

Interpreting High Returns to Microenterprise Investment

Tim Huegerich*

February 27, 2012

Preliminary - Comments welcome

Abstract

The average microenterprise in a developing country has very high marginal returns to capital (on the order of 100% annually), according to several recent studies. Financial constraints are part of the explanation for how these high returns persist. But, given these constraints, it remains to explain the limited self-financing by the households running these businesses. In particular, little is known about the elasticity of intertemporal substitution (EIS) of poor households, which describes how quickly households save in response to rates of return higher than their rate of time preference. If poorer households have a lower EIS, reflecting the difficulty of reducing their already low consumption to finance investment, then they will save slowly in response to high returns despite having a normal rate of time preference, or discount rate. However, separately identifying the EIS from the discount rate is difficult because it requires variation in the returns to capital, which is not present in many data sets. I examine extant data on microenterprises in Sri Lanka, including the results of a field experiment distributing grants to a random selection of firms, which aids in identifying households' varying returns to capital. I document gradual investment by most firms in response to their very high returns. I estimate the EIS of these households to be 0.03, substantially lower than recent estimates around 0.7 for US and UK households. This explains why households are found with such high unexploited returns. The main implication is that households invest cash grants productively with persistent benefits rather than consuming them away shortly after receiving them.

*I have greatly benefited from the guidance of John Kennan. Michael Carter introduced me to this topic area. I received helpful questions and suggestions from Rasmus Lentz, Salvador Navarro, Karl Scholz, Yongseok Shin, Chris Taber, Jim Walker, and other participants in the labor economics summer workshop at UW-Madison. Brad Barham, Jennifer Alix-Garcia, Laura Schechter, and other participants in the UW-Madison Development workshops provided helpful feedback, as well as John Micklewright. Many thanks to Chris Woodruff for providing data and answering multiple inquiries.

1 Introduction

The returns to investments in Sri Lankan microenterprises are very high, averaging above 80% annually, according to de Mel et al. (2008), and a number of other recent studies report comparable estimates in other contexts. These strikingly large returns suggest potentially big implications. One interpretation is that these unexploited high marginal returns show that microenterprise owners lack the patience or self-control to make good investments. The implicit assumption here is that they are at or near their steady state capital level, in which case the marginal rate of return would reflect their rate of time preference. A careful examination of the data, however, shows that is not the case. The poor households that own these businesses are investing in response to the high returns, albeit more slowly than wealthier owners would. In the absence of opportunities to borrow at a lower rate, they must sacrifice their already low current consumption to gradually self-finance their business capital. Thus, they are typically found far from their steady state level of capital despite already having been in business for many years.

A third of Sri Lankan workers are primarily employed in such microenterprises, as is typical of low- and middle-income countries, so the capital levels in these microenterprises affect the incomes of millions of people. Analyzing the causes of the borrowing constraints facing such households is the subject of a large literature, reviewed in Besley (1995) and Karlan and Morduch (2009). The focus of this paper, however, is understanding the investment behavior of households in response to high returns, taking their limited external financing as given. The large and heterogeneous returns among Sri Lankan microenterprises provide a valuable opportunity to study it.

A central question about investment behavior is how sensitive it is to the rate of investment return. The elasticity of intertemporal substitution (EIS) is the measure of this sensitivity, describing how easy it is to substitute future consumption for current consumption. It is a critical parameter in models with credit frictions because, together with the discount rate specifying the rate of return at which investment is zero, it determines how quickly agents with high returns will self-finance. It is also relevant for understanding inequality because it determines the rate of convergence of households with different wealth endowments that are otherwise identical *ex ante*. The intertemporal

preferences that determine investment in microenterprises also determine households' investments in education and health which play a larger role in shaping inequality. The advantage of studying investment in the context of microenterprises, though, is that the returns can be more directly measured than for investments in human capital.

There have been no direct measures of the EIS for poor households, although it is natural to expect poorer households to exhibit a lower EIS because it is relatively more painful to trade off their current low consumption, the common assumption of constant elasticity of substitution (i.e. CRRA) utility for convenience notwithstanding. While it is a myth that very poor households are unable to save anything at all, it is also a fallacy to assume they save at the same rate as wealthy households. At lower levels of consumption expenditure, necessities such as staple foods take up a larger expenditure share so that it is more difficult to sacrifice current consumption for higher future consumption. Under the stark assumption of an absolute minimum subsistence expenditure level, for instance, households approaching that subsistence level have a smaller proportion of their expenditure over which they have the discretion to invest and the EIS approaches 0 as consumption approaches subsistence. However, assuming a strict minimum subsistence level is not necessary for a utility function to exhibit increasing EIS, and there is also no reason to assume that the EIS is constant beyond consumption levels very near absolute subsistence, as does Stone-Geary utility. Rather, the idea is that optimal savings rates will generally be lower and less sensitive to the rate of return at lower consumption levels, other things equal.

The EIS is not separately identified from the rate of time preference without variation in the interest rate, and as a result many studies rely on previous estimates to fix its value. Existing estimates of the EIS use data from the US or UK, so they may not be applicable for much poorer households in developing countries. The Sri Lankan Microenterprise Survey (SLMS) provides a good source of identifying variation in rates of return. Previous studies estimating the EIS have used relatively small variations over time in the economy-wide interest rate. In contrast, the SLMS includes a field experiment in which households were randomly allocated substantial grants, and the induced changes in capital stock allow estimates of household-specific marginal returns to capital. de Mel et al. (2008) focus on the return to capital itself, but their rich data are suitable for going further to study the households' intertemporal preferences as well, with its panel of 11 measurements over three years and comprehensive measurements of changes in capital over time. I estimate the

EIS both by estimating the Euler equation directly and implementing a simulation-based estimation, as each method has different strengths.

The next section presents related literature on measuring the EIS and on microenterprises. Section 3 lays out the model of household business investment behavior, while Section 4 elaborates on the meaning of the EIS. Section 5 describes the data and presents some descriptive statistical analysis, including supportive evidence for the modeling assumptions. Section 6 presents the estimation methods and results. Section 7 examines some implications of the EIS varying with consumption level, and Section 8 concludes.

2 Literature Review

The work closest to this paper is Kaboski and Townsend (2011), which develops a structural model of microenterprise investments to evaluate a natural experiment expanding access to credit in Thailand. They do not claim to identify the EIS, however, as they do not have direct measures of returns to capital. Karlan and Morduch (2009) report an example of the ideal experiment for measuring the EIS. They offer savings accounts with different interest rates and measure the differences in takeup and savings volume. However, they only vary the interest rate by 1% annually and find that the differences are insignificant, consistent with a very low EIS but not permitting an estimate.

As discussed briefly above, modeling poor households' EIS accurately is important for quantitative evaluations of the effects of policies on microenterprise investment, such as subsidizing microcredit. The key qualitative implication of a low EIS, that the steady state capital level is not a good approximation for actual capital levels at any given point in time, means that current assets have medium- to long-term implications even in the absence of nonconvexities. Assuming that the current distribution of capital among microenterprises is stationary, as do Buera et al. (2011), can then lead to very poor estimates of policy impacts. Townsend and Ueda (2006) is another recent study that reports the impact of the EIS on the dynamics of inequality, using 1.0 as a baseline value for the EIS and 0.67 as a robustness check.

Estimating the elasticity of intertemporal substitution (EIS) is the subject of a large literature, largely stemming from interest in rich country asset prices and business cycle fluctuations. Most

estimates follow theoretical work in assuming a constant elasticity of substitution. Recent estimates of the EIS for average U.S. households are around 0.7, as reviewed in Attanasio and Weber (2010). Thus values in the range $[0.67, 1]$ are used to calibrate the EIS in a wide range of papers, even those modeling poor households in developing countries.

A few papers have estimated preferences that allow the EIS to depend on consumption level, but their identification of variation in the EIS may come from unrelated aspects of the data. Attanasio and Browning (1995) and Blundell et al. (1994) model the EIS varying with demographics as well as consumption level to show that allowing heterogeneity in preferences allows the life-cycle consumption model to fit features of the data previously considered incompatible with it. Attanasio and Browning (1995) report estimates that reject the isoelastic specification and find EIS increasing with consumption, using data from the UK Family Expenditure Survey. However, they do not report the implied values for the EIS. Blundell et al. (1994) run a similar exercise with a different specification of utility, and they find the opposite result that EIS decreases with consumption. In any case, both papers impose a relationship between the dependence of the EIS on consumption level and the changes in consumption level induced by changes in family size.¹ It seems likely that identification of the variation in EIS is coming from variation of consumption with family size rather than directly reflecting varying sensitivities to interest rates, which could explain their conflicting results using different function forms. *Atkeson and Ogaki (1996)* is, to my knowledge, the only paper to estimate the EIS for poor households in developing countries, using the ICRISAT data from India. They find an increasing EIS, reporting that the wealthiest villagers in their sample exhibited an EIS 60% higher than the poorest. However, they impose a functional form that forces a specific relationship between intertemporal substitution and intratemporal substitution among different goods, and they report some evidence against this restriction.²

In each of the papers estimating a variable EIS, as well as in other studies using constant elasticity preferences, the source of variation used to identify the EIS is time variation in economy-wide ex-post real interest rates. Unfortunately, it is unclear that these differences in ex-post interest rates reflect variation in the ex-ante interest rate relevant for each household's consumption and

¹See Appendix A for details on the utility functional forms they use.

²Browning and Crossley (2000) show that for any utility function additive across both time and goods, as is the functional form used by Atkeson and Ogaki (1996), the EIS is increasing in consumption unless the expenditure share of each good is constant across wealth levels. See Appendix A for details on the functional form for utility they use.

savings decisions. Low income households may be faced with systematically higher interest rates, confounding any comparison of their consumption growth rates with higher income growth rates. Also, variations in the interest rate are correlated with macroeconomic fluctuations and therefore with the likelihood that households will be liquidity constrained.

A number of other papers provide indirect evidence for an EIS increasing with consumption by showing it is consistent with a number of otherwise anomalous observations. Guvenen (2006) shows how an increasing EIS in the US can explain the different estimates of the EIS obtained by different techniques. Among others, Bliss (2004) points out that an increasing EIS can help explain persistent cross-country inequality. A lower EIS can also explain why poor households would persistently hold debt at very high interest rates, as provided by payday lenders and other forms of fringe banking in the U.S. and elsewhere. Mullainathan and Shafir (2009) argue that poorer households have less “financial slack,” which they define as “the ease with which one can cut back consumption to satisfy an unexpected need.” While they suggest an alternate interpretation, this notion of financial slack corresponds well to the EIS. There is also a wide literature studying the savings behavior of poor households, also reviewed by Besley (1995) and Karlan and Morduch (2009). This work seldom refers to the concept of the EIS, so a contribution of this paper is to make the connection between these empirical studies and the key parameter used to model savings for general equilibrium policy experiments.

Finally, the high rates of return inferred from the Sri Lankan field experiment of de Mel et al. (2008) central to the EIS identification strategy of this paper are in line with other recent estimates of the rates of return in microenterprises.³ Udry and Anagol (2006) use creative methods to bound the returns to capital in Ghana between 60% and 300% annually. Duflo et al. (2008) estimate an average rate of return of 69.5% annually for fertilizer for small Kenyan farms. And Dupas and Robinson (2010) estimate monthly returns of 5.5% for microenterprises run by women in rural Kenya.

This paper may be the first to systematically examine the self-financing behavior of households in response to these high return, but previous studies have ventured some speculation about why these households are not saving more, taking for granted that saving more would be optimal. In

³It also is not contradicted by older literature on the returns to capital in developing countries, as reviewed in Banerjee and Duflo (2005).

particular, de Mel et al. (2008) conclude by asking “...what prevents firms from growing incrementally by reinvesting profits?” and speculating about shocks to households or time-inconsistent preferences. Other papers positing time-inconsistent preferences as an explanation for low savings include Banerjee and Mullainathan (2010) and Dupas and Robinson (2011). Indeed, Dupas and Robinson (2010) and Duflo et al. (2010) provide evidence that seemingly insignificant interventions do boost investment by a small amount, suggesting that investment may otherwise be sub-optimal due to psychological factors or practical difficulties in saving.⁴ The approach of the present paper is instead to rationalize the observed level of investment and thereby examine whether we should necessarily expect anything much different.

3 Model

I model credit-constrained households’ investment behavior as that of agents maximizing utility over consumption at different times, assuming additively separable utility with rate of time preference ρ :

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+\rho} \right)^t u(c_t)$$

This is a form of the life-cycle model of consumption and saving. I follow Kaboski and Townsend (2011) in adopting an infinite horizon, interpreting the household as a dynasty. The time period is a month, so the discount rate ρ is the monthly interest rate at which an agent (facing no uncertainty) would prefer to save or borrow in order to keep consumption constant across time periods. The commonly assumed additively separable form with identical period utility at each time t , $u(\cdot)$, entails some well-known assumptions, including that the discount rate is independent of consumption level and that the coefficient of relative risk aversion is the reciprocal of the elasticity of intertemporal substitution.

Each household has access to a private production technology, the household business. The following assumptions allow a focus on the basic implications of household preferences for capital accumulation without borrowing:

⁴In some cases, the emphasis on such mechanisms leads to the complete omission of the trade-off between current consumption and investment described by the EIS. Duflo et al. (2010) propose a model in which poor farmers facing binding credit constraints simply maximize discounted expected income. They not only use this simplified model to motivate to paper, but they also calibrate it for a rough welfare analysis.

1. The household knows its business's expected productivity, which only changes over time in perfectly foreseen ways common to all households - there are no other persistent changes.
2. Business capital is the household's only asset - no other borrowing or saving is possible. Because of their lack of access to outside financing, the choice of capital for the business is not separable from the consumption and savings choices of the household.
3. Labor supply and other business inputs are unmodeled. Labor supply to the business can be considered perfectly inelastic, while other inputs are employed optimally, their costs netted from business income. In particular, entry and exit of the household business are not modeled. The household expects to continue running its existing business indefinitely.
4. Business capital is liquid and investments are fully reversible.
5. Households have identical preferences.

The data provide some evidence for the reasonableness of the first three assumptions, presented in Section 5.2. The last two assumptions are made now for tractability, with hope of relaxing them in future extensions.

So a household with wealth a can use it for either consumption or capital for the household business. Monthly business income increases with capital as $A_{it}f_i(k_{it})$. The household's initial revenue productivity A_i grows over time at some rate g due to learning-by-doing so that $A_{it} = A_i \exp(gt)$.⁵ The household's revenue productivity is also affected by short-term shocks, whether difficulties in production or fluctuations in demand, that are not known prior to choosing the month's capital level. In particular, suppose that business income is given by $\pi = A_{it}f_i(k_{it})\nu_{it}$, where ν_{it} is an i.i.d. log-normally distributed shock. Capital depreciates at rate δ , and households also receive outside income of y_i^o each month. Then we can write the household's utility of having assets a_{it} at time t recursively as:

⁵Households' differing labor supplies to their businesses may be considered subsumed into their household specific revenue productivities A_i . Labor supply is assumed perfectly inelastic, so it does not change over time and has no impact on intertemporal choices.

$$\begin{aligned}
V_{it}(a_{it}) &= \max_{c_{it}, k_{i,t+1}} u(c_{it}) + \frac{1}{1+\rho} \mathbf{E}_{it} [V_{i,t+1}(k_{i,t+1} + y_i^o + \pi_{i,t+1}(k_{i,t+1})\nu_{i,t+1} - \delta k_{i,t+1})] \\
&\text{s.t. } c_{it} + k_{i,t+1} = a_{it} \\
&\pi_{i,t+1}(k_{i,t+1}) = A_i \exp(g(t+1)) f_i(k_{i,t+1})
\end{aligned}$$

Assume $f(k)$ strictly concave with the marginal return at the initial capital stock at least equal to the discount rate: $A_i f'(k_{i0}) \geq \rho$. Then if we use \sim to signify the optimal value of a choice variable, we can write the Euler equation characterizing the optimal values as:

$$u'(\tilde{c}_t) = \frac{1}{1+\rho} \mathbf{E}_t \left[u'(\tilde{c}_{t+1}) (1 + \pi'_{it}(\tilde{k})\nu - \delta) \right] \quad (1)$$

which just says that, each month, the household chooses capital \tilde{k} so that the marginal costs and marginal benefits of capital are equal. It is known that a steady state capital level k_i^* exists and households starting with less capital will invest over time, approaching the steady state at some rate. We will analyze the rate of capital adjustment more closely in the next section.

4 Elasticity of Intertemporal Substitution

The elasticity of intertemporal substitution (EIS) is defined as the percentage increase in the growth rate of consumption associated with a 1% increase in the marginal rate of substitution (MRS) between consumption in consecutive time periods. In particular, we will consider the Frisch elasticity of substitution⁶ between consumption in two consecutive periods and define the EIS $\phi(c_t, c_{t+1})$ to be positive:

$$\phi(c_t, c_{t+1}) \equiv - \frac{d \ln\left(\frac{c_{t+1}}{c_t}\right)}{d \ln\left((1+\rho) \frac{u'(c_t)}{u'(c_{t+1})}\right)} \Big|_{c_t}$$

For shorthand, we will use $\phi(\tilde{c}_{t+1}) \equiv \phi(\tilde{c}_t, \tilde{c}_{t+1})$. It is straightforward to show that:

$$\phi(c_t, c_{t+1}) = \frac{-u'(c_{t+1})}{c_{t+1} u''(c_{t+1})}$$

⁶Frisch here means holding c_t (and thus $u'(c_t)$) constant as the MRS is varied.

Intuitively, the lower the EIS - that is, the more inelastic is intertemporal substitution - the greater the increase in consumption tomorrow required to compensate for a decrease in consumption today. In particular, as discussed above, we might expect a lower EIS at lower levels of consumption at which a proportional decrease in consumption is relatively more painful.

Now, consider the model above with no uncertainty. Then the Euler equation can be rewritten to show that the MRS between consecutive periods is equal to the return to capital, the marginal rate of technical substitution (MRTS). For clarity of exposition, let's set aside the endogeneity of the MRTS and write it as $1 + r$, substituting $r = \pi'(\tilde{k}) - \delta$:

$$(1 + \rho) \frac{u'(\tilde{c}_t)}{u'(\tilde{c}_{t+1})} = 1 + r$$

Then, substituting into the definition of the EIS, we find that it tells us by what percentage the growth rate of consumption increases in response to a 1% increase in the marginal return to capital:

$$\left. \frac{d \ln\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\right)}{d \ln(1 + r)} \right|_{\tilde{c}_t} = \phi(\tilde{c}_{t+1}) \quad (2)$$

So a lower EIS means that the household's optimal consumption growth rate is less sensitive to differences in the rate of return, which is why we might expect a poorer household's investment level to rise less in response to a higher rate of return.

And here we see that the relationship between the growth rate of consumption and the natural log of the return to capital is linear. Taking the anti-derivative of (2), and using the fact that consumption growth is zero at $r = \rho$ in this case with no uncertainty:

$$\ln\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\right) = \phi(\tilde{c}_{t+1})(\ln(1 + r) - \ln(1 + \rho)) \quad (3)$$

Using first-order Taylor approximations for the natural log functions in equation (3), we see that the growth rate of consumption is approximately equal to the elasticity of intertemporal substitution times the difference between the marginal return to capital and the rate of time preference:

$$\frac{\tilde{c}_{t+1}}{\tilde{c}_t} - 1 = \phi(\tilde{c}_{t+1})(r - \rho) \quad (4)$$

4.1 Allowing Uncertainty

Allowing for uncertainty as in the model in Section 3 adds additional terms to the equation determining consumption growth but leaves the role of the EIS unchanged. Starting from the Euler Equation (1), consider the second-order Taylor expansion of $u'(\tilde{c}_{t+1})$ centered at c_t .⁷ I have again substituted $1 + r \equiv 1 + \pi'_{it}(\tilde{k})\nu - \delta$ for readability, but r here is a random variable unknown at time t :

$$u'(\tilde{c}_t) = \frac{1}{1 + \rho} \left(u'(\tilde{c}_t) \mathbb{E}_t [1 + r] + u''(\tilde{c}_t) \mathbb{E}_t [(\tilde{c}_{t+1} - \tilde{c}_t)(1 + r)] + \frac{1}{2} u'''(\tilde{c}_t) \mathbb{E}_t [(\tilde{c}_{t+1} - \tilde{c}_t)^2(1 + r)] \right) \quad (5)$$

Rearranging and using the formula for covariance $\text{Cov}[X, Y] = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$ and the first-order Taylor approximation $\ln(1 + x) = x$,⁸ we have the following formula for expected consumption growth:

$$\mathbb{E}_t \left[\ln \left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t} \right) \right] = \phi(\tilde{c}_t) \ln(\mathbb{E}_t [1 + r]) + \beta_c \quad (6)$$

where β_c is the sum of three terms, one representing impatience, one representing the risk of investing capital, and one term representing precautionary savings:

$$\begin{aligned} \beta_c = & -\phi(\tilde{c}_t) \ln(1 + \rho) - \frac{\text{Cov}_t \left[\frac{\tilde{c}_{t+1} - \tilde{c}_t}{\tilde{c}_t}, 1 + r \right]}{\mathbb{E}_t [1 + r]} \\ & + \frac{1}{2} \left(\frac{-c_t u'''(\tilde{c}_t)}{u''(\tilde{c}_t)} \right) \left(\frac{\text{Cov}_t \left[\left(\frac{\tilde{c}_{t+1} - \tilde{c}_t}{\tilde{c}_t} \right)^2, 1 + r \right]}{\mathbb{E}_t [1 + r]} + \mathbb{E}_t \left[\left(\frac{\tilde{c}_{t+1} - \tilde{c}_t}{\tilde{c}_t} \right)^2 \right] \right) \end{aligned} \quad (7)$$

So the dependence of expected consumption growth on expected return to capital is the same as in the case without uncertainty, linear with the EIS as the slope. Then, for a given EIS and return to capital, consumption growth is lower with a higher discount rate ρ (more impatience), lower for more risky returns (higher variance of the shock ν), and higher for more variance of consumption growth. Note that the size of the precautionary savings term depends on the coefficient of relative

⁷This derivation follows Dynan (1993), generalizing to an uncertain rate of return.

⁸In particular, using the approximations $\frac{\tilde{c}_{t+1} - \tilde{c}_t}{\tilde{c}_t} \approx \ln \left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t} \right)$ and $\frac{1 + \rho}{\mathbb{E}_t [1 + r]} - 1 = \ln \left(\frac{1 + \rho}{\mathbb{E}_t [1 + r]} \right)$, which are valid for a small consumption growth rate and $\frac{1 + \rho}{\mathbb{E}_t [1 + r]}$ close to 1, respectively.

prudence, $\frac{-cu'''(c)}{u''(c)}$.⁹

5 Data

The Sri Lanka Microenterprise Survey, first reported in de Mel et al. (2008), includes 9 waves of quarterly data on small household businesses, from March 2005 to March 2007, as well as 2 additional follow-up surveys at 6 month intervals which are not yet incorporated in the results here. The sample consists of 408 firms¹⁰ starting with less than \$1000 of capital, excluding land and buildings. Half of the firms are engaged in retail trade, with the rest divided among small-scale manufacturing and services. What is referred to throughout as “capital” for these businesses includes both physical capital, or equipment, and working capital, or inventories, in roughly equal parts. Further details are provided in the Data Appendix.

Households were told after the first round of the survey that there would be a random drawing for prizes of either \$100 or \$200. The average initial capital of the firms was \$275 and the average monthly household income was \$114, so these prizes are substantial. 124 grants were given out immediately after the first wave of the survey, with 104 others selected to receive their prizes after the third wave (but not told so until then). Half of the grants were given in cash, while half were in-kind grants of materials or equipment chosen by the household. Of the in-kind grants, 61% or purchases were of inventories or raw materials, working capital as opposed to physical capital. Thus, they are not primarily using these grants to make large purchases that they would otherwise need to save for a long time to purchase.

de Mel et al. (2008) estimate a treatment effect on real monthly profits of 5.7% of the treatment amount. They estimate the marginal return to capital, using the treatments as instruments for current capital stock and adjusting profits for changes in owner’s labor hours, as between 4.6% and 5.3% per month.¹¹ These estimates of 5% monthly returns correspond to 80% annual returns with

⁹While there is no general relationship between EIS and prudence imposed by additively separable utility, for the commonly used constant EIS, or CRRA, utility with EIS ϕ , the coefficient of relative prudence is $1 + 1/\phi$.

¹⁰Their full baseline sample includes 618 firms in southwestern Sri Lanka, but I follow them in excluding the 210 firms located along the southern coast and directly affected by the December 2004 tsunami for the purposes of measuring the returns to capital for firms under ordinary conditions. The tsunami had no significant effect on national GDP growth, so it is plausible that households not directly affected saw little economic impact. Also, the estimated treatment effects of the random grants were no different for the firms farthest inland than they were for the firms in the middle third of the sample between the coastal and the farthest inland regions.

¹¹The standard error for the direct treatment effect is 2.2. Standard errors on the estimates of the marginal return

compounding.¹² The returns seem less implausible when reported in absolute terms, a \$5 increase in monthly income.

As for effects on investment, firm owners report investing an average of 58% of the cash grants in capital in the months immediately after receiving them. Roughly consistent with this direct self-report, de Mel et al. (2008) estimate a treatment effect of 99% of the grant amount on average capital stock in the subsequent 6-8 quarters after receiving the grant. There is no significant difference in outcomes between those given cash and those given in-kind grants.

The statistical equivalence of the cash and in-kind grants is consistent with households being fully aware of the high marginal returns available to them. If they were unaware that returns were so high, households given cash would likely invest very little of the grant in their businesses while those given in-kind grants would be forced to experiment with investing. Only those given in-kind grants would then discover the high returns and benefit from them.

de Mel et al. (2008) anticipate a number of possible concerns about the experimental design and provide a variety of evidence against the most common objections to these results. One concern is that firms may over-report their profits in response to receiving the grants, to portray themselves as worthy of the grant - or worthy of further assistance. As one of their arguments against this source of error, they point out that there is no effect of the treatment on profit-to-revenue ratios or on the reported markup of the sales price over marginal cost, so if treatment respondents are inflating their profit numbers, they must be adjusting all their other responses in a coordinated way as well. de Mel et al. (2009) reports the results of multiple approaches to assessing the accuracy of the microenterprises' reporting of profits, including having research assistants directly observe the firm's transactions at unannounced times to generate independent measures for comparison. They find that firms systematically under-report revenues by 30% but that their reporting of profits is more accurate on average.

to capital are both 2.3.

¹²While de Mel et al. (2008) provide convincing estimates of the average returns to investment, it is more difficult to determine the risk involved in investing. Is the distribution of returns is highly skewed so that the reported mean is the result of a few lucky ones and many receiving no benefit? They do ask firm owners a series of questions eliciting their subjective perceptions about the uncertainty of profits (though not directly about the uncertainty of marginal returns). When interacted with the coefficient of variation of this subjective distribution, the treatment effect on profits is lower for those expecting more uncertainty, suggesting that high returns do not primarily reflect a risk premium. They also point out that the firms receiving grants do not change their line of business or introduce new products, which suggests they are not using the grants in especially risky ways. Finally, I estimate quantile treatment effects, which show that the treatment does not lead to a skewed distribution of returns. [Calculated but not shown here yet.]

Another concern is that negative spillovers of the grants on control firms may bias the estimated returns upwards. de Mel et al. (2008) do find evidence of negative spillovers of the grants, but only for one industry (bamboo goods) for which several nearby firms happen to have been included in the survey. The results are unaffected when controlling for spillover by either excluding the spillover industry or using distance from treated firms as a measure of potential spillover.

5.1 Descriptive Statistics

The preliminary sample used below is restricted to those households without missing data, a total of 249.¹³ The descriptive statistics shown in Table 1 show that the control group (N=103) and those given grants (N=146) do not have significant differences in initial capital or profits or other indicators of business success, such as education or age of business.

The very large budget share of food expenditures, with a median of 64%, suggests the limited ability of these households to cut back expenditures for self-financing. The median household spends less than 1% on the combined categories of recreation and electronics, further demonstrating the very little slack in the typical household budget. Food share decreases on average by roughly 1 percentage point for each \$100 of monthly household income, while combined expenditure on recreation and electronics increases by roughly 1 percentage point for each \$1000 of additional monthly household income. These observations are consistent with the argument for EIS increasing with consumption level.

The savings rate shown is calculated as the average monthly increase in business capital over a given period divided by the average household income in that period. The savings rate shown for the treatment group is only for the second year, omitting the period in which they received their grants since that dramatically affects their capital levels. That the savings rate for the treatment group is even higher on average is especially remarkable because it shows that these households are not, on average, drawing down their capital stocks after initially stashing the grants in their business capital. It is consistent with the notion that they are gradually investing toward a higher,

¹³In ongoing work, I use the full sample adjusting for attrition and other causes of missing data. Comfortingly, de Mel et al. (2008) report that attrition is relatively mild: with 369 of the original 408 firms remaining in the survey throughout the 9th wave. Including all forms of missing profit or capital data, 11.5% firm-period observations are missing. This rate of missing data is slightly higher for those not receiving grants, 14.3% for the control group compared to 9.6% for the treatment group, but controlling for this differential attrition has little effect on the estimate of returns to capital. Indeed, we might expect the least profitable control firms to be most likely to attrit, which would actually bias the estimate of marginal returns upward.

desired capital level.

One anomaly in Table 1 is that reported household consumption increased by less for treatment households than for control households, despite their larger gains in capital and income. A more detailed breakdown reveals that the big difference in consumption growth rates between the groups is during the first year, the period in which treated households received their grants. In the second year, the consumption of treated households grew at a slightly higher rate on average, but not enough to catch up to control households over the entire two year period.

5.2 Evaluating Modeling Assumptions

The data provide evidence of the reasonableness of the first three modeling assumptions listed in Section 3: changes in a household's business capital are driven by accumulation through self-financing rather than changes in productivity; households do very little borrowing or drawing down other assets to purchase business capital; and ignoring changes in labor supply has little effect on inferred returns to capital.

First, the identification strategy described in Section 6.1 relies on the assumption that changes in capital for a given household over time are driven by the capital choice described in the model rather than by unmodeled changes in the household business's revenue productivity over time. The most direct evidence that income drives capital decisions is that households are observed investing the large majority of the cash grants they are given into capital - firms given cash grants report much larger capital increases than control firms, equivalent to 74% of the grant amounts, within six months of receiving the grants. Aside from that, even control firms are observed investing an average of \$19 per quarter in their businesses, consistent with gradual self-financing.

In addition, households report very limited borrowing opportunities, consistent with their apparent self-financing of capital investment. Only 5% of firms report ever having received a loan for their business, even from family and friends. Households were also asked about the terms at which they could borrow from a moneylender if they needed to. While 29% report access to a monthly interest rate of 3% or lower, 44% report a monthly interest rate of 10% or higher. They also report that they would only be able to borrow a limited amount. 35% report they would only be able to borrow less than \$100 and just 30% report being able to borrow more than \$300.

As for other assets households might use for liquid savings, 80% of households report owning

some gold jewelry, but the data do not include information on the value of such assets or changes in them over time. We do know what recipients of cash grants report spending the money on. They spend an average of 12% on other forms of savings, in addition to the 58% they spent on business capital. However, that their capital increases by an average of 74% of the grant amount, relative to control households, within two quarters of receiving the grant, suggests that they may have soon invested that 12% in their businesses as well. It should also be possible to compare the sum of reported household consumption and capital investment with household income. Unfortunately, the data are not good enough to make sense of this comparison. Total household savings calculated as household income minus household consumption is slightly negatively correlated with capital investment.

We should expect households to increase labor supply to their businesses along with capital,¹⁴ complicating the interpretation of increased profits associated with capital increases. For instance, for the 50% of households that initially have a household member working for a wage in the external labor market, income from this external work may decrease as a result of substitution of labor into the household business. Indeed, as reported in de Mel et al. (2008), the \$100 grants induce an average increase in business owner labor of 4 to 6 hours per week over the 50 hours per week average initially. However, the average effect of the \$200 grants is -1 to 2 hours and not statistically significant, and they find no impact on labor supply to the business from other members of the household. In any case, de Mel et al. (2008) report that adjusting profits by subtracting imputed wages for owner labor hours has very little effect on estimated returns to capital, decreasing their point estimate by just around 1 percentage point. They still estimate returns to capital of 4.6% or 5.3%, the estimates noted above, using two different methods for imputing labor income. For now, I ignore adjustments in labor supply in what follows.

¹⁴Only the extreme cases of perfectly inelastic labor supply or perfectly substitutable capital and labor would imply otherwise.

6 Estimation

6.1 Identification

The profit function $\pi_i(\cdot)$ is identified from observations of capital and business income over time as households accumulate capital. The household-specific flexibility of the business income function is important because accurate measures of households' marginal returns are essential for identifying the EIS. Given each household's marginal return to capital, though, the separate identification of ρ and $\phi(c)$ is straightforward. Conditional on consumption level, $\phi(c)$ is identified by how consumption growth varies with marginal return to capital, as in equation (2). Then, given $\phi(c)$, ρ is identified by the level of the growth rate of capital.

Now, the business income function will be better identified for those households that received grants because they have more observed variation in capital. The preference parameters are then identified primarily from these households' saving behavior. But the knowledge that the control households have the same distribution of returns to capital as the treatment households, by random assignment, allows the investment behavior of control households to help in the identification of preferences as well.

6.2 Estimating a Linearized Euler Equation

We can estimate the approximated Euler equation (6), assuming identical preferences across households since the available data will not allow identification of much heterogeneity across households. Likewise, assume constant EIS utility $u(c) = \frac{1}{1-\gamma}c^{1-\gamma}$, with $\frac{-u'(c)}{cu''(c)} = 1/\gamma = \phi$. Then Equation (6) can be written:

$$\ln\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\right) = \phi \ln(\mathbb{E}_t[1 + r_{t+1}]) + \beta_c + \epsilon_{t+1} \quad (8)$$

where the constant term β_c can be interpreted as in (7) above. The error term ϵ_{t+1} is the response of \tilde{c}_{t+1} to the time $t + 1$ shock and is therefore uncorrelated with all information available to the household at time t , particularly $\mathbb{E}_t[1 + r_{t+1}]$. If there is multiplicative i.i.d. log-normal measurement error in consumption, this will also show up in the error term. Note also that the existence of other assets in the household's portfolio would not affect the validity of the Euler equation for

business capital.

We use the annual observations of total household consumption, along with estimates of the rate of return during the corresponding period, to estimate a form of this equation. Since the data on consumption are annual rather than monthly, we use the annualized form of (8) in which $\ln(E_t[1 + r_{t+1}])$ is simply multiplied by 12. We will also instrument for the rate of return estimates using demographic information related to ability, all of which are known at time t . Instrumenting also helps correct for measurement error in the rate of return estimates. Specifically, I use as instruments the following variables found by de Mel et al. (2008) to have a significant effect on returns to capital: gender, age of business, years of education of the owner, and owner’s score on a digit span intelligence test, as well as dummies for manufacturing and services (with retail being the default industry category). Excluding these variables from the equation for consumption growth imposes the assumption that they do not impact the variance of income or any other factor that affects consumption growth, except through the expected return to capital. This is a potentially problematic assumption I am addressing in ongoing work.

Limiting the sample to grant recipient households, for which there is substantial exogenous variation in capital, I first calculate the change in total household income (observed annually) divided by change in capital for the annual interval spanning their receipt of the grants as a rough measure of that household’s rate of return $\hat{r}_i = \frac{y_{i12}^{tot} - y_{i0}^{tot}}{k_{i13} - k_{i1}}$.¹⁵ Using total household income is preferable to using business income because it includes any decreases in outside income that might come from increasing labor supply to the household business along with increasing capital. The results are shown in Table 2. The point estimate for $\hat{\phi}$ is 0.028 and it is statistically different from 0 at the 5% level. It is also clearly significantly smaller than standard estimates for the US and UK of 0.7 or higher.

This Euler equation approach is robust to alternate assumptions about other assets or borrowing opportunities, as well as to how labor supply to the business affects outside income. However, it depends on a rough measure of marginal returns, and it also uses an approximation of the Euler

¹⁵In ongoing work described below, I specify the monthly business income function for each household, and I will discuss the implied rates of return there. Estimation of the monthly business income function outside of the full estimation is frustrated by the fact that capital stock is only observed at the end of the months for which business income is measured (that is, π_{it}^m and $k_{i,t+1}^m$ are observed for $t = 0, 3, 6, \dots$). But k_{it+1} is correlated with shocks to π_{it} .

In any case, I use this rough estimate of \hat{r}_i for this more direct estimate of the EIS. The levels are likely incorrect, but the differences still appear to be informative.

equation. In particular, we know that the approximation is not good for rates of return much higher than the discount rate. (See footnote 8.) In addition, we have so far used neither the data available on capital investments nor the data for control households not given grants. A simulation-based estimation allows a more satisfactory treatment of some of these limitations, as well enabling explicit policy experiments.

6.3 Production Function Specification

For a simulation-based estimation, the production function must be specified explicitly. Assume $f_i(k) = f_i^0 + k^\alpha$, where f_i^0 is production with no capital for household i , so that business income is $\pi_{it}(k) = A_{it}(f_i^0 + k^\alpha)$. Importantly, the extra degree of heterogeneity provided by f_i^0 ensures that the marginal return to capital is not determined by the levels of business income and capital alone. Rather, the marginal return to capital can be different even for the same capital level and business income level:

$$\pi'_{it}(k) = \alpha \frac{\pi_{it}(k) - A_{it}f_i^0}{k}$$

This gives the specification the flexibility to accurately capture household-specific differences in marginal return to capital, which is essential for identifying the EIS.

It is also important to allow for correlation between the initial capital level observed for a firm and its productivity. Assuming the distribution of productivity is uncorrelated with initial capital would be problematic because we expect that high productivity households will tend to have higher initial capital. Additionally, allowing for correlation between A_i and f_i^0 is needed to allow separate estimation of the level of business income and the marginal return to capital. Thus, I specify the following discrete distribution over (A_i, f_i^0) with $2J$ support points, indexed by $j = 1, \dots, J$ and either L or H : $(A_{1L}, f_{1L}^0), (A_{1H}, f_{1H}^0), \dots, (A_{jL}, f_{jL}^0), (A_{jH}, f_{jH}^0)$. For each j , $A_{jL} < A_{jH}$. The probability of each support point depends on a firm's initial reported capital in the following way. A firm observed with no initial capital has equal probability of either (A_{1L}, f_{1L}^0) or (A_{1H}, f_{1H}^0) , while a firm with the maximum $k_{i1} = \$1000$ has equal probability of either (A_{jL}, f_{jL}^0) or (A_{jH}, f_{jH}^0) , and a firm with initial capital at the appropriate quantile (a $(J - 1)$ -tile) of the initial capital distribution has equal probability of either (A_{jL}, f_{jL}^0) or (A_{jH}, f_{jH}^0) for $1 < j < J$. A firm with initial capital

k_{i1} between two of the quantiles has probability over the 4 possible support points in proportion to the distance between k_{i1} and the corresponding quantile of initial capital.

6.4 Simulated Maximum Likelihood

Using the above specification of the production function and heterogeneity in productivity, as well as the CES utility function, we can simulate the model presented in Section 3 can maximize the simulated likelihood. The depreciation rate δ is fixed at 1% monthly,¹⁶ as the data do not identify this parameter.

The remaining parameters are estimated by simulated maximum likelihood. First, the outside income for each simulated household y_i^o is fixed at the average observed for that household.¹⁷ Monthly business income and capital is simulated over a three year period. Simulated households do not anticipate the treatment grants or expect their possibility, but they arrive in the specified period as unanticipated outside income.¹⁸

Recall that business incomes are observed in the data only every third month, along with the capital stock at the end of those months (that is, π_{it} and $k_{i,t+1}$ for $t = 0, 3, 6, \dots, 24$). This mismatch in timing between observed profits and observed capital obstructs direct analysis of the relationship between profits and capital changes, but it is readily accomodated in a simulation-based approach. Assume the econometrician observes firm profits and capital levels with multiplicative log-Normal measurement errors, $\varepsilon_{it} \sim \text{LogNormal}(0, \sigma_m^2)$: $k_{it} = k_{it}^* \varepsilon_{it}$, $\pi_{it} = \pi_{it}^* \varepsilon_{it}$. Each simulation, indexed by h , is initialized with the observed initial business income, corrected by a realization of simulated measurement error: $\pi_{i0}^h = \pi_{i0} / \varepsilon_{i0}^h$. Then the initial capital is inferred using the inverse profit function: $k_{i0}^h = \pi_{i0}^{-1}(\pi_{i0}^h / \nu_{i0}^h)$. Capital and income are then simulated going forward for $t > 0$: k_{it}^h and π_{it}^h for $t = 1, 2, 3, \dots, 24$. Finally, for a given simulation draw, the likelihood contribution of each observation for $t > 0$ is the pdf of a normal distribution with mean k_{it}^h (or π_{it}^h) and standard deviation σ_m . (The standard normal pdf is denoted ϕ_n to distinguish it from the EIS.) If we use θ as a vector of all the estimated parameters, that is: ρ , ϕ , g , $\{(f_{jL}^0, A_{jL}), (f_{jH}^0, A_{jH})\}_{j=1}^J$, σ_ν , and

¹⁶This is towards the upper end of the range of dpreciation rates found by Schündeln (2007) for Indonesian manufacturing firms. However, they find limited evidence that financially constrained firms experience higher depreciation, consistent with the theoretical result in Udry and Anagol (2006).

¹⁷Outside income y_i^o can only be inferred annually when both business income and total household income are observed.

¹⁸All firms were given a \$25 payment after the fifth survey wave, and this payment is also included as unanticipated income in the simulation.

σ_m , the likelihood contribution of a household can be written as follow, with w indexing waves of the survey since only every third month of data is observed:

$$l_i(\theta) = Pr(k_{i1}|\theta, \pi_{i0}) \prod_{w=1}^8 Pr(k_{i,3w+1}|\theta, \pi_{i0}) Pr(\pi_{i,3w}|\theta, \pi_{i0})$$

In practice, this is calculated by Monte Carlo integration over the shocks and measurement error in the initial profits used to initialize the simulation. For H simulations of the model, the likelihood contribution is calculated as:

$$l_i(\theta) = \sum_{h=1}^H \frac{1}{H} \phi_n \left(\frac{k_{i1} - k_{i1}^h}{\sigma_m} \right) \prod_{w=1}^8 \phi_n \left(\frac{k_{i,3w+1} - k_{i,3w+1}^h}{\sigma_m} \right) \phi_n \left(\frac{\pi_{i,3w} - \pi_{i,3w}^h}{\sigma_m} \right)$$

The results of the simulation-based estimation are not yet complete.

7 Discussion

One way to understand the significance of the very low estimate of the EIS from the Euler equation estimation is to compare its implications with the implications of the standard value currently used for calibrating the EIS, 0.7. For instance, compare the time it would take for a household to double their consumption if they had EIS 0.7 instead of 0.028, the estimate from Table 2. Abstracting away from risk, Equation 4 gives the growth rate of consumption as approximately $\phi(r - \rho)$. So a household with a monthly rate of return of 5% above their rate of time preference ($r - \rho = 0.05$) would choose a consumption growth rate of 3.5% monthly with an EIS of 0.7. Their consumption would double in just 20 months. A household with an EIS of 0.028, however, would choose consumption growth of just 0.14% monthly, doubling their consumption in 500 months, more than 40 years - and this is assuming that their marginal return stays at 5% rather than diminishing.¹⁹ It is this huge difference in convergence rates that may require a qualitative difference in how the savings of poor households are modeled. They may not be found very often near a steady state.

There is not sufficient variation in consumption levels within my sample to estimate significant differences across different consumption levels, much less whether the EIS changes with consumption

¹⁹From Table 1, a direct measure of the savings rate in the data has a median of 0.6% monthly, while the median consumption growth over a two year period is 16%, which corresponds to 0.6% monthly. So the consumption growth observed in the data is of the same order of magnitude as this simplified approximation, although it is somewhat higher due to precautionary savings.

level for a given household. So my data are not informative on whether these poor Sri Lankan households exhibit a lower EIS because they have different preferences from typical US or UK households or because they share the same preferences with the property that the EIS increases with consumption level. To illustrate this idea, I provide an example of a period utility function $u(c)$ such that the EIS matches the measured levels at both the average consumption level in my sample and the average monthly consumption in the U.S.

7.1 Exploring an EIS Increasing in Consumption Level

I propose a non-standard functional form for period utility to allow the intertemporal elasticity of substitution ϕ to vary with the consumption level in a flexible but parsimonious way.²⁰ The marginal utility is defined as $u'(c) = \exp\left(\frac{\gamma}{\lambda}(c^{-\lambda} - 1)\right)$,²¹ so the utility function is:

$$u(c) = \int_0^c \exp\left(\frac{\gamma}{\lambda}(x^{-\lambda} - 1)\right) dx$$

The EIS for this utility specification is:

$$\phi(c) = \frac{1}{\gamma} c^\lambda$$

so the log of the EIS is linear in the log of consumption. Notice that CES, or CRRA, utility is a special case, for $\lambda = 0$.²² The CARA functional form is also a special case, for $\lambda = -1$. To match

To match the EIS at both \$30 per month, $\phi(\$30) = 0.03$, and the U.S. level of \$3000 per month, $\phi(\$3000) = 0.7$, we need $\lambda = 0.68$ and $\gamma = 333$. This implies that with every 10% increase in monthly consumption, this hypothetical individual will exhibit a 6.8% higher EIS. To reiterate, this is a utility function that matches the measured EIS in both Sri Lanka and the U.S., under the assumption that agents in both places share identical preferences. My data do not allow me to distinguish between this hypothesis and the possibility that households in the two places simply

²⁰A review of other utility functions allowing varying EIS in the literature and why they are less appropriate for this exercise is in Appendix A.

²¹This is a generalization of a specification suggested by Bliss (2004) in which the EIS increases linearly with c (i.e. $\lambda = 1$). There is no closed form for $u(c)$ for general λ , but it is not needed for many applications, in which only $u'(c)$ matters.

²²By L'Hopital's Rule, $\lim_{\lambda \rightarrow 0}(c^{-\lambda} - 1)/\lambda = -\ln(c)$, so $\exp(\gamma(c^{-\lambda} - 1)/\lambda) = \exp(-\gamma \ln(c)) = c^{-\gamma}$ the marginal utility for CES/CRRA preferences.

have different preferences, each with constant EIS regardless of consumption level.²³ These two possibilities have very different implications for the future investment behavior of the Sri Lankan households in my sample as they accumulate wealth. If they truly have constant EIS at the very low level of 0.03 regardless of their consumption level, then they would continue to invest very slowly, decreasing their rate of consumption growth as they encounter diminishing returns to capital. However, if their EIS is only low at their current low level of consumption but will increase as they accumulate wealth, they may exhibit an increasing rate of consumption growth even if their return to capital is diminishing.

We can examine this possibility in a simple model. Using the approximation of $3, \frac{\tilde{c}_{t+1}}{\tilde{c}_t} - 1 = \phi(\tilde{c}_t)(r - \rho)$, we can derive a simple formula for how large θ would need to be for the rate of consumption growth to be increasing with wealth. Assume utility has the above form so that $\phi(c) = \frac{1}{\gamma}c^\lambda$. For ease of exposition, take $r - \rho = f(k) = Ak^\alpha$, and assume that consumption for the purposes of determining the EIS is a constant fraction x of income $xf(k)$.²⁴ Whether the rate of consumption growth is increasing or decreasing turns out not to depend on x . Then:

$$\frac{d}{dk} \left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t} \right) = \frac{d}{dk} (\phi(xf(k))f'(k)) = \frac{d}{dk} \left(\frac{1}{\gamma} (xAk^\alpha)^\lambda (\alpha Ak^{-(1-\alpha)}) \right) = \frac{f'(k)(xf(k))^\lambda}{\gamma k} (\lambda\alpha - (1-\alpha))$$

So the consumption growth rate decreases with wealth for $\lambda > \frac{1}{\alpha} - 1$. For instance, if $\lambda = 0.68$ as speculated above, the consumption growth rate would be increasing for any $\alpha > 0.6$. That is, for the rate of increase of EIS with consumption consistent with Sri Lankan and U.S. households having identical preferences, the consumption growth of poor Sri Lankan households will take place at an increasing rate if their marginal return to capital decreases less than 4% for each 10% increase in capital.

8 Conclusions

I find that poor Sri Lankan households exhibit a much lower intertemporal elasticity of substitution than has been previously assumed for poor households in developing countries. If the EIS is

²³The reality may also be some combination of EIS varying with consumption level and heterogeneity in preferences.

²⁴Numerical simulations confirm this is a good approximation for constant EIS for capital levels sufficiently far from the steady state capital level. Changes in the fraction of income consumed, x , will be small with the increasing EIS and decreasing marginal returns counteracting each other. The effect of the changing fraction of income consumed on the EIS can therefore be treated as second-order.

generally much lower at lower consumption levels, this has important implications for understanding poor households' savings and borrowing behavior. The implied low potential for self-financing underscores the importance of access to well-functioning credit markets for the efficient allocation of capital among microenterprises and the consumption growth of poor households, barring more direct approaches to reducing wealth inequality. It also informs expectations about the extent to which low-income households will self-finance high return investments in education or health. The primary policy implication is simply that poor households make productive investments when given the resources to do so. They just invest slowly with their own very limited resources.

References

- ATKESON, A. AND M. OGAKI (1996): "Wealth-varying intertemporal elasticities of substitution: Evidence from panel and aggregate data," *Journal of Monetary Economics*, 38, 507–534. 5, 26
- ATTANASIO, O. AND M. BROWNING (1995): "Consumption over the Life Cycle and over the Business Cycle," *The American Economic Review*, 85, 1118–1137. 5, 27
- ATTANASIO, O. AND G. WEBER (2010): "Consumption and saving: models of intertemporal allocation and their implications for public policy," *NBER Working Paper*. 5
- BANERJEE, A. AND E. DUFLO (2005): "Growth Theory Through the Lens of Development Economics," in *Handbook of Economic Growth*, ed. by P. Aghion and S. N. Durlauf, Elsevier B.V., vol. 1A. 6
- BANERJEE, A. AND S. MULLAINATHAN (2010): "The Shape of Temptation: Implications for the Economic Lives of the Poor," *NBER Working Paper*. 7
- BESLEY, T. (1995): "Savings, credit and insurance," *Handbook of development economics*, 3, 2123–2207. 2, 6
- BLISS, C. (2004): "Some implications of a variable EIS," *Unpublished*. 6, 22
- BLUNDELL, R., M. BROWNING, AND C. MEGHIR (1994): "Consumer demand and the life-cycle allocation of household expenditures," *The Review of Economic Studies*, 57–80. 5, 27

- BROWNING, M. AND T. CROSSLEY (2000): “Luxuries are easier to postpone: A proof,” *Journal of Political Economy*, 108. 5
- BUERA, F., J. KABOSKI, AND Y. SHIN (2011): “The Macroeconomics of Microfinance,” . 4
- DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2008): “Returns to Capital in Microenterprises: Evidence from a Field Experiment,” *Quarterly Journal of Economics*. 2, 3, 6, 7, 12, 13, 14, 16, 18, 27
- (2009): “Measuring microenterprise profits: Must we ask how the sausage is made?” *Journal of Development Economics*, 88, 19–31. 13
- DUFLO, E., M. KREMER, AND J. ROBINSON (2008): “How High are Returns to Fertilizer? Evidence from Field Experiments in Kenya,” *AER Papers and Proceedings*. 6
- (2010): “Nudging farmers to use fertilizer: theory and experimental evidence from Kenya,” . 7
- DUPAS, P. AND J. ROBINSON (2010): “Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya,” *Unpublished*. 6, 7
- (2011): “Why Don’t the Poor Save More? Evidence from Health Savings Experiments,” . 7
- DYNAN, K. (1993): “How prudent are consumers?” *Journal of Political Economy*, 1104–1113. 11
- GUVENEN, F. (2006): “Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective,” *Journal of Monetary Economics*, 53, 1451–1472. 6
- KABOSKI, J. AND R. TOWNSEND (2011): “A structural evaluation of a large-scale quasi-experimental microfinance initiative,” *unpublished*. 4, 7
- KARLAN, D. AND J. MORDUCH (2009): “Access to Finance,” in *Handbook of Development Economics*, ed. by D. Rodrik and M. Rosenzweig, vol. 5. 2, 4, 6
- MULLAINATHAN, S. AND E. SHAFIR (2009): “Savings Policy and Decision-Making in Low-Income Households,” in *Insufficient funds: savings, assets, credit, and banking among low-income households*, Russell Sage Foundation Publications, 121. 6

SCHÜNDELN, M. (2007): “Appreciating Depreciation: A Note on Physical Capital Depreciation in a Developing Country,” . 20

TOWNSEND, R. AND K. UEDA (2006): “Financial Deepening, Inequality, and Growth: A Model-Based Quantitative Evaluation¹,” *Review of Economic Studies*, 73, 251–293. 4

UDRY, C. AND S. ANAGOL (2006): “The Return to Capital in Ghana,” *AER Papers and Proceedings*, 96, 388–93. 6, 20

A Previous Specifications for Varying EIS

Previously used functional forms for marginal utility that allow the EIS to vary with consumption either do not allow sufficient flexibility or have the property that the EIS is not well-defined for some plausible values of consumption.

First, the general form for HARA utility²⁵: $u'(c) = (c - c_s)^{-\gamma}$ allows for an increasing EIS but does not allow much flexibility in the relationship between the EIS and consumption. Consumption below c_s is not allowed and the EIS is only substantially lower than the asymptotic value $1/\gamma$ at rarely observed consumption levels that are very low multiples of the subsistence level:

$$\phi(c) = \frac{1 - \frac{c_s}{c}}{\gamma}$$

Another approach is to model utility over two (or more) goods as the sum of CES utilities for each good: $u'(c_A, c_B) = c_A^{-\gamma_A} + c_B^{-\gamma_B}$, assuming $\gamma_A > \gamma_B$ so that good B is a luxury in the sense that its budget share rises with total expenditure. Then the EIS for total consumption expenditure depends on the budget share $\eta_j(c)$ of each good, rising from $1/\gamma_A$ to $1/\gamma_B$ as total consumption expenditure c rises:

$$\phi(c) = \eta_A(c) \frac{1}{\gamma_A} + \eta_B(c) \frac{1}{\gamma_B}$$

Atkeson and Ogaki (1996) combine subsistence levels and additive multi-good preferences in what they call extended addilog utility, adding additional flexibility. In any case, these specifications

²⁵This is also called Stone-Geary preferences in demand analysis.

impose a precise relationship between budget shares and total EIS, so that budget shares may be used to identify the EIS. The identification of the EIS then comes through the functional form.

Blundell et al. (1994) and Attanasio and Browning (1995) use alternate utility specifications that have Euler equations that are linear or quadratic in logs of consumption. However, they lack flexibility and have the undesirable property of only being well-defined for consumption below a certain maximum \bar{c}_i , which may differ by household. The specification of Blundell et al. (1994) is equivalent to:

$$\phi(c) = \frac{1}{\gamma(1 - \frac{c^{1-\gamma}}{\bar{c}_i^{1-\gamma}})}$$

while the specification of Attanasio and Browning (1995) is equivalent to:

$$\phi(c) = \frac{1}{\gamma(1 - \frac{\ln c}{\ln \bar{c}_i})}$$

In either case, the EIS increases from $\frac{1}{\gamma}$ to ∞ as consumption goes from 0 to \bar{c}_i . Setting $\bar{c}_i = \infty$ is equivalent to constant EIS utility with EIS $\frac{1}{\gamma}$.²⁶

In their estimations, Blundell et al. (1994) and Attanasio and Browning (1995) each allow \bar{c}_i to depend on observable characteristics of households such as family size, to allow for different types of households to have different elasticities. However, when a household's family size (or other characteristic) changes between two periods, the change in \bar{c}_i also affects the marginal rate of substitution between those two periods directly. So differences in \bar{c}_i may be identified by changes in consumption induced by changes in family size rather than by differences in the sensitivity of consumption growth to interest rates for households of different sizes.

B Data

Business income π_{it}^m : “What was the total income the business earned LAST MONTH after paying all expenses including wages of employees, but not including any income you paid yourself. That is, what were the PROFITS of your business LAST MONTH?” I follow de Mel et al. (2008) in

²⁶These interpretations apply only for $\gamma > 0$. Actually, Blundell et al. (1994) estimates $\gamma < 0$, so that their estimates of \bar{c}_i are instead interpreted as a minimum consumption level. In that case, the EIS *decreases* from ∞ to 0 as consumption goes from \bar{c}_i to ∞ .

interpreting this as income before subtracting any profits reinvested in the business as additional capital. The annual household survey also solicits total household income with this question: “What is your total monthly household income now?”

Capital k_{i1}^m is the sum of current inventories (“At market prices, what is the value you calculate of your current inventories [in stock, products for sale, raw materials, products in production, spare parts, or other such materials currently held at your business]?”) and a running total of physical business assets in five categories, “business tool or utensils,” “machinery,” “furniture and equipment,” “vehicles used in the business,” and “other physical assets of the business (excluding inventories).” For initial stocks of physical business assets, the survey asks, “If you had to replace this, how much would it cost you to purchase one in a similar condition?” For additions or repairs/improvements in subsequent waves, the survey asks, “How much did you spend...”

For consumption, the survey asks how much the family spends in a normal week on food, how much spent last month on other non-durables, and how much spent over the past six months on durables. Each category is broken down into a number of subcategories, with a number of examples listed for each subcategory. Households in the full sample report spending an average of 67% on food and 8% on fuel and other household needs. They spend 5% on health, and 4% each on education, clothing/shoes, and transportation. The remaining 10% is divided among communication, housing and furnishings, recreation, electronics and appliances, and services.

All values are adjusted for inflation to 2005 Sri Lankan rupees and then converted to 2005 US\$ (PPP).

Table 1: Descriptive Statistics

	Mean		Std. Dev.	Min	Median	Max
	Control	Treated				
Retail	0.45	0.52		0	0	1
Manufacturing	0.41	0.35		0	0	1
Services	0.15	0.13		0	0	1
Male Owner	0.45	0.53		0	0	1
Owner's Age	42.4	42.0	11.1	20	42	65
Owner's Years Ed.	9.2	9.2	3.1	0	10	16
Age of Business	9.9	11.5	10.4	1	7	55
Initial Capital	272.2	280.5	252.5	1.4	200	952
Initial Profits	37.2	41.1	34.9	2	30	250
Household Income	88.7	82.4	52.8	12	80	450
Consumption/person	23.6	27.1	18.2	3.6	21.8	129.5
Household Size	4.9	4.9	1.7	1	5	10
Food Share	0.611	0.630	0.164	0.075	0.637	0.963
Rec. / Elec. Share	0.027	0.031	0.064	0.000	0.006	0.456
Savings Rate	0.015	0.036	0.120	-0.174	0.006	1.220
% Change Capital	0.554	0.789	0.917	-1.303	0.575	4.868
% Change Hhld Inc.	0.442	0.578	0.670	-1.191	0.483	2.603
% Change Cons.	0.169	0.120	0.557	-1.401	0.163	2.180

Table 2: Euler Equation Estimates

Parameter	Estimates
EIS ϕ	0.028 (0.014)
Constant term β_c	-0.061 (0.099)
N	129

Estimation is by GMM with robust standard errors reported in parentheses. \hat{r}_i is instrumented with business age, dummies for industry category (retail, manufacturing, or service), gender of the owner, digit span score of owner, and years of education of owner.