Economic Quarterly—Volume 94, Number 1—Winter 2008—Pages 17–44

Credit Access, Labor Supply, and Consumer Welfare

Kartik B. Athreya

Recent work has argued that U.S. households have seen a systematic improvement in their ability to borrow against future labor income.[1] First, Narajabad (2007) points out that the "extensive" margin of credit has changed; he calculates that in 1989, 56 percent of households held a credit card, while 29 percent were actively "revolving" debt (i.e., keeping positive balances after the most recent payment to lenders). By 2004, these measures had risen to 72 and 40 percent, respectively. The availability of such credit has been accompanied by its use, suggesting that households are genuinely less constrained at present than they were in the past. Using Survey of Consumer Finances (SCF) data, Narajabad (2007) shows that average debts among those paying interest on credit card debts nearly doubled from 1989 to 2004, jumping from roughly $1,800 per cardholder to $3,300 (in 1989 dollars). When aggregated, these changes are reflected in the striking findings of Krueger and Perri (2006), who show that the ratio of unsecured debt to disposable income *quadrupled* from 2 to 9 percent over the period 1980–2001. Parker (2000) and Iacoviello (forthcoming) provide further details on the increase in household indebtedness. Lastly, and most sensationally, recent events in mortgage markets also suggest that there has been a sharp expansion in credit availability. Notably, both the rapid growth of the aggregate homeownership rate in the late 1990s and the recently high default rates on some types of mortgages suggest that the ability to take highly "leveraged" positions in residential real estate has indeed increased.

I thank Ahmet Akyol and especially Juan Carlos Hatchondo for discussions. I also thank Kevin Bryan, Chris Herrington, and Yash Mehra for comments. I thank Kay Haynes for expert editorial help. The views expressed in this article are those of the author and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. All errors are my own.

[1] Edelberg (2006), Furletti (2003), Krueger and Perri (2006), and Narajabad (2007).

The large changes in borrowing summarized above appear to be consistent with improved information held by lenders at the time of credit extension (see, for example, Athreya, Tam, and Young 2007), as well as a secular decline in the cost of maintaining and issuing credit contracts (see, for example, Athreya 2004). As an empirical matter, Furletti demonstrates strikingly that in 2002, the interest rate conferred on those with the highest credit score was eight percentage-points lower than those with the lowest credit scores. In 1990, by comparison, this premium was essentially nonexistent. Relatedly, Edelberg (2006) notes that there has been a substantial increase in the sensitivity of most loan interest rates to forecasts of default risk. Improvements in the ability of lenders to screen borrowers will have allowed many to access credit, instead of being denied outright. In sum, both theory and evidence strongly suggest that households may now be better able than ever before to use credit markets to smooth consumption.

A direct consequence of better access to credit is allowing households to borrow to finance consumption. However, a perhaps equally important effect, and one that has not received systematic attention thus far, is that better credit access will allow households to more effectively align work effort with productivity. That is, when temporarily unproductive, a household can use credit in lieu of labor effort, and instead work more when it is relatively productive. At a quantitative level, varying labor effort in response to productivity may well be an important channel for consumption smoothing; it has also long been known that idiosyncratic shocks to labor productivity dwarf business cycle-related risks facing U.S. households. It is also agreed that these shocks are, in general, poorly insured.[2]

The use of labor effort itself as a smoothing device, even in the absence of credit markets, has only recently received serious quantitative attention. This line of research includes Pijoan-Mas (2006), Marcet, Obiols Homs, and Weill (2007), Flodén (2006), Flodén and Lindé (2001), Li and Sarte (2006), and Chang and Kim (2005, 2006). Taken as a whole, the preceding body of work suggests that variable labor supply may be an important mechanism by which households maintain smooth paths of consumption. However, aside from the bankruptcy model of Li and Sarte (2006), none of the preceding directly assesses the extent to which changes in credit access will alter labor supply behavior, savings, and consumption. The purpose of this article is to provide some simple experiments aimed at uncovering the interaction between credit markets and labor markets in the presence of idiosyncratic and uninsurable productivity risk.

[2] Storesletten, Telmer, and Yaron (2004) is an important landmark in this literature. The interested reader should also consult the *Review of Economic Dynamics* (2000) interview with Kjetil Storesletten.

I augment the model of household consumption and work effort described in Pijoan-Mas (2006). The latter is a standard model of uninsurable idiosyncratic risk that is augmented to allow for flexible labor supply, but one in which borrowing is prohibited.[3] I ask four specific questions. First, in the presence of flexible labor supply, how do changes in borrowing constraints influence aggregate precautionary savings and the size of the economy? Second, how do changes in borrowing constraints alter the efficiency of the labor input? Third, how do changes in borrowing capacity alter "who" works? Fourth, what are the welfare implications of improvements in credit access, and how are these welfare effects distributed across households?

Why is it useful to address these questions? With respect to the first question, recent work of Marcet, Obiols Homs, and Weill (2007) contains an important insight about precautionary savings in the presence of flexible labor supply. Namely, they point out that at the household level, the ex post effect of increased precautionary savings will be to *reduce* the labor supply. Intuitively, if most households are, on average, wealthier due to the maintenance of a larger stock of wealth, then they may also choose to work less. In turn, aggregate savings may not rise, and can even fall, relative to an economy in which households do not face uninsurable idiosyncratic risk. As a result, a key link between uninsurable risk and the "size" of the economy is broken. Specifically, with inelastic labor, Huggett and Ospina (2001) proved that the economy must be larger in the presence of uninsurable idiosyncratic risk than in its absence.

The second question, on the efficiency of labor supply, is motivated by the observation that when borrowing is possible, a wealth-poor household facing temporarily low productivity may choose to take leisure and instead *borrow* to smooth consumption. Conversely, when borrowing is ruled out, labor supply may be far less sensitive to current productivity. This implication of credit constraints has attracted the attention of development-related research. Recent work of Jayachandran (2007) suggests that in rural India, borrowing limits indeed create nontrivial welfare losses. Similarly, Malapit et al. (2006), and Garcia-Escribano (2003) argue that variations in family labor supply are important for consumption smoothing, especially when households have low asset holdings. In settings in which borrowing is prohibited, Pijoan-Mas (2006) and Flodén and Lindé (2001) both find that the correlation of hours and productivity is near zero, while the ratio of effective hours to labor hours is close to the average productivity of households. If borrowing were possible, both the correlation between effort and productivity, as well as the ratio of

[3] Pijoan-Mas (2006) does study allocations under more generous borrowing limits, but *recalibrates* the model to generate the observed correlations between effort and productivity. This is because he treats borrowing constraints as unobservable. The key point is that the recalibrated elasticity of substitution of labor turns out to be substantially different than in the benchmark. This suggests precisely that borrowing limits are likely to be important in influencing behavior.

"effective" hours to labor hours would likely rise, as households would supply labor primarily when productive.

The third question of "who works hard, and when?" follows naturally from the observation that changes in borrowing constraints will affect households differentially. For example, wealthy households may be fairly insensitive to credit access. Conversely, those who are not as rich but have high current productivity may wish to borrow and work hard. In the absence of credit, however, these households may work fewer hours, as they are unable to offset declines in current leisure with increases in current consumption. The preceding are only two examples of the outcomes that might ensue from changes in credit access. Moreover, at an aggregate level, the behavior of households in the economy will then depend on, and in turn, determine the overall long-run joint distribution of wealth and productivity. Therefore, an emphasis of the present work is to document how changes to credit access alter both the characteristics of worker behavior and the equilibrium joint distribution of wealth, productivity, and effort.[4]

Lastly, the results in this article are useful for organizing one's views on the desirability of increased access to credit. Notably, the model suggests that when credit availability is relatively lax, some households *will* borrow a great deal, and if unlucky in terms of their productivity, will choose to work very hard as a result. However, the model also suggests that ex ante, households prefer the ability to reach high debt levels. Policies that effectively limit the availability of credit may leave borrowers as a class worse off in the long run. The results, therefore, suggest caution in using poor ex post outcomes to decide on the usefulness of an increased ability to borrow. This message is particularly relevant given recent public debate on the desirability of debt relief and mandatory mortgage renegotiation.

The main results are as follows. First, the hardest working households are those who are least wealthy, and most strikingly, also the least productive. Second, credit access can play an important role in reducing high labor effort by low-productivity households. Third, the buffer-stock tendencies of households imply that the distance from the borrowing constraint is often more important than the actual level of wealth in influencing labor effort. Fourth, measures of welfare gains to current consumers show that there are significant benefits from

---

[4] One question that is relevant, but not addressed here, is the extent to which measures of labor supply elasticity are biased by ignoring borrowing constraints when, in fact, they are present. This is valuable for ensuring that models of the type studied here deliver accurate implications when used for policy analysis (see, for example, Domeij and Flodén [2006]). Accurately measuring labor supply elasticities are key for business-cycle related research, as well. A cornerstone of standard models of aggregate economic activity, such as the basic real business cycle model (for example, KPR 88), are the consumers who value consumption and leisure and face productivity shocks. A key parameter governing the behavior of such models is the elasticity of labor supply, which directly dictates the extent to which households, and in turn aggregates, respond to changes in labor productivity.

expansions in credit access and that these gains accrue disproportionately to the relatively poor and relatively rich. The remainder of the article is organized as follows. Section 1 describes the model and equilibrium concept, which closely follows the environment of Pijoan-Mas (2006) and Flodén and Lindé (2001). Section 2 then assigns parameters, and Section 3 presents results. In Section 4, I compute and discuss two measures of consumer welfare gains from relaxing credit limits, and Section 5 contains conclusions and suggestions for future work.

1. MODEL

The model contains three important features. First, households in the model face uninsurable, but purely idiosyncratic productivity risk. Second, households have access to only a single risk-free, noncontingent bond that may be accumulated or sold short. Third, households can vary their labor supply.

Preferences

There is a continuum of ex ante identical, infinitely lived households whose preferences are defined over random sequences of consumption and leisure. The size of the population is normalized to unity, there is no aggregate uncertainty, and time is discrete. Preferences are additively separable across consumption in different periods. Let β denote the time discount rate. Therefore, each agent solves

$$\max_{\{c_t^i, l_t^i\}_{t=0}^{\infty}} E_{\circ} \sum_{t=0}^{\infty} \beta^t u(c_t^i, l_t^i), \tag{1}$$

subject to a budget constraint explained below.

Endowments

Each household is endowed with one unit of time, which it supplies a portion of as labor and uses the remainder for leisure. At the beginning of each period, households receive a cross-sectionally independent productivity shock z_t^i, which leaves them with productivity *level* $q_t^i \equiv e^{z_t^i}$. A useful interpretation of the shocks to productivity is that they are elements of a list of factors that alter the ability of households to convert labor effort into consumption goods. Examples include the health status of workers and even local variations in business conditions. What is precluded from this list are factors that lower the productivity of all workers simultaneously, such as a sharp increase in real prices of inputs such as crude oil.

Market Arrangement

There is a single, competitive, asset market in which agents may trade a one-period-lived, risk-free claim to consumption. The net supply of these claims is interpreted as the aggregate capital stock. Households enter each period with asset holdings a_t^i and face returns on capital and labor of r_t and w_t, respectively. Gross-of-interest asset holdings are, therefore, given by $(1+r_t)a_t^i$. Let private period-t consumption and savings be given as c_t^i and a_{t+1}^i, respectively.

Given that labor supply is endogenous, it is useful to think of the individual household's problem as one in which it first "sells" its entire labor endowment, which yields a labor income of $w_t q_t^i$ and then "buys" leisure l_t^i at its opportunity cost wq_t^i. The household's budget constraint is then given as follows:

$$c_t^i + wq_t^i l_t^i = w_t q_t^i + (1+r_t)a_t^i. \tag{2}$$

$$l_t^i \in [0, 1].$$

Stationary (Constant Prices) Recursive Household Problem

Under constant prices, whereby $r_t = r$ and $w_t = w$, the household's problem is recursive in two state variables, a and z. Suppressing the household index i and time subscripts t in order to avoid clutter, the stationary recursive formulation of the household's problem is as follows:

$$v(a, z) = \max_{c,l,a'} \left[u(c, l) + E(v(a', z')|z)\right] \tag{3}$$

subject to

$$c + wql + a' \leq wq + (1+r)a, \tag{4}$$

where

$$a' \geq \underline{a}. \tag{5}$$

Firms

There is a continuum of firms that take constant factor prices as given and use Cobb-Douglas production. In a stationary equilibrium, the aggregate capital stock K and the aggregate labor supply measured in productivity units L will be constant. Let the stationary joint distribution of assets and labor productivity be denoted by μ. The aggregate effective labor input is then given as:

$$L = \int q(z)(1 - l(a, z))d\mu.$$

By contrast, aggregate *hours* worked are given as:

$$H = \int (1 - l(a, z))d\mu.$$

Notice that in general, L and H will differ, precisely because hours worked and productivity may move together when labor is elastically supplied. As stated at the outset, a measure of the efficiency of labor supply will be the deviation of the ratio $\frac{L}{H}$ from the mean of log productivity, which is set to unity. Denoting the stationary marginal distribution of household assets by $\mu_a \equiv \int_z d\mu(a, z)$, aggregate savings is given by

$$K = \int a d\mu_a.$$

Aggregate output then arises from a Cobb-Douglas production function that combines effective hours and capital:

$$Y = F(K, L). \tag{6}$$

Finally, denote the depreciation rate by δ, and current aggregate consumption by C. This implies that the economy-wide law of motion for the capital stock is given by

$$K' = (1 - \delta)K + F(K, L) - C. \tag{7}$$

I will restrict attention to stationary equilibria where aggregate capital, output, and consumption are all constant, which then implies that

$$\begin{aligned} K' &= K, \ and \\ C &= F(K, L) - \delta K. \end{aligned} \tag{8}$$

Equilibrium

A stationary recursive competitive general equilibrium for this economy, given parameters, is a collection of (i) a constant interest rate, r and wage rate, w; (ii) decision rules for the household, $a' = g_a^*(a, z|w, r)$, $l = g_l^*(a, z|w, r)$; (iii) aggregate/per-capita demand for capital and effective labor by firms, $K^*(w, r)$, and $L^*(w, r)$, respectively; (iv) supply of capital and effective labor by households, $K(w, r)$ and $L(w, r)$, respectively; (v) a transition function $P(a, z, a', z')$ induced by z and the optimal decision rules; and (vi) a measure of agents $\mu^*(a, z)$ of households over the state space that is stationary under $P(a, z, a', z')$, such that the following conditions are satisfied:

1. Households optimize, whereby $g_a^*(a, z|w, r)$ and $g_l^*(a, z|w, r)$ solve equation (1).

2. Firms optimize given prices, whereby K and L satisfy

$$\begin{aligned} w &= F_L(K, L), \ and \\ r &= F_K(K, L) - \delta. \end{aligned}$$

3. The capital market clears

$$K(w, r) = K^*(w, r). \tag{9}$$

4. The labor market clears
$$L(w, r) = L^*(w, r).$$
5. The distribution of agents over states is stationary across time
$$\mu^*(a', z') = \int P(a, z, a', z')\mu^*(da, dz). \tag{10}$$

2. PARAMETERIZATION

In this section, I describe the parameters chosen for the problem. Given parameters, I use standard discrete state-space dynamic programming to solve the households' problem for given prices, and Monte Carlo simulation to compute aggregates.[5]

Preferences

I follow Pijoan-Mas (2006) in assuming standard time-separable utility with exponential discounting over sequences of consumption and leisure. Within any given period, utility is additively separable in consumption and leisure. The latter assumption is made primarily to remain close to the setting of Pijoan-Mas (2006). These preferences also have the feature that the marginal rate of substitution between consumption and leisure is invariant to the levels of consumption and leisure; this avoids introducing changes in behavior arising solely from changes in the long-run location of the wealth distribution that result from the relaxation of borrowing constraints. More precisely, households solve
$$\max_{\{c_t^i, l_t^i\}_{t=0}^{\infty}} E_\circ \sum_{t=0}^{\infty} \beta^t \left(\frac{c_t^{1-\sigma} - 1}{1-\sigma} + \lambda \frac{l_t^{1-\nu} - 1}{1-\nu}\right),$$
subject to the budget constraints described earlier in equations (4) and (5). The parameters β, σ, λ, and ν summarize preferences and are set following Pijoan-Mas (2006). In particular, I set $\beta = 0.945$, $\sigma = 1.458$, $\lambda = 0.856$, and $\nu = 2.833$. The choices for the discount factor β and the risk-aversion coefficient on consumption σ are standard in the literature and stem from long-run observations on interest rates. Relative to a standard model without valued

[5] I use 700 unevenly spaced grid points for capital and the method of Tauchen (1986) to generate an 11-state Markov chain to approximate the productivity process. I then simulate the economy for 200,000 periods to compute aggregates. All code is available on request. The interested reader should consult Nakajima (2007), which describes how to do discrete-state dynamic programming, and Nakajima (2006), which contains a helpful description of the algorithm used to solve the present model.

leisure, the parameters λ and ν are new. These parameters govern, respectively, the average amount of time spent working and the aversion to fluctuations in leisure. In particular, the larger λ is, the more leisure a household takes on average, and the larger ν is, the more a household will seek to avoid fluctuations in leisure.

Endowments

The parameter z, which denotes the log of labor productivity, evolves over time according to an AR(1) stochastic process,

$$z_t^i = \rho z_{t-1}^i + \varepsilon_t^i. \tag{11}$$

The random variable ε_t^i represents the underlying source of productivity risk and is assumed to be i.i.d. with standard deviation σ_ε. The parameter ρ determines the persistence of the shock. The mean of ε_t^i is set so that $Eq_t^i \equiv E\exp(z_t^i)$=1. I follow Pijoan-Mas (2006) to assign values of $\rho = 0.92$, and $\sigma_\epsilon = 0.21$. For computational purposes, I use the method of Tauchen (1986) to locate a 11-state Markov chain and associated transition matrix, which jointly approximate the process for productivity.

Technology

The consumption good in the economy is produced by an aggregate production technology that is Cobb-Douglas in aggregate effective labor input and physical capital. Thus,

$$Y = K^\alpha L^{1-\alpha}.$$

The single parameter governing production, α, is assigned according to capital's share of national income, as is standard (see e.g., Cooley 1995) and is, therefore, set $\alpha = 0.36$.

Borrowing Constraints

I will focus exclusively on equilibria in which prices are constant, and in which all borrowing is risk-free. I, therefore, abstract from fluctuations in interest rates, as well as the possibility of loan default. Given these restrictions, it is relevant to first locate the largest debt level that can be repaid with certainty in this economy. Let $\underline{\epsilon}$ be the lowest realization of productivity that is possible. Since the household is endowed with one unit of time, our insistence that all debt be repaid with certainty implies that it must be possible to repay a debt, even if it requires working full time. Denote the largest limit under which debt remains risk-free, by $\underline{b}_{nat}$ to follow the "natural borrowing limit" terminology introduced in Aiyagari (1994). For the present economy, $\underline{b}_{nat}$ is given by

$$\underline{b}_{nat} = -\frac{w\epsilon}{r}.$$

For standard preferences, including those that will be used in this article, households will never allow this borrowing limit to bind. This is simply because any plan that involves a positive probability of a state in which the marginal utility of leisure is infinite can be improved on by one that involves less consumption smoothing and less debt. The limit $\underline{b}_{nat}$ is clearly an upper bound on indebtedness among those studied here and will allow us to understand the implications of limits that are more stringent.

Modeling An Improvement in Credit Access

Credit access can improve in several mutually compatible ways. For example, transaction costs arising from the resources required to forecast borrowers' default risk may have been much higher in the past than they currently are. In turn, such a change would induce borrowing by lowering the interest rate faced by those who borrow, which is a topic explored in Athreya (2004). Second, if default is a possibility, and lenders may know more about borrowers now than in the past, credit risk may be better priced and thereby allow low-risk borrowers to avoid being treated like high-risk borrowers. In related work, Athreya, Tam, and Young (2007) evaluate this possibility. My goal here is to abstract from both default risk and transactions costs and, instead, evaluate the simplest form of an expansion in credit. I, therefore, study five economies in which the borrowing limit is increased by equal increments, from a benchmark value of 0 to a maximal level that approximates the natural borrowing limit. That is, $\underline{b} = \{0, -1, -2, -3, -4\}$. Given that I use the normalizations that (i) $Eq = 1$, (ii) households across all economies work approximately one-third of their time, and (iii) wages are near unity, the borrowing limits explored here cover a wide range from zero ($\underline{b} = 0$) to approximately 12 times median household labor income ($\underline{b} = -4$).

3. FINDINGS

The central experiment that I perform is to compare allocations and prices arising from the five different levels of the borrowing constraint defined earlier. The benchmark environment is taken to be one in which households are unable to borrow. That is, $\underline{b} = 0$. The remaining outcomes cover four levels of borrowing limits, up to a level $\underline{b}_5$ that is very close to the natural borrowing limit. All other parameters, including notably the stochastic process for labor productivity, are held fixed throughout the analysis.

Turning first to the behavior of economy-wide aggregates, Panels A and B of Table 1 summarize outcomes. There are several implications arising from

Table 1 Aggregates

Panel A						
Borr. Limits/Agg.	r	w	Y	K	C	Corr(a, z)
$\underline{b}_1$	0.0368	1.1884	0.6677	2.0051	0.5013	0.4976
$\underline{b}_2$	0.0410	1.1656	0.6563	1.9101	0.4978	0.4612
$\underline{b}_3$	0.0434	1.1531	0.6507	1.8584	0.4964	0.4336
$\underline{b}_4$	0.0448	1.1460	0.6488	1.8390	0.4961	0.4142
$\underline{b}_5$	0.0455	1.1425	0.6462	1.8133	0.4957	0.4048
Panel B						
Borr. Limits/Agg.	CV$_{cons}$	L	H	L/H	CV$_{labor}$	Corr(H, z)
$\underline{b}_1$	0.4149	0.3597	0.3644	0.9872	0.1149	0.0492
$\underline{b}_2$	0.4312	0.3599	0.3633	0.9905	0.1398	0.0640
$\underline{b}_3$	0.4475	0.3606	0.3636	0.9917	0.1608	0.0628
$\underline{b}_4$	0.4608	0.3611	0.3638	0.9925	0.1764	0.0621
$\underline{b}_5$	0.4694	0.3617	0.3644	0.9926	0.1853	0.0597

the interaction of labor supply and borrowing constraints for aggregates. A first finding is that, as with inelastic labor supply (e.g., Huggett 1993), the equilibrium interest rate rises monotonically with borrowing capacity. The fact that relaxing credit constraints leads the interest rate to rise is evidence of the "insurance," or consumption-smoothing benefits, conferred by the availability of a simple debt instrument. That is, when credit constraints are relaxed relative to the prevailing limit, all households will be able to use borrowing from each other to smooth consumption, and must rely less on accumulating claims in the capital stock alone. In equilibrium, this incentive forces the interest rate to rise to clear asset markets. This is noteworthy because debt has relatively poor insurance properties, as it requires borrowers to repay a fixed amount unrelated to their current circumstances. The rise of equilibrium interest rates with borrowing capacity is also a reflection of the "buffer-stock" behavior of households. Buffer-stock behavior refers to the feature of optimal decisionmaking under uncertainty and borrowing constraints whereby households preserve a reserve of either savings (if borrowing is altogether prohibited), or borrowing capacity, if the latter is allowed. In turn, as borrowing constraints are relaxed, households are, in effect, given a larger buffer, all else being equal, and so choose to hold fewer assets on average. However, the model is one in which some households are temporarily lucky in their productivity, while others are unlucky. Those who are lucky will choose to both work hard and save the proceeds. To the extent that the net effect of increased borrowing capacity is that households in the aggregate wish to hold fewer assets, the interest rate at which household savings exactly equals the increased borrowing demands of the average household must rise.

As displayed in Table 1, Panel A, in a steady-state equilibrium with interest rates higher than the benchmark economy, both the demand for capital by firms and output are lower. The stock of capital falls, by more than 10 percent, as borrowing constraints approach the natural limit. However, notice that output levels fall by substantially less. In particular, the decline in aggregate output is fairly small, approximately 3 percent. This is a direct reflection of the relatively low marginal product of capital in the benchmark equilibrium where borrowing is ruled out. Additionally, borrowing constraints seem to have only small effects on the aggregate efficiency of the labor input, as measured by the ratio of effective hours to raw hours. As borrowing constraints rise from $\underline{b}_1$ to $\underline{b}_5$, this ratio rises monotonically, by approximately one-half of one percentage point from 0.9871 to 0.9926.

The behavior of the economy in response to relaxed credit constraints is, thus far, analogous to that of an economy in which labor is supplied inelastically. Therefore, where precisely does the ability of households to vary work-effort manifest itself? A first measure lies in the volatility of household labor effort. The column "CV_{labor}" in Table 1, Panel B displays the ratio of the standard deviation of household labor effort to its mean. The clear pattern is that of a very large increase, a near-doubling, in variability of labor effort as households are allowed to borrow more. This suggests that households use labor supply less to buffer consumption than to take advantage of temporarily high productivity.

A second clear change in aggregate labor supply behavior arising from an increased ability to borrow is the large decrease in the correlation between wealth and labor supply seen in Table 1, Panel A. The nearly 20 percent decrease in the cross-sectional correlation of current assets and current labor supply is another reflection of the use by households, of labor for efficient production rather than constant consumption smoothing. In the economies studied, the high persistence of labor productivity means that the lucky are also the wealthy. When borrowing is ruled out, households that are productive have two reasons to work. First, the relative price of leisure is high. Second, the value of accumulating a buffer stock is high. In turn, it would be expected that once borrowing is made relatively easy, the former incentive remains, while the latter diminishes.

In contrast to the decline in correlation between wealth and labor hours arising from a relaxation of credit constraints, the correlation between productivity and labor supply generally rises with credit limits. Most noticeable, perhaps, is the low *level* of the correlation between hours and productivity; the level is approximately 0.06, very close to that level of 0.02 measured in the data by Pijoan-Mas (2006). Along this dimension, the model produces realizations under all specifications of the borrowing limit. In fact, the original work of Pijoan-Mas (2006) was aimed at demonstrating that incomplete asset markets could make labor effort insensitive enough to variations in productivity to match observations. The results in the present article suggest that relaxed borrowing constraints are not enough to substantially alter this result.

Interestingly, Table 1 shows that average number of hours worked as well as the average efficiency of the labor supplied remain fairly constant. The former, therefore, implies that credit constraints in this economy do not have strong effects on the total hours supplied, but as I show later, do matter for the timing of those hours. The same feature is true of the "effective" labor supply of households. This is a reflection of the fact that even though households may work more when productive, and less when not, the complex interaction of labor supply and household wealth results in there being a very weak relationship between borrowing capacity and the aggregate efficiency of the labor input. In particular, two things are worth noting. First, the preferences used in this article are not consistent with balanced growth as they display wealth effects. In turn, as wages fall, the substitution effect leading households to work less may be offset by a wealth effect that leads them to choose less leisure. Second, as will be discussed later, changes in borrowing constraints generate large changes in equilibrium wealth distributions. These effects appear to be offsetting for aggregate hours.

To get a clearer sense of how borrowing matters for labor supply, it is useful to study households grouped by wealth levels. In Table 2, I use the cutoffs defined by the quintiles of the benchmark wealth distribution, denoted Q_i, $i = 1, \dots, 5$. This way, a given wealth percentile always refers to a particular *level* of wealth, which allows one to disentangle the effects of borrowing constraints from the effects of changes in the wealth distribution that occur when credit limits are changed. A first finding is that the effect of borrowing constraints on the behavior of households *does* depend on wealth, especially for low-wealth households. In Panel A of Table 2, I display the labor hours supplied by households across (benchmark) wealth quintiles for households receiving the lowest productivity shock. It is immediately apparent that poor households supply substantially more hours when borrowing is ruled out than when it is allowed. As wealth rises, however, changes in borrowing constraints have much smaller effects on labor supply. The fact that wealth-poor households work so much when relatively unproductive when they cannot borrow, and much less when they can, is direct evidence that labor supply is an important device for smoothing consumption, at least for low-wealth, low-productivity households. From Panel B, it is clear that for households with 25th percentile productivity, labor supply varies less with both borrowing constraints and wealth across all wealth quintiles. This pattern is seen again in Panel C of Table 2, which covers median-productivity households. In sum, borrowing constraints alter the relationship between productivity and hours for the wealth-poor, but not for the wealth-rich.

The behavior of equilibrium outcomes is partially determined by the decisions households would take for wealth and productivity levels that are rarely,

Table 2 Labor Effort By Wealth and Productivity

			Panel A: Lowest Productivity		
Borr. Limits/Wealth Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
$\underline{b}_1$	0.0669	0.0473	0.0255	0.0080	0.0004
$\underline{b}_2$	0.0369	0.0303	0.0178	0.0093	0.0008
$\underline{b}_3$	0.0309	0.0281	0.0123	0.0071	0.0008
$\underline{b}_4$	0.0285	0.0235	0.0174	0.0061	0.0004
$\underline{b}_5$	0.0263	0.0215	0.0173	0.0085	0.0008
			Panel B: 25th Percentile Productivity		
Borr. Limits/Wealth Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
$\underline{b}_1$	0.1284	0.1080	0.0844	0.0549	0.0141
$\underline{b}_2$	0.1023	0.0880	0.0746	0.0479	0.0144
$\underline{b}_3$	0.0877	0.0806	0.0718	0.0460	0.0108
$\underline{b}_4$	0.0822	0.0768	0.0684	0.0468	0.0097
$\underline{b}_5$	0.0792	0.0741	0.0662	0.0499	0.0089
			Panel C: Median Productivity		
Borr. Limits/Wealth Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
$\underline{b}_1$	0.3472	0.3266	0.3024	0.2648	0.1865
$\underline{b}_2$	0.3232	0.3177	0.2960	0.2632	0.1804
$\underline{b}_3$	0.3235	0.3162	0.2944	0.2638	0.1754
$\underline{b}_4$	0.3224	0.3147	0.2960	0.2630	0.1725
$\underline{b}_5$	0.3209	0.3141	0.2998	0.2603	0.1723

or even never, observed. An example of this: even though the natural borrowing limit will never bind, the possibility that households may experience shocks, which require borrowing, leads them to be cautious. Therefore, it is instructive to study household decision rules, in particular for labor effort. Figure 1 contains optimal asset accumulation as a function of current wealth and productivity, across borrowing limits, while each panel of Figure 2 collects optimal labor supply. In both Figures 1 and 2, interest rates and wages are held fixed at their benchmark values (i.e., those obtained under borrowing limit $\underline{b}_1$), so that the effect of borrowing limits on decisionmaking is isolated.

I display results for three productivity levels that correspond approximately to the 25th percentile, 50th percentile, and 75th percentile of productivity. Three key points are apparent. First, the qualitative shape of optimal labor effort does not depend on the extent of borrowing capacity. In all three panels, the most productive households work substantially more than the least productive, except very near the borrowing constraint. Second, households with relatively low productivity are much more sensitive to increases in wealth than those with high productivity. Specifically, low productivity households reduce their labor supply with increases in wealth much more rapidly than their higher productivity counterparts. Third, what determines the sensitivity of labor effort to assets is the proximity to the borrowing constraint. In other words, being poor, per se, does not necessarily increase labor effort, but being

Figure 1 Optimal Savings Across Productivity Levels

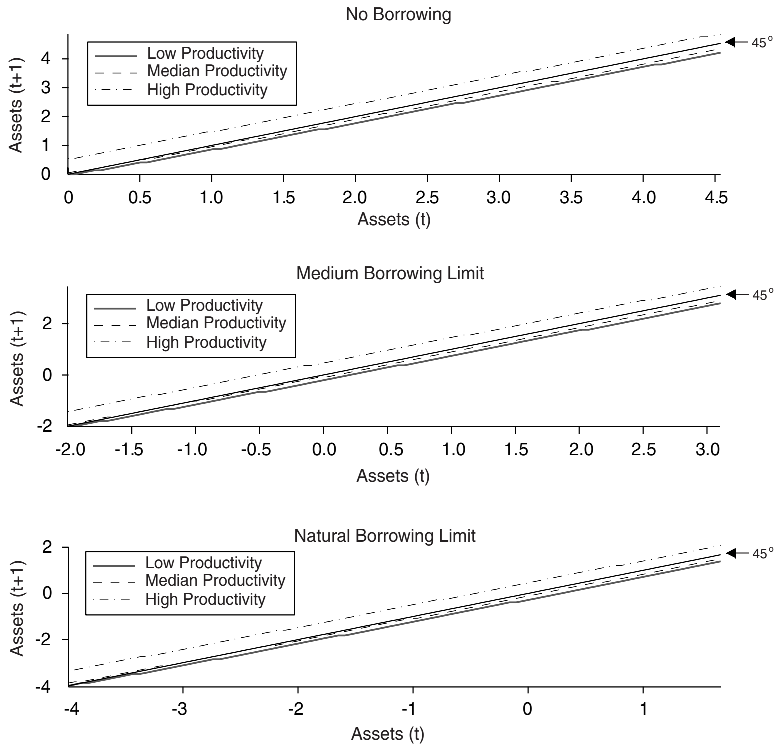

close to a borrowing constraint does. In fact, under both the medium borrowing limit and the natural limit, it is the households with the lowest productivity that work the hardest when near the borrowing limit. This is direct evidence of an inefficient use of time by households. Under complete markets, households would work most when most productive, not when least productive.

In order to better understand the role played by borrowing limits on labor effort, see Figure 3. The three panels of this figure point to three findings. First, labor supply depends on the proximity to the borrowing constraint, rather than on wealth itself. For example, in the top panel, households have received the "low" (25th percentile) level of productivity. At a level of zero wealth, when borrowing is prohibited, households work much longer than when either of the other two borrowing limits are imposed. A second feature illustrating the importance of the distance from the borrowing constraint is that in each panel of Figure 3, the wealth level at which a given labor supply is chosen "shifts" to the left by approximately the amount of the increase in borrowing constraint.

Figure 2 Optimal Labor Supply Across Productivity Levels

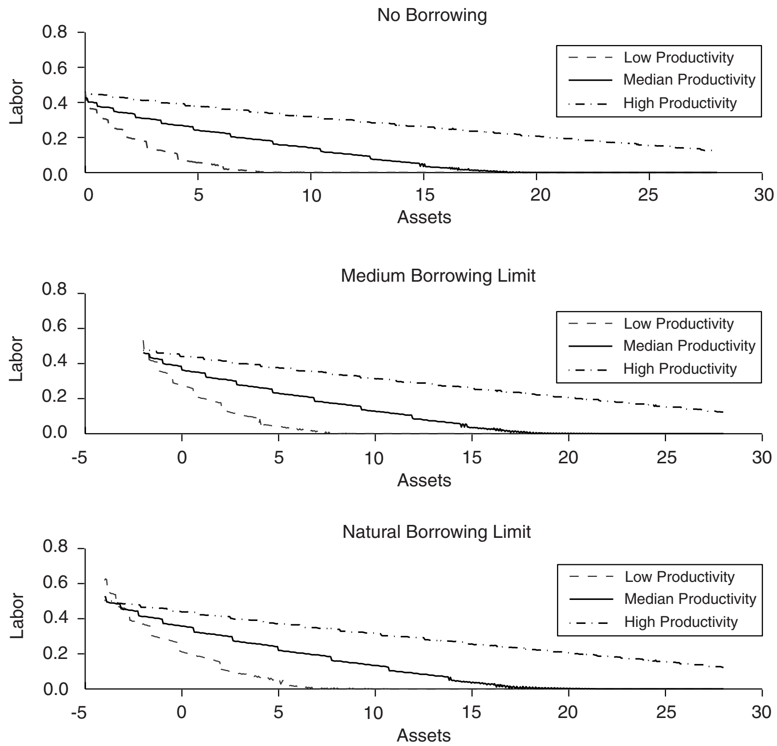

A second finding is that the importance of borrowing limits diminishes as productivity rises, as seen in the increasing similarity of labor supply decisions across wealth levels as productivity rises. Under the relatively high persistence of productivity shocks used in the model and thought to characterize U.S. household experience, high current productivity leads households to expect high future productivity. Conversely, a currently low-productivity household can reasonably expect more of the same in the future. Borrowing is then unlikely to provide a riskless stream of consumption, and households, therefore, respond by working harder. In sum, borrowing constraints alter the behavior of the low-productivity poor the most. A natural interpretation of this finding is that borrowing constraints create a set of workers who cannot "afford" *not* to work, even when they are extremely unproductive.

The preceding discussion described household behavior for arbitrary combinations of productivity and wealth. However, it is possible that precisely because households would "have" to work hard when close to the borrowing limit if they were unlucky, many might save at high enough rates to avoid

Figure 3 Optimal Labor Supply Across Borrowing Limits

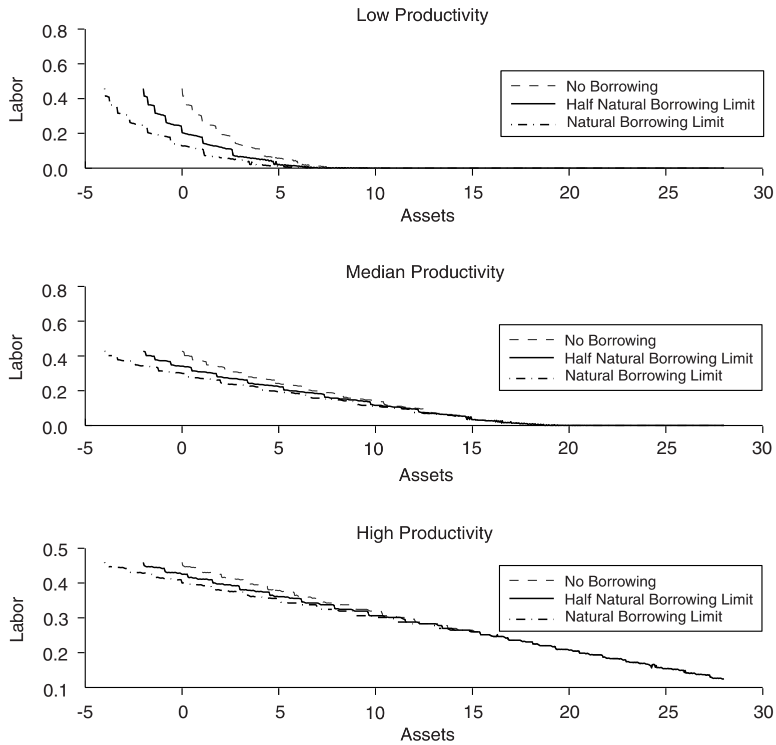

spending much time in such situations. In turn, observed labor supply might appear fairly insensitive to wealth. The outcomes documented in Table 3 are important because they show that the behavior embedded in the decision rules does indeed influence realized equilibrium outcomes. Table 3 contains three measures aimed at answering the question of “who works hard.” In each panel of the table, within a given row, borrowing limits are held fixed, while the columns represent quintiles of labor effort. For example, the first row, first column entry of Table 3, Panel A, gives the mean level of productivity of households who work the least, in the sense of being the lowest quintile of labor effort. The mean wealth level for the same subset of households is given by the analogous entry in Panel B. Similarly, the first row, fifth column entries of Panels A and B give the mean productivity and wealth of the hardest working 20 percent of households in the model when borrowing is not allowed. Panel C of Table 3 collects the conditional means of labor effort for households by productivity quintile. Here, it can be seen that for the least productive households (the column under the heading “Q1”) labor effort falls systematically as borrowing limits are expanded. Conversely, for the highest productivity households, labor supply increases as borrowing limits are extended. Moreover, given that productivity is lognormal, the increased effort of the highest productivity households further increases the "effective" labor supply to the economy.

The findings here suggest the following. One, in general, the hardest working are the poorest, especially those close to the borrowing constraint. Two, when borrowing is ruled out, the efficiency of those in the top quintile of hours is only about three-fourths (76.82 percent) of mean productivity. This is a striking indicator that the potential for inefficiently high (from a first-best perspective) supply of labor by the relatively unproductive highlighted in Figures 2 and 3 is a phenomenon that is actually realized in equilibrium. Three, as borrowing constraints are relaxed, this measure improves substantially and then stabilizes. This suggests that a move from no-borrowing to being able to borrow roughly two to three times the annual income ($\underline{b}_1 = -1$) generates large gains in the productivity of the labor input, with subsequent increases being less important.[6]

4. BORROWING LIMITS AND CONSUMER WELFARE

Economists' interest in the ability of consumers to borrow ultimately stems from a view that credit constraints may have important implications for welfare. However, measuring welfare gains arising from the relaxation of credit constraints under uninsurable risks is not as straightforward as it may seem. First, welfare can be measured by directly comparing the value functions for a household across any two specifications of the borrowing constraint, and then expressing the gains or losses in terms of differences in constant or "certainty equivalent" levels of consumption. Specifically, given a household state $(\hat{a}, \hat{z})$, let $V^{(i)}(\hat{a}, \hat{z})$ be the maximal utility attainable under a borrowing constraint $\underline{b}_i$. In the model, households derive utility from both consumption and leisure. Therefore, in order to convert utility into constant levels of consumption, we use the preferences over *consumption alone,* with the same curvature parameter $\sigma = 1.458$, and discount factor $\beta = 0.945$. We then compute the certainty equivalent as the scalar $ce(\hat{a}, \hat{z})$ that solves:

$$\sum_{t=0}^{\infty} \beta^t \left(\frac{ce(\hat{a}, \hat{z})^{1-\sigma} - 1}{1-\sigma} \right) = V^{(i)}(\hat{a}, \hat{z}),$$

which requires

$$ce(\hat{a}, \hat{z}) = [V^{(i)}(\hat{a}, \hat{z})(1-\beta)(1-\sigma) + 1]^{\frac{1}{1-\sigma}}.$$

[6] Another way to see this is that as borrowing limits expand, while the hardest working households are increasingly poor, as seen in the first column of Panel B, mean wealth does not fall one-for-one with borrowing limits.

Table 3 Who Works Hard? Mean Productivity and Mean Wealth by Labor Effort Quintiles

Borr. Limit/Effort Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
Panel A: Mean Productivity					
$\underline{b}_1$	0.8806	1.0257	1.0793	1.1591	0.7682
$\underline{b}_2$	0.8667	1.0439	1.0615	1.0939	0.9102
$\underline{b}_3$	0.8697	1.0380	1.0972	1.1538	0.9374
$\underline{b}_4$	0.8645	1.0491	1.0727	1.3187	0.9329
$\underline{b}_5$	0.8607	1.0612	1.1025	1.3492	0.9280
Panel B: Mean Wealth					
Borr. Limit/Effort Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
$\underline{b}_1$	5.1679	2.5128	1.3645	0.7993	0.1974
$\underline{b}_2$	5.7685	2.6604	1.2875	0.3535	−0.5227
$\underline{b}_3$	6.0268	2.5543	1.2419	0.4925	−1.1160
$\underline{b}_4$	6.2467	2.6087	1.0796	0.8819	−1.5769
$\underline{b}_5$	6.3246	2.6935	1.1539	0.8914	−1.8765
Panel C: Mean Effort					
Borr. Limit/Productivity Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
$\underline{b}_1$	0.3768	0.3539	0.3550	0.3616	0.3747
$\underline{b}_2$	0.3712	0.3507	0.3548	0.3638	0.3766
$\underline{b}_3$	0.3693	0.3508	0.3557	0.3654	0.3777
$\underline{b}_4$	0.3676	0.3514	0.3562	0.3666	0.3784
$\underline{b}_5$	0.3667	0.3521	0.3572	0.3673	0.3787

For any two borrowing limits $\underline{b}_i$ and $\underline{b}_j$, the difference $\Delta ce^{ij}(\hat{a}, \hat{z}) \equiv ce^{(i)}(\hat{a}, \hat{z}) - ce^{(j)}(\hat{a}, \hat{z})$ is then a measure of the effect of welfare effects of changes in borrowing constraints.

Of course, in economies with uninsurable risk, this measure will differ across households, as the latter differ in their asset levels a and productivity levels z. Therefore, in order to get an aggregate measure of welfare gains or losses, a weighted average is useful. Given $\Delta ce^{ij}(a, z)\ \forall a, z$, and the *current* long-run distribution of assets and productivity, $\mu^i(a, z)$, that prevails under a given borrowing limit, the average difference in certainty equivalents across two policies i and j is:

$$E_{\mu^i}[\Delta ce^{ij}] \equiv \int \Delta ce^{ij}(a, z) d\mu^i(a, z). \tag{12}$$

In sum, $E_{\mu^i}[\Delta ce^{ij}]$, gives the average gain or loss across inhabitants of an economy that will be experienced by an immediate move from the extension of borrowing limits from $\underline{b}_i$ to $\underline{b}_j$, given their current state.[7] One appropriate context for the use of this criterion is when borrowing limits $\underline{b}_i$ and $\underline{b}_j$ have prevailed for a long time in two different places, such as countries i and j,

[7] This idea originates in Benabou (2002) and is also applied in Seshadri and Yuki (2004).

for example. $E_{\mu^i}[\Delta ce^{ij}]$ then gives the average of the gains experienced by each household in country i if only they (or a subset of households of measure zero) were moved, *with their current wealth and productivity*, to country j.

An alternative welfare measure to the preceding is obtained by computing the weighted average of maximal utility a household could obtain if it began with a given level of assets and productivity. A common procedure for choosing the weights, originating in Aiyagari and McGrattan (1998), is to assign households a state according to the long-run distribution under borrowing limit $\underline{b}_j$, denoted $\mu^j(a, z)$. As before, converting these differences in expected utility into units of constant consumption yields a tangible measure of long-run or "steady state" welfare gains and losses $E_{\mu^j}[\Delta ce^{ij}]$. I denote this measure as:

$$E_{\mu^j}[\Delta ce^{ij}] \equiv \int \Delta ce^{ij}(a, z) d\mu^j(a, z). \tag{13}$$

Notice that the neither the measure in equation (12) nor that in equation (13) takes account of the transitional dynamics of wealth during the adjustment to the new steady state, and will, therefore, be potentially misleading. However, because the latter measure uses the long-run distribution under a proposed policy to weight welfare gains, it also does not control for *long-run* changes in the joint distribution of households over the state arising from changes in credit availability. For example, if constraints were relaxed relative to the present, in the long run there may be many more households holding large debts than before. In such a case, weighting the value functions by the distribution under relaxed borrowing limits will understate the welfare gains accruing to households who decumulated wealth in the aftermath of the policy change. In particular, an improved ability to borrow will lead many households to reduce their reserve of assets, which allows them a jump in consumption along the transition. It is beyond the scope of the current article to compute the welfare gains inclusive of the transition, but the two measures reported here are quite useful polar cases.

The preceding discussion makes clear that the central difference between the two measures above lies in the distribution used to weight households. The measure $E_{\mu^j}[\Delta ce^{ij}]$ has perhaps most relevance for generations arriving in the distant future, whose state-vectors will be drawn from the long-run distribution associated with the permanent imposition of the proposed change in borrowing constraints. It is useful to note that, under some circumstances, the model used here may be interpreted as consisting of (altruistically linked) overlapping generations of households. The implied per-period discounting of future generations by current ones is $\beta < 1$. However, a policymaker who values future generations the *same* as present ones (i.e., has an effective discount rate of $\beta = 1$) will view those born in the future as being at the mercy

Table 4 Borrowing Limits and Welfare, General Equilibrium

Borrowing Limits/Welfare	$\frac{E_{\mu^i}[\Delta ce^{1j}]}{E_{\mu^i}[c^1]}$	$\frac{E_{\mu^j}[\Delta ce^{1j}]}{E_{\mu^i}[c^1]}$
$\underline{b}_1$	–	–
$\underline{b}_2$	1.05%	0.36%
$\underline{b}_3$	1.56%	0.74%
$\underline{b}_4$	1.80%	1.01%
$\underline{b}_5$	1.93%	1.29%

of their ancestors' debt choices.[8] When large debts are feasible to incur, there may be many in the future who are destitute early in life. In turn, even though each of those households would be better off for any given value of the state, there may be so many low-wealth households under a lax credit constraint that overall average welfare decreases.

With the preceding discussion in mind, Table 4 presents the welfare consequences of more relaxed credit limits. All welfare changes are expressed in terms of the ratios of $E_{\mu^i}[\Delta ce^{1j}]$ and $E_{\mu^j}[\Delta ce^{1j}]$ to mean consumption under the tightest borrowing limit $\underline{b}_1$, given by $E_{\mu^i}[c^1]$.

The striking thing to note is that welfare grows much faster with the relaxation of borrowing constraints according to the welfare measure that uses the current distribution (i.e., the one prevailing prior to a policy change) than when measured using the long-run distribution following from a policy change. For example, a move from $\underline{b}_1$ to $\underline{b}_3$ appears more than twice as desirable under the former criterion than under the latter. What accounts for the difference? The answer lies in the changes in wealth accumulation induced by changes in borrowing limits. In the top panel of Figure 4, I present the distributions of assets obtained under the benchmark borrowing limit $\underline{b}_1$, an intermediate limit $\underline{b}_3$, and the most relaxed limit under consideration, $\underline{b}_5$. Notice that the latter contains a great deal of indebtedness, relative to the other cases. This feature is a striking implication of the "buffer-stock" behavior of these households. More ability to borrow simply pushes many households to hold wealth that keeps roughly at the same distance to the (now relaxed) borrowing constraint. In turn, any weighted average of utilities reflects the lower utility gains experienced by a systematically poorer population. However, such a measure ignores the increased consumption enjoyed en route to the new steady-state by all households that became able to borrow more. Finally, and naturally perhaps, I find that the gains relative to the no-borrowing benchmark are largest

[8] Limited liability for debts incurred by previous generations is a very widespread legal practice, and one that is potentially important in preventing such outcomes. Under this form of intergenerational limited liability, the weighted average using the current wealth distribution is perhaps more sensible.

Figure 4 The Wealth Distribution Across Borrowing Limits

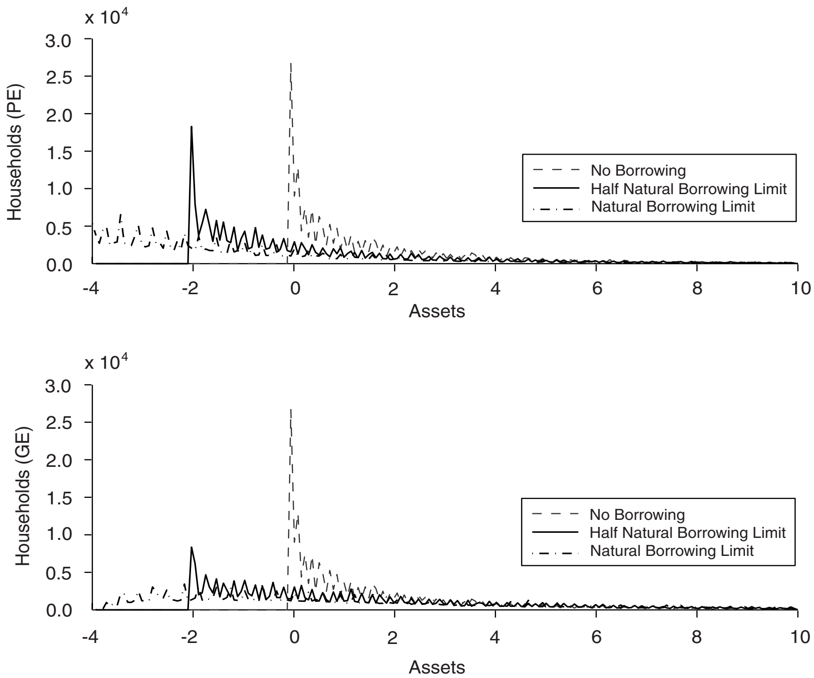

for initial relaxations in the constraints and, subsequently, grow much more slowly.

The Importance of General Equilibrium

In an incomplete-insurance economy, prices themselves are a source of risk. For example, a higher interest rate is good for households who receive good shocks, as they are likely to wish to save income. Conversely, high interest rates are bad for those who are unlucky, as they will find borrowing expensive. Therefore, it is useful to provide measures of welfare gains and losses coming from experiments in which the economy is treated as small and open. In such a setting, prices (wages and interest rates) can be viewed as being determined outside the economy.

Table 5 presents the welfare implications of relaxing credit limits when interest rates and wages are held fixed at their benchmark levels, i.e., when $\underline{b}_1$ is imposed.

In this case, the results are much larger in size than before for both measures of welfare, but most striking is the fact that the second measure shows

Table 5 Borrowing Limits and Welfare, Partial Equilibrium

Borrowing Limits/Welfare	$\frac{E_{\mu^i}[\Delta ce^{1j}]}{E_{\mu^i}[c^1]}$	$\frac{E_{\mu^j}[\Delta ce^{1j}]}{E_{\mu^i}[c^1]}$
$\underline{b}_1$	–	–
$\underline{b}_2$	1.21%	−3.25%
$\underline{b}_3$	2.03%	−6.26%
$\underline{b}_4$	2.63%	−8.97%
$\underline{b}_5$	3.09%	−11.28%

that welfare *falls* as credit limits expand. How can this be? The answer is that expansions in credit generate much more extreme changes in the long-run wealth distribution in partial equilibrium than in general equilibrium. This is seen by comparing the top and bottom panels of Figure 4. In partial equilibrium, the incentives of all households to borrow more under relaxed constraints is not met by a higher interest rate or by lower wages. In turn, the wealth distribution shifts even further to the left as households are allowed to acquire larger debts. Using the current distribution then gives households access to more credit at the relatively low benchmark interest rate and high benchmark wage, which is why the welfare gains are larger than in general equilibrium. However, precisely because the average household is much poorer in the long run under relaxed constraints, outcomes look much worse from the perspective of a household being assigned an initial state according to the long-run distribution.

The Distribution of Welfare Changes

A key aspect of the model used in this article is that it generates heterogeneity in current wealth, and as a result, in consumption and leisure, as well. Therefore, welfare gains from the relaxation of credit limits will differ across households. In order to provide insight into the gains or losses accruing to particular subsets of households, Table 5, Panel A gives the average difference in certainty equivalent across borrowing limits for households within each quintile of wealth, as defined by the benchmark economy's wealth distribution. That is, the welfare gain to households in quintile-k is

$$\text{Welfare Gain (quintile-}k\text{)} = \frac{E_{\mu^i_k}[\Delta ce^{1j}]}{E_{\mu^i_k}[c^1]}, \tag{14}$$

where μ^i_k is the distribution of the household state given that wealth lies within the kth quintile.

Table 6 collects a set of welfare gains organized by household wealth. Panel A displays partial equilibrium results, and Panel B, general equilibrium outcomes. The results are interesting along several dimensions. First, in both Panels A and B, it is clear that all households gain systematically from an increased ability to borrow. However, under partial equilibrium, the gains are largest by far for the wealth-poorest of households, and then fall steadily as households become wealthier. This is perhaps natural; richer households would seem to have less to gain directly from any increase in the ability to borrow. After all, such households are unlikely to need credit in the near future.

Once interest rates and wages are allowed to adjust to changes in borrowing capacity, the results change in a striking way. First, the welfare gains themselves are in general substantially smaller, and second, the biggest beneficiaries of a move to relaxed credit limits are currently wealthy. Why is this? Recall from Table 1 that an increase in credit limits leads to (i) a higher long-run interest rate and (ii) a lower long-run wage. How will this affect households of different wealth levels? A currently poor household that is likely to need to borrow will prefer, all else being equal, paying a lower interest rate *and* earning a higher wage. Its rich counterpart will want, by contrast, a higher interest rate, and will also care less about a fall in the wage; for the latter, capital income is the most important part of overall earnings. In the middle quintiles, these effects partially offset and result in smaller gains. As a result, there is a U-shaped relationship between welfare gains and wealth in general equilibrium. By contrast, under partial equilibrium, there are no price effects at all, which, therefore, leads welfare gains to shrink monotonically (but remain positive) as credit limits expand.[9] A useful interpretation of the findings above is that for a small open economy, the biggest beneficiaries of an expansion in credit will be the wealth-poor, while for a large closed economy, the currently rich can be expected to gain the most.

5. CONCLUDING REMARKS

In this article, I studied the interactions between credit markets, labor markets, and uninsurable idiosyncratic risk. The analysis proceeded by evaluating allocations across a variety of specifications of the ability of households to borrow against future income. The main results are as follows. First, the hardest working households are those who are least wealthy, and most strikingly, also the least productive. Second, credit access can play an important role in reducing high labor effort by low-productivity households. Third, the buffer-stock ten-

---

[9] The welfare gains are qualitatively and quantitatively very similar when households are ranked by current productivity, and therefore are not presented here. This result is natural given that productivity shocks are highly positively correlated with wealth (see Table 1, Panel A) and are highly persistent. Therefore, the wealth-poor value access to credit, while the wealth-rich value a higher return on savings. Correspondingly, welfare gains are again U-shaped across productivity quintiles in general equilibrium and positive, but monotone-decreasing in partial equilibrium.

Table 6 Welfare Gains by Wealth Quintile

Panel A: Across Benchmark Wealth Quintiles/Partial Equilibrium

Borr. Limits/Wealth Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
$\underline{b}_1$	–	–	–	–	–
$\underline{b}_2$	2.46%	1.43%	1.05%	0.71%	0.39%
$\underline{b}_3$	3.90%	2.46%	1.81%	1.29%	0.71%
$\underline{b}_4$	4.81%	3.22%	2.39%	1.74%	0.97%
$\underline{b}_5$	5.60%	3.76%	2.84%	2.08%	1.17%

Panel B: Across Benchmark Wealth Quintiles/General Equilibrium

Borr. Limits/Wealth Quintile	Q_1	Q_2	Q_3	Q_4	Q_5
$\underline{b}_1$	–	–	–	–	–
$\underline{b}_2$	1.57%	0.64%	0.47%	0.56%	2.01%
$\underline{b}_3$	2.11%	0.90%	0.64%	0.86%	3.27%
$\underline{b}_4$	2.21%	1.00%	0.71%	1.02%	4.05%
$\underline{b}_5$	2.35%	1.01%	0.73%	1.09%	4.46%

dencies of households imply that the distance from the borrowing constraint is often more important than the actual level of wealth in influencing labor effort. Fourth, measures of the welfare gains to current consumers show that there are significant benefits from expansions in credit access, and that these gains accrue disproportionately to the relatively poor and relatively rich.

There are many directions for future research along the lines developed here that appear productive. Two of these are as follows. First, a potentially fruitful avenue for future work is to augment the present model to include aggregate risk. This would allow for the coherent analysis of so-called "wealth effects," that have occupied the attention of numerous atheoretical studies and have been influential in the decisions of atheoretically-oriented policymakers. As it is, the model presented in this article suggests that aggregate relationships between *endogenous* variables such as consumption and wealth are the result of aggregating the behavior of households that differ substantially in their productivities, and more crucially, in their marginal propensities to work, consume, and save.

An important caveat to these results is that the expansion of credit was treated in this article as exogenous. The important work of Alvarez and Jermann (2000) demonstrates that it is quite possible that the same forces that lead households to want to borrow more may also allow them to do so. Krueger and Perri (2006), for example, apply this logic suggesting that when defaulters can be excluded from asset markets altogther, increases in income risk simultaneously make credit more beneficial and borrowing more feasible. The present work can be seen as measuring the effect on allocations arising solely from an increased ability to borrow, while abstracting from the additional effect on credit availability arising from a change in households' underlying environment.

A second line of research suggested by the results is that if recent financial innovation has genuinely altered household borrowing capacity, this in turn may imply a secular increase in the long-run average real interest rate.[10] An implication of a recent class of models of monetary policy is the desirability of consistently targeting a nominal rate that mirrors the underlying real interest rate in a nonmonetary economy. Thus, it may be useful to extend the model used here to allow for monetary policy.

REFERENCES

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics* 109 (3): 659–84.

Aiyagari, S. Rao, and Ellen McGrattan. 1998. "The Optimum Quantity of Debt." *Journal of Monetary Economics* 42 (3): 447–69.

Alvarez, Fernando, and Urban J. Jermann. 2000. "Efficiency, Equilibrium, and Asset Pricing with Risk of Default." *Econometrica* 68 (4): 775–98.

Athreya, Kartik B. 2004. "Shame As It Ever Was: Stigma and Personal Bankruptcy." Federal Reserve Bank of Richmond *Economic Quarterly* 90 (2): 1–19.

Athreya, Kartik B., Xuan S. Tam, and Eric R.Young. 2007. "A Quantitative Theory of Information and Unsecured Credit." Mimeo, Federal Reserve Bank of Richmond and University of Virginia.

Benabou, Roland. 2002. "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica* 70 (2): 481–517.

Chang, Yongsung, and Sun-Bin Kim. 2005. "On the Aggregate Labor Supply." Federal Reserve Bank of Richmond *Economic Quarterly* 91(1): 21–37.

Chang, Yongsung, and Sun-Bin Kim. 2006. "From Individual to Aggregate

[10] Of course, the U.S. is a (large) open economy, and the 1990s and 2000s saw large increases in the purchase of U.S. corporate and government debt by China and others. All else being equal, these purchases may well have kept real interest rates down. The results of the present article merely suggest that barring such changes, more borrowing capacity by U.S. households to borrow from each other implies a higher equilibrium real rate of interest.

Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy." *International Economic Review* 47 (1): 1–27.

Cooley, Thomas F., ed. 1995. *Frontiers of Business Cycle Research.* Princeton, NJ: Princeton University Press.

Domeij, David, and Martin Flodén. 2006. "The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates are Biased." *Review of Economic Dynamics* 9 (2): 242–62.

Edelberg, Wendy. 2006. "Risk-Based Pricing of Interest Rates for Consumer Loans." *Journal of Monetary Economics* 53 (8): 2,283–98.

Flodén, Martin. 2006. "Labor Supply and Saving Under Uncertainty." *Economic Journal* 116 (513): 721–37.

Flodén, Martin, and Jesper Lindé. 2001. "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics* 4 (2): 406–37.

Furletti, Mark. 2003. "Credit Card Pricing Developments and Their Disclosure." Federal Reserve Bank of Philadelphia, Payment Cards Center Discussion Paper No. 03-02.

Garcia-Escribano, Mercedes. 2003. "Does Spousal Labor Smooth Fluctuations in Husband's Earnings?" PhD diss. University of Chicago.

Huggett, Mark. 1993. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *Journal of Economic Dynamics and Control* 17 (5–6): 953–69.

Huggett, Mark, and Sandra Ospina. 2001. "Aggregate Precautionary Savings: When is the Third Derivative Irrelevant?" *Journal of Monetary Economics* 48 (2): 373–96.

Iacoviello, Matteo. Forthcoming. "Household Debt and Income Inequality, 1963–2003." *Journal of Money, Credit and Banking*.

Jayachandran, Seema. 2007. "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries." *Journal of Political Economy* 114 (3): 538–75.

Krueger, Dirk, and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies* 73 (1): 163–93.

Li, Wenli, and Pierre-Daniel Sarte. 2006. "U.S. Consumer Bankruptcy Choice: The Importance of General Equilibrium Effects." *Journal of Monetary Economics* 53 (3): 2,067–83.

Malapit, Hazel Jean, Jade Eric Redoblado, Deanna Margarett Cabungcal-Dolor, and Jasmin Suministrado. 2006. "Labor Supply

Responses to Adverse Shocks Under Credit Constraints: Evidence from Bukidnon, Philippines." PMMA Working Paper No. 2006-15.

Marcet, Albert, Francesc Obiols Homs, and Philippe Weill. 2007. "Incomplete Markets, Labor Supply, and Capital Accumulation." *Journal of Monetary Economics* 54 (8): *2,621–35.*

Nakajima, Makoto. 2006. "Note on Heterogeneous Agent Model: Labor-Leisure Choice and Fiscal Policy." Available at: http://www.compmacro.com/makoto/200701econ552/note/note_im_aiyagari_ext.pdf (accessed December 11, 2007).

Nakajima, Makoto. 2007. "Note on Neoclassical Growth Model: Value Function Iteration+Discretization." Available at: http://www.compmacro.com/makoto/200701econ552/note/note_ngm_disc.pdf (accessed December 11, 2007).

Narajabad, Borghan. 2007. "Information Technology and the Rise of Household Bankruptcy." Mimeo, Rice University.

Parker, Jonathan. 2000. "Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate." In *NBER Macroeconomics Annual 1999,* vol. 14, eds. Ben Bernanke and Julio Rotemberg. Cambridge, MA: MIT Press: 317–70.

Pijoan-Mas, Josep. 2006. "Precautionary Savings or Working Longer Hours?" *Review of Economic Dynamics* 9 (2): 326–52.

Seshadri, Ananth, and Kazuhiro Yuki. 2004. "Equity and Efficiency Effects of Redistributive Policies." *Journal of Monetary Economics* 51 (7): 1,415–47.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing Over the Life Cycle." *Journal of Monetary Economics* 51 (3): 609–33.

Tauchen, George. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economic Letters* 20 (2): 177–81.