4.103		-10.014		1.412
p-value sum = 0		0.2343575	0.067	0.062		0.009		0.585

Notes: Includes regional fixed effects. Standard errors are clustered at the region level for the first four columns. The dependent variables are from the NSS and includes the 55th round. Urban banks per capita and urban consumption is measured in the same way as rural banks and consumption. Not all regions have an urban sample in all surveys. The p-value of the sum = 0 is the value from the F-stat of the test with the null that the sum of all lags equals 0. The p-value that L3+L5+L10+L15 = 0 is from the F-stat with the null that the sum of lags after the first equals 0.

over the period, and they make up a small but increasing fraction of the household budget, it is also useful to know if the same patterns hold for durable consumption by itself. Columns E and F show that the same pattern of high initial consumption and lower later consumption holds, although the lag coefficients themselves are quite volatile.

Construction of new urban branches did not generally mean an increase in access as used in the model since most urban areas already a bank branch. As shown in columns G and H, the pattern is not the same in urban areas: new branches do not seem to have an initial effect, but the long-term effect does seem to be positive.

5 Conclusion

Credit is inherently intertemporal, and this paper makes it clear that it is necessary to study its effects in an intertemporal context. Not doing so leads to an incomplete and biased view of the effects of credit. In particular, although the long-term effect of better access to credit may be a decrease in mean consumption, consumers will tend to consume substantially more initially, and so any estimate about the effects of gaining access to credit will depend on when these changes are measured. Approaches that do not deal with the dynamics will tend to produce unstable estimates that are easy to misinterpret and so should be treated with a great deal of caution.

One advantage of a complete model is that it helps gives structure to the empirics, tightly specifying what matters, what does not, and what we should be trying to identify. For example, even if banks or microfinance institutions go into high income areas, fixed effects can separate out the effects of income differences, so the endogenous choices of banks do not necessarily affect identification. Similarly, it is straightforward to deal with different growth rates of income, as long as there is enough data to take out the trend. In trying to understand the effects of financial access, it is not necessarily better instruments that matter, but the hard work of collecting more data over a long time.

The evidence from the branch banking expansion in India suggests that the dynamic effects described by the model are important. While the initial effect on consumption is large and positive, lags after the first are negative and so the total effect of the expansion falls over time. The estimates of the total effect remain positive, however, although not always significantly so. One reason for the positive long-term effect may be that not enough time has passed to move to the new stationary distribution. The simulations suggest that it is reasonable for it to take up to 30 years to approach the new stationary distribution. More hopefully, given the resources employed expanding access, perhaps credit does have some beneficial income effects by allowing new investments to take place.

There is a less optimistic interpretation of the evidence as well. As noted by Burgess and Pande (2005), the default rate for loans from rural branches was around 40% in the 1980s, during

the biggest increases in the number of rural branches. Social lending policies also targeted “priority sectors” with preferential treatment. More recently, the government of India forgave or reduced the loans of 43 million small rural borrowers, at a cost of approximately 1.6% of GDP in June 2008 (The Economist, 2008). In such an environment, loans may have become gifts. The theory still describes the path of mean consumption, since the effect initially of a large gift is nearly identical to a change in the liquidity constraint. The only difference comes in the long term: gifts are eventually consumed, and the community returns to the old distribution.

The model, although a standard one for considering intertemporal consumption, leaves out some possibly important effects that credit may have on productive activities. By leaving out some complications, I am able to prove results generally, and so establish that leaving out the dynamic effects is not an innocuous assumption. Indeed, the model suggests that it is only when the dynamic consumption effects are dealt with we are able to conclude anything about possible production effects from credit. An important but subtle point is that focusing on measuring the investments of households which gain access to credit may not help differentiate between the effects of credit on consumption and production, since investments may just be a way of storing a buffer of wealth until it is needed. The empirical results do not rule out that bank branches may have a positive effect on production or earnings, but instead suggest that the consumption shifting effect is important, and possibly dominant. While some recent work takes the challenging step of integrating the two approaches (Kaboski and Townsend, 2008, 2009), this paper makes clear that a complete re-evaluation of the effects of financial access is necessary since so much of the literature has ignored important dynamic effects.

References

- Acemoglu, Daron and Fabrizio Zilibotti. 1997. "Was Prometheus Unbound by Chance? Risk, Diversification, and Growth." *Journal of Political Economy* 105 (4):709–751.
- Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109 (3):659–84.
- Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. 2009. "The miracle of microfinance? Evidence from a randomized evaluation." Available <http://econwww.mit.edu/files/4162> (accessed 14 August 2009).
- Banerjee, Abhijit V. 2001. "Contracting Constraints, Credit Markets and Economic Development." MIT Department of Economics Working Paper 02-17 (Available at SSRN: <http://ssrn.com/abstract=306990>).
- Banerjee, Abhijit V. and Andrew F. Newman. 1993. "Occupational Choice and the Process of Development." *Journal of Political Economy* 101 (2):274–298.
- Bencivenga, Valerie R. and Bruce D. Smith. 1991. "Financial Intermediation and Endogenous Growth." *Review of Economic Studies* 58 (2):192–208.
- Burgess, Robin and Rohini Pande. 2005. "Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment." *The American Economic Review* 95 (3):780–795.
- Caballero, Ricardo J. 1991. "Earnings Uncertainty and Aggregate Wealth Accumulation." *The American Economic Review* 81 (4):859–871.
- Carroll, Christopher D. 1997. "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112 (1):1–55.
- . 2001. "A Theory of the Consumption Function, with and without Liquidity Constraints." *Journal of Economic Perspectives* 15 (3):23–45.
- Carroll, Christopher D., Karen E. Dynan, and Spencer D. Krane. 2003. "Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets." *The Review of Economics and Statistics* 85 (3):586–604.
- Carroll, Christopher D., Robert E. Hall, and Stephen P. Zeldes. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 1992 (2):61–156.
- Clarida, Richard H. 1987. "Consumption, Liquidity Constraints and Asset Accumulation in the Presence of Random Income Fluctuations." *International Economic Review* 28 (2):339–351.
- Critchfield, Tim, Tyler Davis, Lee Davison, Heather Gratton, George Hanc, and Katherine Samolyk. 2004. "The Future of Banking in America, Community Banks: Their Recent Past, Current Performance, and Future Prospects." *FDIC Banking Review* 16 (2). (<http://www.fdic.gov/bank/analytical/banking/2005jan/article1.html>).

- Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5):1221–1248.
- . 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Baltimore: Johns Hopkins University Press.
- . 2003. "Prices and Poverty in India, 1987-2000." *Economic and Political Weekly* 38 (4):362–368.
- Deaton, Angus and Valerie Kozel. 2005. "Data and Dogma: The Great Indian Poverty Debate." *The World Bank Research Observer* 20 (2):177–199.
- Deaton, Angus and Guy Laroque. 1992. "On the Behavior of Commodity Prices." *Review of Economic Studies* 59 (1):1–23.
- Deaton, Angus and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy* 102 (3):437–467.
- Dichter, Thomas. 2007. "A Second Look at Microfinance: The Sequence of Growth and Credit in Economic History." *CATO Institute Development Policy Briefing Paper* .
- Fuchs-Schündeln, Nicola. 2008. "The Response of Household Saving to the Large Shock of German Reunification." *American Economic Review* 95 (5):1798–1828.
- Galor, Oded and Joseph Zeira. 1993. "Income Distribution and Macroeconomics." *Review of Economic Studies* 60:35–52.
- Gertler, Paul, David I. Levine, and Enrico Moretti. 2009. "Do microfinance programs help families insure consumption against illness?" *Health Economics* 18 (3):257–273.
- Giné, Xavier and Robert M. Townsend. 2004. "Evaluation of financial liberalization: a general equilibrium model with constrained occupation choice." *Journal of Development Economics* 74:269–307.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. "Consumption Over the Life Cycle." *Econometrica* 70 (1):47–89.
- Granger, C.W.J. and P. Newbold. 1974. "Spurious Regressions in Econometrics." *Journal of Econometrics* 2 (2):111–120.
- Greenwood, Jeremy and Boyan Jovanovic. 1990. "Financial Development, Growth, and the Distribution of Income." *The Journal of Political Economy* 98 (5):1076–1107.
- Gross, David B. and Nicholas S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics* 117 (1):149–185.
- Hadar, Josef and William R. Russell. 1969. "Rules for Ordering Uncertain Prospects." *The American Economic Review* 59 (1):25–34.

- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *The Journal of Political Economy* 103 (2):360–399.
- Jappelli, Tullio and Marco Pagano. 1994. "Saving, Growth, and Liquidity Constraints." *The Quarterly Journal of Economics* 109 (1):83–109.
- Kaboski, Joseph P. and Robert M. Townsend. 2008. "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative." BREAD working paper 194.
- . 2009. "The Impacts of Credit on Village Economies." MIT Department of Economics Working Paper Series 09-13.
- Kazianga, Harounan and Christopher Udry. 2006. "Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso." *Journal of Development Economics* 79 (2):413 – 446.
- Khandker, Shahidur R. 2005. "Microfinance and Poverty: Evidence Using Panel Data from Bangladesh." *The World Bank Economic Review* 19 (2):263–286.
- King, Robert G. and Ross Levine. 1993. "Finance and Growth: Schumpeter might be right." *Quarterly Journal of Economics* 108 (3):717–737.
- Kiyotaki, Nobuhiro and John Moore. 1997. "Credit Cycles." *Journal of Political Economy*, 105 (2):211–248.
- Kochar, Anjini. 2005. "Social Banking and Poverty: A Micro-empirical Analysis of the Indian Experience." Stanford Center for International Development, Stanford University. Unpublished.
- Kroll, Yoram and Haim Levy. 1982. "Stochastic Dominance: A Note." *The Journal of Finance* 37 (3):871–875.
- Krusell, Per, Toshihiko Mukoyama, Aysegul Sahin, and Jr. Anthony A. Smith. 2009. "Revisiting the Welfare Effects of Eliminating Business Cycles." *Review of Economic Dynamics* 12 (3):393–402.
- Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *The Quarterly Journal of Economics* 112 (2):443–477.
- Leeladhar, V. 2008. "Consolidation in the Indian Financial Sector." RBI Bulletin (available www.rbi.org.in/scripts/BS_ViewBulletin.aspx?Id=9312, accessed 28 May 2008).
- Levine, Ross. 2005. *Handbook of Economic Growth*, vol. 1, chap. Finance and Growth: Theory and Evidence. Amsterdam: Elsevier, 865–934.
- Menon, Nidhiya. 2006. "Long-Term Benefits of Membership in Microfinance Programmes." *Journal of International Development* 18:571–594.
- Miranda, Mario J. and Paul L. Fackler. 2002. *Applied Computational Economics and Finance*. Cambridge: MIT Press.

- Morduch, Jonathan. 1998. "Does Microfinance Really Help the Poor? New Evidence from Flagship Programs in Bangladesh." Princeton University Research Program in Development Studies Working Paper 211.
- Panagariya, Arvind. 2008. *India: The Emerging Giant*. Oxford: Oxford University Press.
- Pitt, Mark M. and Shahidur R. Khandker. 1998. "The Impact of Group Based Credit Programs on Poor Households in Bangladesh: Does the Gender of Participants Matter?" *Journal of Political Economy* 106 (5):958–996.
- Rabault, Guillaume. 2002. "When do borrowing constraints bind? Some new results on the income fluctuation problem." *Journal of Economic Dynamics and Control* 26:217–245.
- Roodman, David and Jonathan Morduch. 2009. "The Impact of Microcredit on the Poor in Bangladesh: Revisiting the Evidence." Center for Global Development Working Paper 174.
- Rosenzweig, Mark R. and Kenneth I. Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *The Journal of Political Economy* 101 (2):223–244.
- Schechtman, Jack and Vera L. S. Escudero. 1977. "Some results on an 'Income Fluctuation Problem'." *Journal of Economic Theory* 16:151–166.
- Shah, Mihir, Rangu Rao, and P. S. Vijay Shankar. 2007. "Rural Credit in 20th Century India: Overview of History and Perspectives." *Economic and Political Weekly* :1351–1364.
- Srinivas, M. N. 1976. *The Remembered Village*. Delhi: Oxford University Press.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and risk sharing over the life cycle." *Journal of Monetary Economics* 51 (3):609 – 633.
- Tarozzi, Alessandro. 2007. "Calculating Comparable Statistics From Incomparable Surveys, With an Application to Poverty in India." *Journal of Business & Economic Statistics* 25:314–336.
- The Economist. 2008. "Waiving, not drowning: India writes off farm loans. Has it also written off the rural credit culture?" *The Economist* 3 July.
- Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* 62 (3):539–591.
- Townsend, Robert M. and Kenichi Ueda. 2006. "Financial Deepening, Inequality and Growth: A Model-Based Quantitative Evaluation." *Review of Economic Studies* 73:251–293.

A Mathematical Appendix

The following definitions and propositions are variations from Rabault (2002), who expands slightly on the original Schechtman and Escudero (1977) results. They are used in the proofs that follow.

Definition 1. *An admissible policy for a given initial condition w_0 is a sequence $\{c_t\}$ which satisfies the accumulation equation and budget constraints.*

Generally we are interested in a policy function which creates such an admissible sequence $c_t = c(w_t, \alpha)$. Since the expected discounted sum of period utilities will not necessarily converge, a more general optimality condition is used:

Definition 2. *For an initial w_0 and admissible policies $\{c_t\}$ and $\{\tilde{c}_t\}$, $\{c_t\}$ overtakes $\{\tilde{c}_t\}$ if there exists a finite S_0 such that $E_0 \left[\sum_{t=0}^S \beta^t (u(c_t) - u(\tilde{c}_t)) \right] > 0$ for all $S > S_0$. For an initial condition w_0 , an admissible policy is optimal if it overtakes any other admissible policy.*

In addition to the model, three additional assumptions ensure that the asset accumulation process will not explode so there can be a stationary distribution in which consumers will occasionally exhaust their resources:

Assumption 1. *The rate of return is not large enough to overcome impatience: $\beta(1+r) < 1$.*

Assumption 2. *The coefficient of absolute risk aversion tends to zero as consumption goes to infinity: $\lim_{c \rightarrow \infty} u''(c)/u'(c) = 0$.*

Assumption 3. *The lower bound of the non-borrowable income process is greater than 0: $y_l > 0$.*

Then under these assumptions Rabault (2002) shows the following proposition which is a special case of a more general result that allows for potentially unbounded marginal utility. Similar results in Schechtman and Escudero (1977) and Clarida (1987) require slightly more restrictive assumptions for utility.

Proposition 6 (Rabault (2002)). *Under the model and assumptions 1, 2, and 3, the Markov process for disposable wealth w_t induced by the optimal policy has a unique invariant distribution, and the borrowing constraints bind infinitely often along the optimal path.*

Since the income process is bounded away from zero, not being able to borrow ($\alpha = 0$) does not affect the existence of a stationary distribution.

Corollary 2. *The existence of a stationary distribution does not depend on the value of α .*

Following Deaton (1991), define the operator T which takes a function $\phi_i : W \times A \rightarrow \mathbb{R}$ to a function ϕ_{i+1} where $W = [y_l - h, \infty)$ and $A = [0, 1]$ as $\phi_{i+1}(w_t, \alpha) = \lambda(y_l)$ if $w_t < y_l - \alpha h$ and otherwise:

$$\phi_{i+1}(w_t, \alpha) = \max\{\lambda(w_t + \alpha h), (1+r)\beta E_t[\phi_i((1+r)(w_t - \lambda^{-1} \circ \phi_{i+1}(w_t, \alpha)) + rh + y_{t+1}, \alpha)]\} \quad (5)$$

where λ is marginal utility. If T has a fixed point $T\phi = \phi$, then the associated policy is $c(w_t, \alpha) = \lambda^{-1} \circ (\phi(w_t, \alpha))$. Then combining the results of Deaton and Laroque (1992) and Schechtman and Escudero (1977):

Proposition 7. *Under assumptions 1 and 3 for a given α , T has a unique fixed point in the space of non-negative continuous non-increasing functions of w , and the associated policy is optimal.*

Having defined the mapping T , the following two lemmas establish that the fixed point is non-increasing as a function of α , a result that I use in proving proposition 2.

Lemma 1. *T takes continuous non-increasing functions of α to continuous non-increasing functions.*

Proof. Let $G(q, w, \alpha) = (1+r)\beta E_t[\phi_{i+1}((1+r)(w_t - \lambda^{-1}(q)) + rh + y_{t+1}, \alpha)]$, and note that if ϕ_{i+1} is continuous and non-increasing then G is continuous and non-increasing in each of its arguments (since λ^{-1} is strictly decreasing, $-\lambda^{-1}(q)$ is an increasing function of q). For a given (w, α) , $\phi_{i+1} = \lambda(y_l)$ if $w_t < y_l - \alpha h$, and is trivially continuous and non-increasing. $\lambda(w_t + \alpha h)$ is continuous and strictly decreasing in both arguments. Otherwise, ϕ_{i+1} is defined as the solution q of $G(q, w, \alpha) - q = 0$. Using the mean value theorem, if G is continuous then ϕ_{i+1} is continuous. If $G(q^0, w, \alpha^0) - q^0 = 0$, for $\alpha^1 > \alpha^0$, $G(q^0, w, \alpha^1) - q^0 \leq 0$, since G is non-increasing. Then the solution q^1 to $G(q^1, w, \alpha^1) - q^1 = 0$ is $q^1 \leq q^0$ since $G(q, w, \alpha) - q$ is strictly decreasing in q . \square

Lemma 2. *The fixed point $T\phi = \phi$ is a continuous and non-increasing function of α .*

Proof. By standard arguments T is a contraction on the space of continuous bounded functions $C(W \times A)$ with a supremum norm. Since T maps non-increasing functions to non-increasing functions, which are a closed sub-set of $C(W \times A)$, ϕ is non-increasing and continuous in α . \square

It is sometimes useful to rewrite the the consumer's problem so that $\rho = r$ which reduces the number of parameters, which can be done by redefining h .

Lemma 3. *The consumer's problem with illiquid income ρh and borrowing constraint $B = (\rho/r)h$ is isomorphic to a problem with $\rho = r$ and $B = \tilde{h}$ for some \tilde{h} .*

Proof. Let $\tilde{h} = (\rho/r)h$. Then illiquid income is $r\tilde{h}$, and $B = \tilde{h}$, and the problem is unchanged. \square

Finally, there are different results in the literature which make different assumptions to bound marginal utility. The following lemma establishes that the different approaches are equivalent.

Lemma 4. *Let $c(w)$ be the optimal policy in the consumer's maximization problem when $y_l > 0$ and $u'(y_l) < \infty$. Then in the alternate problem with period utility $\tilde{u}(c) = u(c + y_l)$, so that $\tilde{u}'(0) < \infty$, income process $\tilde{y}_t = y_t - y_l \geq 0$, and the additional constraint that $\tilde{c} \geq 0$, the optimal policy is $\tilde{c}(w) = c(w) - y_l$.*

Proof. Without loss of generality, take $h = 0$ (a change of variables setting $z_t = w_t + \alpha h$ and $\hat{y}_t = y_t + r(1 - \alpha)h$ transforms the problem). If the policy $c(w)$ meets the first order conditions for the original problem, then $\tilde{c}(w)$ meets the first order conditions for the modified problem without the additional constraint that $\tilde{c} \geq 0$. To show this constraint does not modify the solution, it is necessary to show that in the original problem $c(w) \geq y_l$. Suppose for some w_t , $c(w_t) < y_l$. We know $w_t \geq y_l$, since $c \leq w$, and this implies $w_{t+1} = (1+r)(w_t - c) + y_{t+1} \geq (1+r)(w_t - c) + y_l > (1+r)(w_t - y_l) + y_l = w_t - r(w_t - y_l) \geq w_t$. Under the proposed policy, $w_{t+1} > w_t$ in all possible realizations of y_{t+1} . The first order conditions require that $u'(c(w_t)) = \beta E_t[u'(c(w_{t+1}))]$ since by the supposition $c(w_t) < w_t$. But this implies that $c(\cdot)$ is decreasing for some range of w since $w_{t+1} > w_t$, which violates that the optimal policy is non-decreasing by Theorem 3.5 in Schechtman and Escudero (1977). \square

Corollary 3. *The results of Clarida (1987) with $u'(0) < \infty$ and $y_l \geq 0$ and $c \geq 0$ apply to the model with $y_l > 0$ and $u'(0) = \infty$.*

A.1 Proof of proposition 1

Proof. By lemma 4, I can use results based on the alternate assumption that $u(\cdot)$ is bounded. Then the operator in the Bellman equation is a contraction, and the value function exists and is unique (Clarida, 1987). Let $c(\cdot, \alpha^1)$ be the optimal (overtaking) policy for $\alpha^1 < 1$ and $c(\cdot, \alpha^0)$ the optimal policy for α^0 with $\alpha^1 > \alpha^0$. By assumptions 1 and 3 the limiting policy is the optimal policy (Schechtman and Escudero, 1977), and so by definition of overtaking optimal $V(w, \alpha^1) > \tilde{V}(w, \alpha^1)$ where $\tilde{V}(\cdot, \alpha^1)$ is total utility obtained from following some non-optimal admissible policy (these comparisons are meaningful since V is bounded since the period return is bounded). $c(\cdot, \alpha^0)$ is an admissible policy in the problem with α^1 . Then $V(w, \alpha^1) > V(w, \alpha^0)$, unless the optimal policy for α^0 is also the optimal policy for α^1 . That is, since $c(\cdot, \alpha^0)$ is optimal, any other admissible policy is worse unless it is the same policy. To show they are different it is necessary to show that relaxing the borrowing constraint changes the optimal policy with α^0 which occurs if the policy for α^1 is not admissible in the problem with α^0 . But with assumption 2, the borrowing constraint binds in finite time almost surely. Then for some w reached by the policy with α^1 , $c(w, \alpha^1) = w + B\alpha^1 > w + B\alpha^0$ and so is inadmissible in the problem with α^0 . \square

A.2 Proof of proposition 2

Proof. By lemmas 1 and 2, $\phi(w, \alpha)$ the fixed point of the operator T is continuous and non-increasing in α . To show that $c(w, \alpha) = \lambda^{-1} \circ (\phi(w_t, \alpha))$ is strictly increasing in α when $w \geq y_l - \alpha^0 h$, it is necessary to show that $\phi(w_t, \alpha)$ is strictly decreasing in α whenever $w_t \geq y_l - \alpha h$. For a given α^0 and $w_t \geq y_l - \alpha^0 h$, take $\alpha^1 > \alpha^0$. Then by definition of the fixed point:

$$\phi(w_t, \alpha) = \max\{\lambda(w_t + \alpha h), (1+r)\beta E_t[\phi((1+r)(w_t - \lambda^{-1} \circ \phi(w_t, \alpha)) + rh + y_{t+1}, \alpha)]\}. \quad (6)$$

By proposition 6 for any $\alpha \in A$ and initial w , the process $w_{t+1} = (1+r)(w_t - \lambda^{-1} \circ \phi(w_t, \alpha)) + rh + y_{t+1}$ is a renewal process, and the borrowing constraint will bind in finite time almost surely. Then for any initial w_t there exists some finite number of periods S such that after S periods the support at time t of w_{t+S} puts positive measure that $\lambda(w_{t+S} + \alpha h) > (1+r)\beta E_{t+S}[\phi((1+r)(w_{t+S} - \lambda^{-1} \circ \phi(w_{t+S}, \alpha)) + rh + y_{t+S+1}, \alpha)]$. For w_{t+S} for which the liquidity constraint binds, $\phi(w_{t+S}, \alpha^0) > \phi(w_{t+S}, \alpha^1)$. Then at $t+S-1$, for any w_{t+S-1} for which the borrowing constraint binds with positive probability at $t+S$ or binds at $t+S-1$, $\phi(w_{t+S-1}, \alpha^0) > \phi(w_{t+S-1}, \alpha^1)$ since ϕ is non-increasing for any w for which the borrowing does not bind this period or with positive probability next period, and strictly decreasing otherwise. The same reasoning applies to each of the S periods, and so $\phi(w_t, \alpha^0) > \phi(w_t, \alpha^1)$. \square

A.3 Proof of proposition 3

Proof. Without loss of generality, let $\alpha^1 = 1$ and $\alpha^0 = 0$ and $r = \rho$ by lemma 3. The steady state distribution must satisfy:

$$F(w', \alpha) = \int G(w' - (1+r)(w - c(w, \alpha)) - rh) dF(w, \alpha).$$

Define the mapping T_α as:

$$T_\alpha \tilde{F}(w') = \int G(w' - (1+r)(w - c(w, \alpha)) - rh) d\tilde{F}(w).$$

So T_α maps the distribution of wealth \tilde{F} which obtains in a given period to the one which obtains in the next period.

Start with the steady state distribution for $\alpha = 1$, which by definition satisfies $F(w, 1) = T_1 F(w, 1)$, and apply the mapping T_0 . By proposition 2, $c(w, 1) > c(w, 0)$. Since G is a non-decreasing function, and $w - c(w, 1) < w - c(w, 0)$:

$$T_0 F(w, 1) \leq T_1 F(w, 1) = F(w, 1)$$

and the inequality is strict for some w , since G must be increasing over some interval.

The amount saved $w - c(w, 1)$ is a non-decreasing function by Schechtman and Escudero (1977, Theorems 3.5 and 3.2 and property 1.3). This means that the kernel $\kappa(w', w) = G(w' - (1 + r)(w - c(w, 1)) - rh)$ is a non-increasing function of w . It is not, however, necessarily differentiable in w since $c(w, \alpha)$ need not be differentiable in w . Applying the results of Kroll and Levy (1982), which extends the first order stochastic dominance results of Hadar and Russell (1969) to non-differentiable functions, we have that $T_0^2 F(w, 1) \leq T_0 T_1 F(w, 1) = T_0 F(w, 1)$.²⁹

By the same reasoning $T_0^n F(w, 1) \leq T_0^{n-1} F(w, 1)$ and each step is stochastically dominated by $F(w, 1)$. In the limit:

$$F(w, 0) = \lim_{n \rightarrow \infty} T_0^n F(w, 1) \leq F(w, 1)$$

where the inequality holds strictly for some w from the first step.

By stochastic dominance $\bar{w}_0 > \bar{w}_1$ where $\bar{w}_\alpha = \int w dF(w, \alpha)$. In equilibrium:

$$\bar{w}_0 = (1 + r)(\bar{w}_0 - \bar{c}_0) + \bar{y} + rh$$

$$\bar{w}_1 = (1 + r)(\bar{w}_1 - \bar{c}_1) + \bar{y} + rh.$$

Solving for \bar{c}_0 and \bar{c}_1 then implies $\bar{c}_0 > \bar{c}_1$. □

A.4 Proof of proposition 4

Proof. Without loss of generality, consider subpopulation j and its relationship with the “root” subpopulation with $\mu = 1$, and denote its income distribution G^1 and illiquid wealth h^1 . Take the functional operator T^j defined as in equation 5 which maps ϕ_i^j to ϕ_{i+1}^j (here, as in equation 5, i is an index of functions, not of subpopulations). Then $T^j \phi^j = \phi^j$ is the unique fixed point of T^j . Using the definition of the fixed point ϕ^1 gives:

$$\begin{aligned} (\mu^j)^{-\gamma} \phi^1(w_t / \mu^j) &= \max\{(\mu^j)^{-\gamma} (w_t / \mu^j + \alpha h^1)^{-\gamma}, \beta(1 + r) \\ &\quad E_t[(\mu^j)^{-\gamma} \phi^1((1 + r)(w_t / \mu^j - (\phi^1(w_t / \mu^j))^{-1/\gamma}) + rh^1 + y_{t+1}^1)]\} \\ &= \max\{((w_t + \alpha h^j)^{-\gamma}, \beta(1 + r) \\ &\quad E_t[(\mu^j)^{-\gamma} \phi^1(((1 + r)(w_t - ((\mu^j)^{-\gamma} \phi^1(w_t / \mu^j))^{-1/\gamma}) + rh^j + y_{t+1}^j) / \mu^j)]\} \end{aligned}$$

which implies that $(\mu^j)^{-\gamma} \phi^1(w_t / \mu^j)$ is a fixed point of T^j . Since the fixed point is unique $\phi^j(w_t) = (\mu^j)^{-\gamma} \phi^1(w_t / \mu^j)$.

²⁹For differentiable functions the argument is simpler, and can be shown directly. Integrating by parts and combining gives that: $\int_{-\infty}^{\infty} \kappa(w', w) dT_0 F(w, 1) - \int_{-\infty}^{\infty} \kappa(w', w) dT_1 F(w, 1) = \kappa(w', w)[T_0 F(w, 1) - T_1 F(w, 1)]_{-\infty}^{\infty} - \int_{-\infty}^{\infty} \partial \kappa(w', w) / \partial w [T_0 F(w, 1) - T_1 F(w, 1)] dw \geq 0$ since the first part is zero since $T_0 F(w, 1)$ and $T_1 F(w, 1)$ are distributions, and the inequality comes from the combinations of $\partial \kappa(z', z) / \partial z \leq 0$ and $[T_0 F(w, 1) - T_1 F(w, 1)] \leq 0$.

By proposition 6 there is a unique steady state distribution of wealth for each subpopulation. This distribution must satisfy:

$$F^j(w') = \int G^j(w' - (1+r)(w - c_j(w)) - rh_j) dF^j(w)$$

where $c_j(w) = (\phi^j(w))^{-1/\gamma}$. Then transforming into G^1 we have:

$$\begin{aligned} F^j(w') &= \int G^j(\mu_j(w'/\mu_j - (1+r)(w/\mu_j - (\mu_j^\gamma \phi^j(w))^{-1/\gamma}) - rh_j/\mu_j)) dF^j(w) \\ &= \int G^1(w'/\mu_j - (1+r)(w/\mu_j - (\mu_j^\gamma \phi^j(w))^{-1/\gamma}) - rh_j/\mu_j) dF^j(w) \quad (7) \\ &= \int G^1(w'/\mu_j - (1+r)(w/\mu_j - (\phi^1(w/\mu_j))^{-1/\gamma}) - rh_1) dF^j(w), \end{aligned}$$

which is the stationary transformation for $F^1(w/\mu_j)$. So in steady state $F^1(w/\mu_j) = F^j(w)$, and so the distribution of w_t^j is multiplicatively shifted by μ_j .

By the same argument, the same holds true for the stationary distribution after a change in the liquidity constraints since it is possible to redefine the variables to attain a zero liquidity constraint without changing the problem. Moreover, the same relationship holds at all points in the transition path. From equation 7, if $F_t^j(w) = F_t^1(w/\mu_j)$, then so will F_{t+1}^j , since the transformation from one period to another preserves the relationship. So starting from the stationary distribution in which the relationship holds, it will hold along the transition path as well.

If $F_t^j(w) = F_t^1(w/\mu_j)$ at time t , then the distribution of consumption has the same relationship. For any liquidity constraint, $c_j(w) = (\phi^j(w))^{-1/\gamma} = (\mu_j^{-\gamma} \phi^1(w/\mu_j))^{-1/\gamma} = \mu_j c_1(w/\mu_j)$. Then the relationship for the distribution of consumption for subpopulation j is given by: $J_t^j(c) = F_t^j(c_j^{-1}(c)) = J_t^1(c/\mu_j)$.

□

A.5 Proof of proposition 5

Proof. The proof follows a standard argument, and so I merely sketch it. A more complete version is available in the working paper version of this paper. The consumer's problem is nearly identical with the difference that h is replaced by h_t . (Use lemma 3 to rewrite so $\rho = r$ if necessary.) If μ_t is not stationary, it no longer makes sense to try to examine stationary distributions of consumption and wealth. With CRRA utility the ratio of consumption and wealth to μ_t is still stationary (under the assumptions of proposition). Written in ratios, the problem meets all the requirements for stationarity in proposition 6 and so follows the path in propositions 2 and 3. □