Essays in Applied Microeconomics

by

Timothy J. Huegerich

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Economics)

at the

UNIVERSITY OF WISCONSIN-MADISON

2012

Date of final oral examination: 07/27/2012

The dissertation is approved by the following members of the Final Oral Committee:

John Kennan, Professor, Economics
Christopher Taber, Professor, Economics
Rasmus Lentz, Associate Professor, Economics
James R. Walker, Professor, Economics
Laura Schechter, Associate Professor, Agricultural and Applied Economics

Acknowledgment

I have greatly benefited from the guidance of John Kennan. Michael Carter introduced me to the topic area of the first essay. 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.

Table of Contents

Abstract

The average microenterprise in a developing country may have very high marginal returns to capital, 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, they will save slowly in response to high returns despite having a normal rate of time preference. 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 tentatively estimate the EIS of these households to be 0.03, substantially lower than recent estimates around 0.7 for US and UK households. The main implication is that households invest cash grants productively with persistent benefits rather than consuming them away shortly after receiving them.

The second essay examines the impact of labor market frictions on the equilibrium relationship between wages and non-pecuniary job characteristics. It analyzes a model with firm-side heterogeneity in both absolute productivity and marginal cost of providing job amenities. In a search equilibrium in which all firms make equal profits, firms providing greater non-pecuniary job amenities have incentives to offer higher overall utility to workers. Thus workers at lower amenity levels are impeded from advancing to preferred positions with better amenities. This positive correlation between amenity level and utility contrasts with the frictionless framework in which all workers weakly prefer their current job to all available alternatives.

Chapter 1

Interpreting High Returns to Microenterprise Investment

1.1 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. Their finding in de Mel et al. (2012) that the effects of the random grants persist for five years are consistent with a low EIS and difficult to explain otherwise. I estimate the EIS by estimating the Euler equation directly.

The next section presents related literature on measuring the EIS and on microenterprises. Section 1.3 lays out the model of household business investment behavior, while Section 1.4 elaborates on the meaning of the EIS. Section 1.5 describes the data and presents some descriptive statistical analysis, including supportive evidence for the modeling assumptions. Section 1.6 presents the estimation methods and results. Section 1.7 examines some implications of the EIS varying with consumption level, and Section 1.8 concludes.

1.2 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. (2012) using 0.67 for the EIS, can then lead to very poor estimates of policy impacts. Townsend and Ueda (2006) is another recent study for which the EIS assumed for very poor households likely affects the results substantially. In a study of the dynamics of inequality in Thailand, they calibrate the EIS to 1.0 for the baseline value, checking 0.67 for robustness. Modeling poorer households as having much lower EIS while holding other things equal would imply more inequality since they would save less in response to higher return projects, with potentially large implications for comparing the calibrated model with the data.

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. 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. [1]

In each of the papers estimating a variable EIS, as well as in other studies using

[1] Browning 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.1 for details on the functional form for utility they use.

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 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.

In the standard model of consumption and savings presented in the following section, the EIS and relative risk aversion are necessarily reciprocals of one another and cannot be separately specified, but this is only the result of using a convenient and conventional preference specification. In principle, these are distinct aspects of preferences that should be specified separately as in Epstein and Zin (1989). While it is possible for relative risk aversion to be precisely the reciprocal of the EIS, there is no theoretical reason to expect that is so. Rather, it is an empirical question. My data do not provide sufficient information about risk to separately identify risk preferences, but rather my estimation is aimed at measuring the EIS. Thus the results here are not intended to be interpreted as measuring risk aversion, even though they do imply such measures, strictly speaking.

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.[2] 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 (2012a) estimate monthly returns of 5.5% for microenterprises run by women in rural Kenya. However, other recent studies using similar methods such as Karlan et al. (2012) find lower or statistically insignificant returns to capital, showing that high average returns to microenterprise investment are not universal. The aim of this paper is to understand what allows very high marginal returns to persist when they are made available by external factors rather than to explain why high marginal returns are present in some contexts and not in others.

This paper may be the first to systematically examine the self-financing behavior of households in response to these high returns, 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 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 (2012b). Indeed, Dupas and Robinson (2012a) 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.[3] 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.

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

[3]In 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.

1.3 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:

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 1.5.2. 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 current consumption c_{it} or invest it as next month's capital for the household business, $k_{i,t+1}$. The next month's business income increases with capital according to $f_i(k_{i,t+1})$. The household's revenue productivity grows over time at some rate g due to learning-by-doing so that $A_{it} = A_i \exp(gt)$.[4] 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. Overall, business income is given by $\pi = A_{it} f_i(k_{it})\nu_{it}$, where ν_{it} is assumed 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:

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

Assume $f_i(k)$ strictly concave with the marginal return at the initial capital stock at least equal to the discount rate: $A_i f_i'(k_{i0}) \geq \rho$. 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:

$$
u'(\tilde{c}_t) = \frac{1}{1+\rho}\mathrm{E}_t\left[u'(\tilde{c}_{t+1})(1 + \pi'_{i,t+1}(\tilde{k}_{i,t+1})\nu_{i,t+1} - \delta)\right] \tag{1.1}
$$

which just says that, each month, the household chooses the new capital level $\tilde{k}_{i,t+1}$ 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]Households' 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.

1.4 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). In particular, we will consider the Frisch elasticity of substitution[5] between consumption in two consecutive periods and define the EIS $\phi(c_t, c_{t+1})$ to be positive:

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

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

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

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:

$$\frac{d\ln(\frac{\tilde{c}_{t+1}}{\tilde{c}_t})}{d\ln(1+r)}\Big|_{\tilde{c}_t} = \phi(\tilde{c}_{t+1}) \tag{1.2}$$

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

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

the natural log of the return to capital is linear. Taking the anti-derivative of (1.2), and using the fact that consumption growth is zero at $r = \rho$ in this case with no uncertainty:

$$\ln(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}) = \phi(\tilde{c}_{t+1})(\ln(1+r) - \ln(1+\rho)) \tag{1.3}$$

Using first-order Taylor approximations for the natural log functions in equation (1.3), 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:

$$\frac{\tilde{c}_{t+1}}{\tilde{c}_t} - 1 = \phi(\tilde{c}_{t+1})(r - \rho) \tag{1.4}$$

1.4.1 Allowing Uncertainty

Allowing for uncertainty as in the model in Section 1.3 adds additional terms to the equation determining consumption growth but leaves the role of the EIS unchanged. Starting from the Euler Equation (1.1), consider the second-order Taylor expansion of $u'(\tilde{c}_{t+1})$ centered at c_t.[6] I have again substituted $1 + r_{t+1} \equiv 1 + \pi'_{it}(\tilde{k})\nu - \delta$ for readability, but r_{t+1} here is a random variable unknown at time t:

$$u'(\tilde{c}_t) = \frac{1}{1+\rho}\Big(u'(\tilde{c}_t)\mathrm{E}_t\left[1 + r_{t+1}\right] + u''(\tilde{c}_t)\mathrm{E}_t\left[(\tilde{c}_{t+1} - \tilde{c}_t)(1 + r_{t+1})\right] \tag{1.5}$$

$$+\frac{1}{2}u'''(\tilde{c}_t)\mathrm{E}_t\left[(\tilde{c}_{t+1} - \tilde{c}_t)^2(1 + r_{t+1})\right]\Big) \tag{1.6}$$

Rearranging and using the formula for covariance $\mathrm{Cov}\left[X,Y\right] = \mathrm{E}\left[XY\right] - \mathrm{E}\left[X\right]\mathrm{E}\left[Y\right]$ and the first-order Taylor approximation $\ln(1+x) = x$,[7] we have the following formula for expected consumption growth, highlighting its relationship to the expected rate of return:

$$\mathrm{E}_t\left[\ln\left(\frac{\tilde{c}_{t+1}}{\tilde{c}_t}\right)\right] = \phi(\tilde{c}_t)\ln\left(\mathrm{E}_t\left[1 + r_{t+1}\right]\right) + \beta_c \tag{1.7}$$

[6] This derivation follows Dynan (1993), generalizing for an uncertain rate of return.

[7] In 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 $\frac{1+\rho}{\mathrm{E}_t\left[1+r_{t+1}\right]} - 1 \approx \ln\left(\frac{1+\rho}{\mathrm{E}_t\left[1+r_{t+1}\right]}\right)$, which are valid for a small consumption growth rate and $\frac{1+\rho}{\mathrm{E}_t\left[1+r_{t+1}\right]}$ close to 1, respectively.

with β_c representing the sum of three other terms, one representing impatience, one representing the risk of investing capital, and one representing precautionary savings:

$$\begin{aligned}\beta_c &= -\phi(\tilde{c}_t)\ln(1+\rho) - \frac{\mathrm{Cov}_t\left[\frac{\tilde{c}_{t+1}-\tilde{c}_t}{\tilde{c}_t}, 1+r\right]}{\mathrm{E}_t\left[1+r\right]} \\ &+\frac{1}{2}\left(\frac{-c_t u'''(\tilde{c}_t)}{u''(\tilde{c}_t)}\right)\left(\frac{\mathrm{Cov}_t\left[\left(\frac{\tilde{c}_{t+1}-\tilde{c}_t}{\tilde{c}_t}\right)^2, 1+r\right]}{\mathrm{E}_t\left[1+r\right]} + \mathrm{E}_t\left[\left(\frac{\tilde{c}_{t+1}-\tilde{c}_t}{\tilde{c}_t}\right)^2\right]\right) \quad (1.8)\end{aligned}$$

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 prudence, $\frac{-cu'''(c)}{u''(c)}$.[8]

1.5 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 firms[9] 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 ran-

[8]While 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 $1 + 1/\phi$.

[9]Their 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.

dom 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% of purchases were of inventories or raw materials, working capital as opposed to physical capital. Thus, large purchases that households would otherwise need to save for a long time to purchase do not make up the majority of in-kind grants.

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.[10] These estimates of 5% monthly returns correspond to 80% annual returns with compounding.[11] 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

[10] The standard error for the direct treatment effect is 2.2. Standard errors on the estimates of the marginal return to capital are both 2.3.

[11] While 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 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.

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.

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.

1.5.1 Descriptive Statistics

The preliminary sample used below is restricted to those households without missing data for profits or capital in the 9 waves of the initial survey period, a total of 249.[12] The descriptive statistics shown in Table A.1 show that the control group (N=103) and those given grants (N=146) have comparable initial capital, profits, and 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 bud-

[12] In 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.

get. 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 investment rate shown is calculated as the average monthly increase in business capital over a two year 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, desired capital level.

One anomaly in Table A.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.

1.5.2 Evaluating Modeling Assumptions

The data provide evidence of the reasonableness of the first three modeling assumptions listed in Section 1.3: 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 1.6.1 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,[13] 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.

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

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.

1.6 Estimation

1.6.1 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 (1.2). Then, given $\phi(c)$, ρ is identified by the level of the growth rate of capital.[14]

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.

Although this argument establishes the identification of the full model, I currently only estimate the EIS in what follows. In fact, it is an advantage of the Euler equation approach described below that it is robust to alternate assumptions about other assets or borrowing opportunities, as well as to how labor supply to the business affects outside income.

[14]Note that a model in which the rate of time preference were also allowed to depend on c, $\rho(c)$, is identified just as well. For households with the same consumption level, c, but different rates of return, $\phi(c)$ and $\rho(c)$ are separately identified as above. Then the same follows at a different consumption level c', so that the changes of both $\phi(c)$ and $\rho(c)$ with consumption level may be identified. Given the small sample size of my data, I do not expect to be able to estimate both a variable ϕ and a variable ρ. I focus on a variable ϕ because of the focus of this paper on the EIS and because I find the intuition for a variable ϕ more compelling than for a variable ρ.

1.6.2 Reduced Form Estimation

The EIS is typically estimated by using the Euler equation (1.1) as an orthogonality condition for GMM. Assuming rational expectations means that, for the realized values of the relevant variables, the equation will hold on average over time. Given the short panel for consumption available in the data, I need to impose an additional assumption to use cross-sectional variation rather than longitudinal variation to estimate the EIS in this section. I need to assume that expectational errors in (1.1) average to zero across households in a given period. This is consistent with the model in Section 1.3 in which the shocks are independently distributed across households, but it would be violated by any correlation in shocks across households.[15]

Though we cannot use the approximated Euler equation (1.7) directly to get an estimate for the EIS by semi-parametric regression, we can run an analogous regression that allows us to plausibly bound the estimate of the EIS from above. Recall we are 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 (1.7) can be written:

$$\mathrm{E}_t\left[\ln\left(\frac{\tilde{c}_{it+1}}{\tilde{c}_{it}}\right)\right] = \phi\ln\left(\mathrm{E}_t\left[1+r_{it+1}\right]\right) + \beta_c$$

where the term β_c collects all terms other than the expected rate of return, detailed in (1.8) above. Since I only have annual data for (monthly) consumption,[16] consider the analogous expression for the annual growth rate of monthly consumption:

$$\mathrm{E}_t\left[\ln\left(\frac{\tilde{c}_{it+12}}{\tilde{c}_{it}}\right)\right] = \phi\ln\left(\mathrm{E}_t\left[\prod_{s=t+1}^{t+12}(1+r_{is})\right]\right) + \beta_{c12} \qquad (1.9)$$

But I only have one estimate of the monthly return to capital for each household in this exercise, $\hat{r}_i$ discussed below, so the analogue is to substitute the annualized form $(1+\hat{r}_i)^{12}$ for $\prod_{s=t+1}^{t+12}(1+r_{is})$. What I actually estimate is the following model, where Z_{it} are instruments for the rate of return discussed below:

[15] See Attanasio and Low (2004) and references therein for further discussion of this assumption, which is generally avoided by using longitudinal rather than cross-sectional asymptotics.

[16] Only non-durable consumption is used because it corresponds to the model assumption that utility is time separable.

$$\ln\left(\frac{\tilde{c}_{it+12}}{\tilde{c}_{it}}\right) = \beta_1 \mathrm{E}\left[12\ln(1+r_i)|Z_{it}\right] + \beta_0 + \epsilon_{it+12} \quad (1.10)$$

Comparing to (1.9), the household's expected rate of return is replaced by the rate of return conditional on observed characteristics Z_{it}. Also note that the log of the expectation of the annual rate of return has been replaced with the expectation of the log, which is strictly less by Jensen's inequality. This will bias estimates of β_1 upward from the true value of the EIS ϕ so that these estimates could be viewed as an upper bound on the EIS.[17] The error term ϵ_{it+12} includes expectational errors and measurement error in consumption. It also includes any differences across households in the leftovers term β_{c12} detailed in equation (1.8), which includes impatience and the second-order terms representing risk aversion and precautionary savings, or higher order terms not included in the log-linear approximation. In particular, then, consistency requires that these leftover terms are uncorrelated with the instruments, as discussed by Attanasio and Low (2004), who use with Monte Carlo exercises to examine the magnitude of the bias induced by the failure of this assumption.

I calculate the realizations of returns to capital as 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 $\hat{r}_i = \frac{y^{tot}_{i12} - y^{tot}_{i0}}{k_{i13} - k_{i1}}$. I limit the sample to grant recipient households, for which there is substantial exogenous variation in capital to make this sort of estimate plausible. Yet this rough measure attributes all changes in income to the difference in capital. If incomes rise over the year for other reasons such as increased experience, this will be an overestimate of the returns to capital and the estimate of β_1 will be biased downward. Using total household income is preferable to using only business income to calculate returns because it includes any decreases in outside income that might come from increasing labor supply to the household business along with increasing capital. However, I also report results for the corresponding measure of returns using business profits only, $\hat{r}_i = \frac{\pi_{i12} - \pi_{i0}}{k_{i13} - k_{i1}}$.

I instrument for the realized rates of return using demographic variables related to ability, all of which are known in the initial period, so that β_1 is the coefficient on the expected rate of return. Instrumenting also helps correct for measurement error in the rate of return estimates. Specifically, I use as instruments the following variables found

[17] This may explain why the estimates in Table 3 of Attanasio and Low (2004) are 0.04 to 0.08 higher than the true value of 0.67 in their Monte Carlo exercise.

by de Mel et al. (2008) to have a significant effect on returns to capital: gender, years of education of the owner, and owner's score on a digit span intelligence test, as well as the age of the business and 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 consumption or any other factor that affects consumption growth, except through the expected return to capital. In particular, if values of Z_{it} associated with higher rates of return are associated with lower consumption growth for other reasons, such as greater volatility in consumption due to a relatively tight liquidity constraint, the estimate of β_1 will be biased downward, making this estimate of the EIS appear lower than the true value. [18] Thus more work is needed to assess the likelihood and plausible size of this potential bias that casts the main result below in question.

The main results are shown in Table A.2. For the preferred specification using total household income, the point estimate $\hat{\beta}_1$ 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. For the specification using only profits to calculate household-specific returns to capital, however, the point estimate is negative, though it is not statistically different from the point estimate using total income due to the relative imprecision of the estimate. Results from OLS regressions for both measures of returns to capital are also shown. In either case there is no statistically significant correlation between consumption growth and the measured returns to capital. This is consistent with either measurement error in the returns to capital or simply the model implication that consumption growth is less responsive to shocks in the realized returns to capital than to differences in expected returns.

Yet the results of the first stage regression of the log returns to capital on the instruments, shown in Table A.3, call into question the IV estimates of β_1. For the total income measure of returns, the F test says the instruments overall are significant at the 10% level, though only the dummy for manufacturing businesses and the owner's education are individually significant. It is concerning, however, that the coefficient on education is negative, contrary to the results of de Mel et al. (2008) who found that education is positively correlated with returns to capital measured more carefully using the randomized assigment of their field experiment. They also found significantly higher returns for males, as opposed to the negative point estimate here, and no significant

[18] On the other hand, if values of Z_{it} associated with higher rates of return are also associated with higher consumption growth for other reasons, such as higher ability owners having less volatility in income, the estimate of β_1 will be biased upward.

differences in returns among industry sectors, contrary to the significantly higher returns for manufacturing over retail here. As for the first stage results for the measure of returns using profits only, none of the instruments have significant coefficients and the F test does not reject that the instruments are unrelated to that measure of returns. Moreover, the point estimates are often significantly different from the corresponding coefficients for the total income measure, so they do not provide any reassurance that the latter are sensible.

Overall, the instrumental variables estimate using total income to measure returns to capital yields a plausible estimate of the EIS, but I regard this exercise as merely suggestive given the plausible suspicion of downward bias and lack of robustness to using only profits to measure returns.

1.7 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 A.2. Abstracting away from risk, Equation 1.4 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 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

[19] From Table A.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.

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.

1.7.1 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 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 have different preferences, each with constant EIS regardless of consumption level.[23] These two possibilities have

[20] A review of other utility functions allowing varying EIS in the literature and why they are less appropriate for this exercise is in Appendix A.1.

[21] This 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.

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

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

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 equation (1.3), $\frac{\tilde{c}_{t+1}}{\tilde{c}_t} - 1 = \phi(\tilde{c}_t)(r - \rho)$, 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:

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

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.

1.8 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 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 microenter-

[24] Numerical simulations confirm this is a good approximation for constant EIS for capital levels sufficiently far from the 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.

prises 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.

Chapter 2

The Implications of Search Frictions for Measuring Workers' Preferences for Job Characteristics

2.1 Introduction

In a labor market with perfect information, jobs with worse characteristics from the perspective of the worker (working conditions, safety risks, and so on) must have a fully compensating wage differential to attract any workers, as first theorized by Adam Smith. Rosen (1974) formalized the relationship between job amenities (good characteristics) and wages, showing that the wage is a function of job amenities, the slope of which must at each point equal the worker's marginal rate of substitution between consumption and job amenity and the firm's technical rate of substitution. This wage function is commonly referred to as a hedonic wage and can be understood as a fixed base wage minus an implicit price charged to the worker for job amenities.

In a survey article, Rosen (1986) acknowledges that the classic theory of compensating wage differentials assumes perfect information while at the same time noting that search and information costs "sustain significant wage variability among measurably identical jobs and workers" (fn. 2, p. 643). He argues that a model with search frictions will nevertheless tend toward the same relationship between amenities and wages in equilibrium as the perfect information outcome.

> The search for a job and investment in information is in many ways a search for the type of allocations described here. Hence the theory must be considered as one of longer run tendencies and of equilibrium behavior in the steady state of a more complex dynamic process. (p. 643)

However, recent papers have shown that search frictions can plausibly lead to labor market equilibria with very different configurations of wages and amenities, in which the slope of the wage-amenity relationship need not equal the slope of the worker's indifference curve. These papers adapt equilibrium search models, such as that of Burdett and Mortensen (1998), in which there is wage dispersion in equilibrium even with homogeneous workers. Firms face a tradeoff between the cost of offering a higher wage to workers and the difficulty of finding and retaining workers with a low wage. In equilibrium firms each maximize profit at different wage choices. For each, the gain from lowering their wage offer would be exactly offset by the resulting increase in costs of hiring and turnover. It is straightforward conceptually to extend this idea to multidimensional compensation, a model in which the worker values characteristics of jobs as well as the wage. Workers search for jobs based on the total utility they would get, accounting for both the wage and amenities. Their optimal acceptance strategies will involve reservation utility levels rather than reservation wages, for instance. And by the same reasoning as in the single-dimensional compensation case, there will be a dispersion of utilities offered to workers in equilibrium.

The contribution of this paper is to characterize the search equilibrium when firms are heterogeneous in both productivity and cost-efficiency in providing an amenity. I also show that if free entry adjusts the demand for the output goods of different firm types, and thus their revenue productivity, such that all firms earn equal expected profits, then whether the slope of the wage-amenity relationship will be greater or less than the workers' indifference curve depends on whether the amenity is a normal or inferior good, or neither. In particular, if the amenity is a normal good and different firm types earn equal expected profits, the firms offering higher amenity levels will also be offering higher overall utility levels, causing a downward bias in estimates of workers' willingness to pay for the amenity by traditional hedonic wage methods.

In the next section, I discuss the existing literature on compensating wage differentials with search frictions and explain my contribution in more detail. In Section 2.3, I lay out the search framework that gives rise to equilibrium utility dispersion. In Section 2.4, I extend the model to allow for general firms types and prove the main result of the paper, the utility ordering of firm types. I also present some representative examples. Section 2.5 develops the model for a continuous distribution of firm types. Then Section 2.6 presents the implications of assuming that firms tend to earn equal expected profits across types, and Section 2.7 concludes.

2.2 Literature Review

Lang and Majumdar (2004) consider the case of homogenous firms, as well as identical workers. In a frictionless labor market, the equilibrium would consist of a single wage-amenity pair at a point of tangency between the workers' indifference curves and firms' isoprofit curves. With utility dispersion, however, there are a range of wage-amenity pairs in equilibrium with a positive relationship between wages and amenities. This is because each firm will provide utility to workers in the most cost effective way, trading off between using higher wages or better amenities to make their jobs attractive. Under the standard assumptions (job amenities are normal goods for workers and firms have increasing marginal costs of providing amenities),[1] firms offering higher utility to workers will thus offer both a higher wage and better amenities than the lower utility firms, leading to the positive relationship in equilibrium, the opposite of a compensating wage differential. This is analogous to the well known potential for bias in estimating hedonic wage functions that can result from unobserved heterogeneity in worker ability, which is cited by Rosen (1986) as "the fundamental reason why low paying jobs tend to be the 'worst' jobs" (p. 671). However, in this model, workers are of identical ability, and the only difference between workers is in how lucky they are when searching for jobs.

Hwang et al. (1998) explore the implications of firm heterogeneity for the equilibrium distribution of wages and amenities. Now, in a frictionless labor market with homogeneous workers and multiple types of firms offering different levels of an amenity, the hedonic wage function will trace out an indifference curve of the workers. This is a best case scenario from the perspective of estimating the preferences of the workers for job characteristics. In an equilibrium with utility dispersion, however, wage-amenity pairs are spread among a range of the workers' indifference curves, rather than tracing out a single utility level. In the special case considered by Hwang et al. (1998), the

[1] I follow Lang and Majumdar (2004) in their notion of the amenity as a normal good throughout. It means that, hypothetically, if the worker had a fixed wealth each moment that she could use to either purchase the amenity at a constant price per unit or allocate toward the composite consumption good (i.e. the wage), she would purchase more of the amenity if she had more wealth. See the Appendix for a formal discussion.

firm types that offer the higher levels of amenity also offer discretely higher utility jobs overall, which means that wage differentials between jobs are not fully compensating. The slope of the relationship between wages and amenity will tend toward zero, less steep than the workers' indifference curves, and may even become positive.

Concern about search frictions as a potential source of significant bias when estimating worker preferences using the traditional hedonic wage model led Gronberg and Reed (1994) to propose an alternative approach to estimating worker preferences based on job duration data rather than the relationship between wages and amenities alone. Although compensating wage differentials can be distorted in a market with search frictions, the on-the-job search model itself predicts a systematic relationship between the utility of a job and how long a worker will stay at the job that can be used instead to estimate worker preferences. Gronberg and Reed (1994) apply their method to working conditions such as whether the job requires kneeling or stooping. It has also been applied to commuting time by van Ommeren et al. (2000), night shift work by Manning (2003), and risk of injury by Dale-Olsen (2006). In each case, the workers' marginal willingness to pay (MWP) for job amenities are estimated to be higher by the job durations method than by the traditional hedonic wage method, which has been interpreted as confirmation of the claim of Hwang et al. (1998) that search frictions always cause the observed slope of the hedonic wage relationship to be less than the true slope of workers' indifference curves, that is, their MWP. However, it is as yet unclear whether these estimates should be considered more accurate than the traditional method,[2] as various objections can be raised to the basic job durations method.[3]

Another approach to estimating worker preferences that is not new but has been given renewed attention in light of the challenge of search frictions to traditional hedonic wage regressions is the method of estimating preferences of workers by looking directly at their choices between jobs, as seen in their voluntary job transitions. Each instance of a worker voluntarily changing jobs gives unambiguous information about the worker's preferences between them. Bonhomme and Jolivet (2009) present an innovative recent application of this approach based on moving costs rather than search frictions. It may be the first empirical analysis of the difference between the performance of the cross-sectional hedonic wage approach and an alternative method for estimating

[2]For example, Manning (2003) reports an estimate using the job durations method that implies night shift workers in the UK would be willing to pay over 90% of their salary to switch to a day shift job!

[3]For example, the job durations method relies on the assumption of random search. If, say, securing a high utility job makes it easier to find other high utility jobs due to a signalling effect, such jobs could have a systematically shorter tenure, contrary to the identifying assumption that average tenure length is monotonically increasing with job utility.

preferences within a single modelling framework.

In this paper, I adapt the models of Hwang et al. (1998) and Lang and Majumdar (2004) and allow for a more general form of firm heterogeneity, retaining the assumption of homogeneous workers for clarity and tractability.[4]There is a well-developed literature for frictionless markets on the identification and estimation of preferences and technology with heterogeneity in both,[5] but there has been less work on identification in the search frictions case. Thus I consider the most basic case, identifying worker preferences only for homogeneous workers and heterogeneous firms. The theoretical contribution here is characterizing in full generality which firm types will offer higher utility jobs in equilibrium. I also extend the model to allow for a continuous distribution of firm types, following the model of Burdett and Mortensen (1998) for the single-dimensional compensation case. However, in the case of multi-dimensional compensation, allowing a continuum of firm types makes for a qualitatively different and substantially more intuitive distribution of wages and amenities. With the resulting model, the full range of possible wage-amenity distributions comes into view, and it is possible to characterize in a general way how the distribution of firm heterogeneity determines the equilibrium relationship between wages and amenities, including the conditions under which the special cases explored by Hwang et al. (1998) and Lang and Majumdar (2004), respectively, would arise.

With the range of theoretical possibilities in view, the question arises of what sort of relationship between wages and amenities we should expect to actually see. I show that if we conjecture that different types of firms earn equal profits in expectation, the special case of Hwang et al. (1998) in which higher amenity jobs tend to offer higher utility is indeed the result. If this reasoning is correct, it explains the serious bias in the traditional hedonic wage methodology that empirical work using the job durations method suggests. It is also worth noting that this hypothesis that higher amenity jobs are more desirable than lower amenity jobs distinguishes the frictions theory from the unobserved worker productivity heterogeneity explanation for more positive relationships between amenities and wages. In order for the latter to explain such a correlation between overall utility and amenity level, it would be necessary to posit a seemingly ad hoc correlation between unobserved ability and preferences for the amenity.

[4]Lang and Majumdar (2004) consider a case with two types of worker, differing in their taste for a binary amenity, and homogeneous firms. They characterize the equilibrium for both the case in which firms can observe the worker type and make their job offers contingent on it and the case in which firms must blindly offer the same jobs to both types.

[5]See Heckman et al. (2010) and references therein.

2.3 Search Framework

The following equilibrium search framework combines elements taken from Hwang et al. (1998) and Lang and Majumdar (2004). Jobs are search goods for the workers, characterized by a wage w and a scalar amenity level x, interpreted here as a good. Firms post job offers specifying w and x, seeking to maximize their expected profit per worker contacted.[6] Workers search for jobs in an undirected way both when unemployed and on-the-job. Suppose there is an infinite number, a continuum, of both workers and firms.

The flow value for a worker of being employed at a job is given by $v(w, x)$, defining each worker's preferences over jobs. Assume v is $\mathbb{C}^2$ and strictly quasiconcave. The value of being unemployed is b. Job offers arrive for each worker at rate λ_U in unemployment or λ_E when employed. Jobs are lost at a constant rate δ.

To determine their optimal search behavior, workers need only know the distribution of job offer utilities, $F(v)$, which is assumed to be common knowledge. Hwang et al. (1998) show that the optimal job acceptance strategy[7] when unemployed in this setting is to accept all jobs with value greater than

$$v^* = b + (K_U - K_E) \int_{v^*}^{\overline{v}} \frac{1 - F(v)}{1 + K_E[1 - F(v)]} dv \tag{2.1}$$

where $K_U = \lambda_U/\delta$ and $K_E = \lambda_E/\delta$ may be interpreted as quantifying the competitiveness of the labor market and inversely related to the magnitude of search frictions for unemployed and on-the-job search, respectively. For employed workers, the optimal job acceptance strategy is simply to accept any offer better than their current one.

Let U be the steady-state level of unemployment, and let $G(v)$ be the steady-state proportion of employed workers with jobs worth v or less. First, we will see that no firm has an incentive to offer a job with value less than v^* in equilibrium, so we simplify

[6]Nothing about this analysis limits its applicability to labor markets. The extension to other hedonic markets is straightforward, with consumers analogous to workers and suppliers to employers. The key assumption is that suppliers fix the characteristics of their good ex ante, but may sell any quantity that the market bears. This is one of three cases discussed by Rosen (1974) and seems the most amenable to being modeled in an undirected, stationary search framework. It is important to understand the distinction between this case and that studied recently by Heckman et al. (2010) in which each firm only demands one worker. In this framework, firms of different productivities can survive together in equilibrium, but all offer wages according to the equilibrium hedonic wage function–there is no wage dispersion with equally productive workers.

[7]I follow Hwang et al. (1998) in presenting the results of the model in which there is no discounting, or $r = 0$. Allowing for $r > 0$ is straightforward, but it makes the formulas somewhat more unweildy and does not affect the qualitative features of the model.

the exposition by taking $F(v^*) = 0$. Then the flow out of unemployment is simply $U\lambda_U$ and the flow into unemployment is $(1-U)\delta$. Setting these equal, we find that in steady state

$$U = \frac{1}{1+K_U} \tag{2.2}$$

Similarly, the flow of workers into jobs of value v or lower, $U\lambda_U F(v)$, and the flow of workers out of such jobs by on-the-job search and job loss is $G(v)(1-U)\left(\lambda_E\left[1-F(v)\right]+\delta\right)$. Thus in the steady state

$$G(v) = \frac{F(v)}{1+K_E\left[1-F\left(v\right)\right]}. \tag{2.3}$$

From the firms' perspective, both the probability that a randomly contacted worker will accept their job offer and the expected amount of time the worker will stay with them depend on the utility of the job they offer. The probability that a random worker will accept their offer of utility v is simply the proportion of workers that would prefer the job to their current state $U + (1-U)G(v)$. Once a worker accepts the job, she will leave the job at rate: $\delta + \lambda_E\left[1-F(v)\right]$ due either to an exogenous job separation or to finding a better job, so the expected duration of employment is the inverse. The product of the probability of a contacted worker accepting and the expected duration of employment tells us the expected duration of employment from contacting a worker:

$$\begin{aligned} m(v) &= \frac{U+(1-U)G(v)}{\delta+\lambda_E\left[1-F\left(v\right)\right]} \\ &= \frac{1}{\left(1+K_E\left[1-F\left(v\right)\right]\right)^2}\frac{1+K_E}{1+K_U} \end{aligned} \tag{2.4}$$

where the second expression follows from (2.2) and (2.3).[8] Intuitively, $m'(v) > 0$ reflects the firms' incentives to offer higher utility jobs, and the dependence of $m(v)$ upon $F(v)$ shows how the equilibrium job offer distribution shapes these incentives.

Firms each have a production technology that is constant returns to scale in the labor input, with revenue product per worker given by $\phi_j(x)$, decreasing in x, with j indexing the firm type. Assume ϕ is $\mathbb{C}^2$ with $\phi''_{xx} < 0$, reflecting an increasing marginal cost of providing the amenity. Then, the expected profit of a firm of type j is the product of the flow profit $\phi_j(x) - w$ and the expected duration of employment resulting from a

[8]This presentation of the firm's problem is adapted from Section 2.2 of Mortensen (2003).

worker contact:

$$\pi_j = [\phi_j(x) - w]\, m(v(w,x)). \tag{2.5}$$

Conditional on v, w and x only affect the flow profit per worker. Thus, the firm's problem may be decomposed into two steps: (1) choose v, and (2) choose w and x to maximize the flow profit per worker given v. First consider the latter. Let

$$\begin{aligned} \widetilde{\pi}_j(v) = \quad & \max\nolimits_{w,x}\ \phi_j(x) - w \\ & \text{s.t. } v(w,x) = v \end{aligned} \tag{2.6}$$

Thus we see that the firm will choose (w,x) so that

$$\phi_j'(x) = -\frac{v_x(w,x)}{v_w(w,x)}. \tag{2.7}$$

That is, the firm chooses the amenity level x so that the marginal revenue product of x is equal to the marginal rate of substitution of the worker. The assumptions on v and ϕ are sufficient to guarantee the existence of a unique solution for a given v. Thus $\widetilde{\pi}_j(v)$ is the maximum flow profit a firm can make from employing a worker at value v. That is, $\pi_j = \widetilde{\pi}_j(v) \cdot m(v)$. Notice $\widetilde{\pi}_j'(v) < 0$ reflects the firms incentives to lower costs by offering lower utility.

Now the unique equilibrium offer distribution $F(v)$ may be derived in the following way. First, it must satisfy the following conditions, proven by Hwang et al. (1998):

(a) No positive mass of firms offers the same job value v.

(b) There are no gaps in the support of the job value distribution $F(v)$.

(c) The lowest offered job bundle has value $v = v^*$.

If any of these were not true, there would be an opportunity for a firm to increase profit per worker by a discrete jump with little or no cost. If (a) were not true, one of the firms in the positive mass at v could increase the value of their offer slightly for a discrete jump in $m(v)$. If (b) or (c) were not true, offering a lower-valued bundle in the gap would increase $\widetilde{\pi}_j(v)$ while maintaining the same relative position in the offer distribution, that is, with the same offer acceptance probability and average employment duration, the same m.

Now consider the case in which firms are homogeneous. In order for these identical firms to make different offers in equilibrium, each must make equal profit. Thus, we can derive the offer distribution from the condition that profits from a job of value v

must equal the profit from offering $\underline{v} = v^*$:

$$\begin{aligned} \widetilde{\pi}(v)m(v) &= \widetilde{\pi}(v^*)m(v^*) \\ \widetilde{\pi}(v)\frac{(1+K_E)/(1+K_U)}{(1+K_E\left[1-F\left(v\right)\right])^2} &= \widetilde{\pi}(v^*)\frac{(1+K_E)/(1+K_U)}{(1+K_E\left[1-F\left(v^*\right)\right])^2} \end{aligned}$$

Since $F(v^*) = 0$, we have

$$F(v) = \frac{1+K_E}{K_E} - \frac{1+K_E}{K_E}\left(\frac{\widetilde{\pi}(v)}{\widetilde{\pi}(v^*)}\right)^{\frac{1}{2}} \tag{2.8}$$

And the highest job value $\overline{v}$ is determined by $F(\overline{v}) = 1$, which yields:

$$\widetilde{\pi}(\overline{v}) = \frac{1}{(1+K_E)^2}\widetilde{\pi}(v^*) \tag{2.9}$$

An example of the relationship between wages and amenity levels that results can be seen in Figure 1a. As Lang and Majumdar (2004) argue will typically be the case for homogeneous firms, wages and amenities are positively related. The example is described in more detail in the following section.

2.4 Discrete Firm Types

Suppose there are n types of firm, characterized by $\phi_j(v)$, $j = 1, \ldots, n$, and let γ_j denote the proportion of firms of type j or lower. A reasonable single crossing condition, defined with respect to the given $v(w, x)$ specification of workers' preferences, will facilitate our solution:

$$\frac{\widetilde{\pi}_i'(v)}{\widetilde{\pi}_i(v)} < \frac{\widetilde{\pi}_j'(v)}{\widetilde{\pi}_j(v)} \quad \forall\, i < j,\ \forall\, v \tag{2.10}$$

This says that for firms of higher type j, the proportional, or percentage, decrease in their profit per worker from increasing the worker's utility v is everywhere higher (or less negative) than for lower type firms. In other words, lower type firms increase their profit by a greater percentage than higher type firms by decreasing the utility given to the worker. The ordering depends on a combination of the firms' productivity in producing the output good and their marginal cost of providing the job amenity for the worker, but it is misleading to think of the ordering as ranking them according to a composite productivity measure. Instead, the ordering of types describes their relative incentives to offer high utility jobs and attract workers more easily versus offering lower utility jobs and increasing the profit flow, as I will now show formally.

The following is a generalization of a proposition in (Hwang et al., 1998) that allows a complete characterization of the steady state equilibrium for general firm heterogeneity characterized by (2.10):

Proposition 1

Higher type firms offer jobs of at least as much utility as lower type firms. That is, for $i < j$, if v_i and v_j are the values of jobs offered by firms of types i and j, respectively, then $v_i \leq v_j$.

Proof. First, notice that

$$\int_{v_i}^{v_j} \frac{\widetilde{\pi}'(v)}{\widetilde{\pi}(v)} dv = \ln \frac{\widetilde{\pi}(v_j)}{\widetilde{\pi}(v_i)}, \tag{2.11}$$

which, along with the condition (2.10), implies that

$$v_i \leq v_j \Leftrightarrow \frac{\widetilde{\pi}_i(v_j)}{\widetilde{\pi}_i(v_i)} \leq \frac{\widetilde{\pi}_j(v_j)}{\widetilde{\pi}_j(v_i)} \tag{2.12}$$

Now, if a firm of type i chooses to offer value v, it must be that

$$\widetilde{\pi}_i(v_i) \cdot m(v_i) \quad \geq \quad \widetilde{\pi}_i(v_j) \cdot m(v_j) \tag{2.13}$$

so we have

$$\frac{\widetilde{\pi}_i(v_j)}{\widetilde{\pi}_i(v_i)} \leq \frac{m(v_i)}{m(v_j)} \tag{2.14}$$

Likewise

$$\frac{m(v_i)}{m(v_j)} \leq \frac{\widetilde{\pi}_j(v_j)}{\widetilde{\pi}_j(v_i)} \tag{2.15}$$

Thus

$$\frac{\widetilde{\pi}_i(v_j)}{\widetilde{\pi}_i(v_i)} \leq \frac{\widetilde{\pi}_j(v_j)}{\widetilde{\pi}_j(v_i)} \tag{2.16}$$

But then by (2.12), we must have $v_i \leq v_j$. □

Thus higher types will always offer higher utility jobs in equilibrium because they have less to lose from doing so. In other words, lower types have more to gain from offering lower utility jobs.

This proposition, along with conditions (a)-(c) above, implies the following distribution of wage offers. Assume that all types are capable of making positive profit, that is, that $\widetilde{\pi}_j(v^*) \geq 0$ for all j. Then the firms in each type j will offer jobs of value $v \in [\underline{v}_j, \overline{v}_j]$, with $\underline{v}_1 = v^*$, and $\overline{v}_j = \underline{v}_{j+1}$ for all $j < n$. These intervals and the distribution of offers can be determined in the following way.

We derive the distribution of offers by firms of type $j = 1$ in the same way as we did for the case of all homogeneous firms. The profits from a job offer of value $v \in [\underline{v}_1, \overline{v}_1]$ must be equal to the profit from offering $\underline{v}_1 = v^*$:

$$\widetilde{\pi}_1(v)m(v) \quad = \quad \widetilde{\pi}_1(v^*)m(v^*)$$

So we have

$$F(v) = \frac{1+K_E}{K_E} - \frac{1+K_E}{K_E}\left(\frac{\widetilde{\pi}_1(v)}{\widetilde{\pi}_1(v^*)}\right)^{\frac{1}{2}} \tag{2.17}$$

for $v \in [\underline{v}_1, \overline{v}_1]$. And the highest job value $\overline{v}_1$ is determined by $F(\overline{v}_1) = \gamma_1$, which yields:

$$\widetilde{\pi}_1(\overline{v}_1) = \frac{(1+K_E(1-\gamma_1))^2}{(1+K_E)^2}\widetilde{\pi}_1(v^*) \tag{2.18}$$

Since we have $\underline{v}_j = \overline{v}_{j-1}$, we can repeat this argument for $j = 2, \ldots, n$ and find that for $v \in [\underline{v}_j, \overline{v}_j]$,

$$F(v) = \frac{1+K_E}{K_E} - \frac{1+K_E(1-\gamma_{j-1})}{K_E}\left(\frac{\widetilde{\pi}_j(v)}{\widetilde{\pi}_j(\underline{v}_j)}\right)^{\frac{1}{2}} \tag{2.19}$$

so that $\overline{v}_j$ is determined by

$$\widetilde{\pi}_j(\overline{v}_j) = \frac{(1+K_E(1-\gamma_j))^2}{(1+K_E(1-\gamma_{j-1}))^2}\widetilde{\pi}_j(\underline{v}_j) \tag{2.20}$$

Finally, then, we have a complete expression for $F(v)$ in terms of v^* which we can put together with (2.1) to find an expression for v^* in terms of model parameters only, for which it can be shown that a unique solution exists. In general, the resulting expression must be solved numerically, but since we are interested here in the potential wage-amenity distributions that can result from this model, we will simply take v^* as given, with the understanding that it is determined by b, K_U, K_E, and the distribution of firm technologies. It is particularly helpful to note that b and K_U *only* affect the wage-amenity distribution through their influence on v^*. Thus the shape of the wage-amenity distribution depends only on the resulting v^*, along with K_E, the firms' technologies, and the workers' preferences. Further, it can be seen that K_E has little effect on the shape of the wage-amenity relationship, as determined by $F(v)$, but instead the main impact of K_E, in addition to affecting v^*, is on the distribution of workers in steady state within the utility dispersion $G(v)$.[9]

2.4.1 Illustrative Examples

As an example, specify worker preferences by $v(w,x) = wx^\alpha$ and firm technology by $\phi_j(x) = \rho_j - \frac{\sigma_j}{2}x^2$, with some $\{(\rho_1,\sigma_1),\ldots,(\rho_n,\sigma_n)\}$. Roughly speaking, types with higher ρ have higher productivity in producing the outside good and types with lower σ have lower costs (higher productivity) in producing the job amenity and will tend to offer more of it. Then (2.7) says that each firm will offer some wage-amenity pair that satisfies $w = \frac{\sigma_j}{\alpha}x^2$ depending on its type j, and we have $\widetilde{\pi}_j(v) = \rho_j - \left(1+\frac{\alpha}{2}\right)\left(\frac{\sigma_j}{\alpha}\right)^{\frac{\alpha}{\alpha+2}} v^{\frac{2}{\alpha+2}}$. It can be shown that the single crossing condition (2.10) is satisfied for a given α if

$$\frac{\rho_i}{\sigma_i^{\frac{\alpha}{\alpha+2}}} < \frac{\rho_j}{\sigma_j^{\frac{\alpha}{\alpha+2}}} \quad \forall\, i < j$$

Figure 1 shows samples of the equilibrium wage-amenity distribution for different distributions of firm heterogeneity. For comparison, $\alpha = 2$, $v^* = 0.7$, and $K_E = 15$ are held constant in all four cases, and the firm technologies are chosen so that the highest utility offered by any firm is also roughly the same ($\overline{v} = 1$) across cases. Thus, the utility dispersion $F(v)$ is roughly the same in each, highlighting the effects of different distributions of firm heterogeneity.

Figure 1a is a case with just one type of firm, $\rho = 2\sqrt{2}$ and $\sigma = 2$. All firms have identical technology, and they all offer wage-amenity pairs with $w = \frac{2}{2}x^2 = x^2$ in a certain range. They all receive the same profits on average, with the low-v firms enjoying higher profit flows from worker but finding it more costly to find and keep workers, while the high-v firms have lower profit flows from workers but have an easy time hiring and lower turnover. Notice that the workers' marginal rate of substitution, what we refer to as the marginal willingness to pay for the amenity (MWP), is given by

$$MWP = \frac{v_x}{v_w} = \frac{\alpha w}{x} = 2x$$

so it is roughly in the range of 1.8 (at the lowest utility jobs) to 2.0(at the highest). In a standard hedonic wage regression, the MWP would be estimated as the negative of the slope of the hedonic wage function. If that method were naively applied to this example,

[9]This explanation is intended to correct the misleading diagrams in (Hwang et al., 1998) that mean to show how the K parameters affect the shape of the resulting wage-amenity distribution, but do not explain that the K parameters only matter for the most part in how they affect the difference between v^* and $\overline{v}$.

a negative estimate of around -2 would be the result, exactly the negative of the true MWP. However, it is not too surprising to find that the MWP could not be identified in this case if we had data on the wage-amenity pairs only. (Recall $w = \frac{\sigma}{\alpha}x^2$ so that all we could identify is $\frac{\sigma}{\alpha}$.) This is because this case is analogous to a perfect information hedonic wage model in which there is only one data point, at the particular (w, x) pair where the homogenous workers and firms all locate.

This is an example of the homogeneous firms case thoroughly examined by Lang and Majumdar (2004), who show that if the amenity is a normal good as it is for the Cobb-Douglass preferences here, the relationship between wages and amenity will be positive.

It is worth noting that a wage-amenity relationship similar to that displayed in Figure 1a could also result from a situation with heterogenous firms but little variation in their cost of producing the amenity σ relative to the magnitude of the utility dispersion (i.e. $\overline{v} - v^*$). In all such cases, a naive hedonic wage regression will result in a MWP estimate of the wrong sign.

Figure 1b is an example of the special case of firm heterogeneity explored by Hwang et al. (1998). Here, the firm types that have lower marginal cost of providing the amenity, σ_j, and thus offer higher levels of x are the "higher types" that offer higher valued jobs than the firms that provide lower levels of x.[10] So instead of tracing out one of the workers' indifference curves as would be expected in the perfect information hedonic wage model without utility dispersion, the relationship between x and w traced out has a systematically more positive slope than the indifference curves. This effect can be so severe that the overall relationship between wages and amenity is a positive one, but for a different reason than in Figure 1a.

Figure 1c is the opposite case, in which the firms that are less efficient at providing x have higher productivity per worker ρ so that they are the firms offering the highest value jobs. And we see that, for this example, the relationship traced out between x and w has a systematically more negative slope than the indifference curves. Ex ante, it is unclear why this should be any more or less likely than the case in Figure 1b.

Figure 1d is perhaps a more typical case, in which there is no systematic ordering of the types according to their marginal cost of providing the amenity. While far from trac-

[10]The firm types in (Hwang et al., 1998) always have this property because they only consider heterogeneity in σ. All firms share the same ρ so the ones with lower σ are necessarily the higher types. This leads them to the unambiguous conclusion that estimates of willingness to pay by the traditional hedonic wage model will always be biased downward by the presence of producer heterogeneity. However, this assumption that firms will only differ in σ is a departure from previous hedonic wage literature, and they do not offer a justification for it. They do not acknowledge that this is a special case.

ing out the shape of an indifference curve, it is ambiguous whether the average slope of the relationship between x and w is higher or lower than the slope of the indifference curves. This example gives some idea of the full range of distributions of wages and amenities that is possible.

2.4.2 Job Durations Approach to Estimating Preferences

Given that the observed wage-amenity relationship need not be tangent to the workers' indifference curves, the traditional hedonic wage approach has little hope of accurately estimating workers' MWP for an amenity. Using job durations, however, the workers' preferences are identified in this model, and the various cases in Figure 1 can be distinguished in the data.

Gronberg and Reed (1994) were the first to propose an alternative method of estimating MWP based on the durations of job spells that they claim can correctly estimate worker preferences when the observed wage-amenity pairs are generated by an equilibrium search model such as the model described above. In this model, the hazard rate for ending a job spell is a function of $v(w, x)$:

$$h(v(w,x)) = \delta + \lambda_E \left[1 - F(v(w,x))\right] \tag{2.21}$$

Thus

$$\frac{\partial h}{\partial x} = \lambda_E \frac{\partial[-F(v)]}{\partial v}\frac{\partial v}{\partial x} \tag{2.22}$$

and

$$\frac{\partial h}{\partial w} = \lambda_E \frac{\partial[-F(v)]}{\partial v}\frac{\partial v}{\partial w} \tag{2.23}$$

so that we have

$$\frac{\frac{\partial h}{\partial x}}{\frac{\partial h}{\partial w}} = \frac{\frac{\partial v}{\partial x}}{\frac{\partial v}{\partial w}} = MWP. \tag{2.24}$$

In other words, the hazard rate is simply a monotonic transformation of the workers' utility function. Jobs with equal hazard rates trace out the workers' indifference curves. Identification thus depends only on which jobs are observed.

2.5 Continuous Firm Types

Instead of a finite number of firm types, assume there is a joint distribution of firm types $H(\epsilon, \sigma)$, each with technology given by $\phi(x; \epsilon, \sigma)$. Suppose the unconditional distribution

$H(\epsilon)$ is continuous. For given preferences $v(w,x)$, assume ϵ is a sufficient statistic for

$$\frac{\widetilde{\pi}'_v(v;\epsilon,\sigma)}{\widetilde{\pi}(v;\epsilon,\sigma)} = \psi(v;\epsilon). \tag{2.25}$$

The equivalent of our discrete type single crossing condition (2.10) is

$$\psi'_\epsilon(v;\epsilon) > 0 \quad \forall\, \epsilon, v \tag{2.26}$$

For a given value v, then, assume that σ alone determines the firm's choice of the amenity level x, so that we can implicitly define $x(v,\sigma)$ without reference to ϵ.

Notice that since there is not a positive mass of any one ϵ type of firm, it must be that each ϵ type maps to one unique choice of v.[11] The following derivation of this mapping is an adaptation of Burdett and Mortensen (1998). Consider the profit maximization problem of an individual firm, which takes $m(v)$ as given.

$$\max_v \ \widetilde{\pi}(v|\epsilon,\sigma)m(v) \tag{2.27}$$

The first order necessary condition is

$$\widetilde{\pi}'_v(v|\epsilon,\sigma)m(v) + \widetilde{\pi}(v|\epsilon,\sigma)m'(v) = 0 \tag{2.28}$$

or

$$-\frac{m'(v)}{m(v)} = \frac{\widetilde{\pi}'_v(v;\epsilon,\sigma)}{\widetilde{\pi}(v;\epsilon,\sigma)} = \psi(v,\epsilon) \tag{2.29}$$

which verifies that a firm's choice of v depends on ϵ and not σ. A second order sufficiency condition is

$$\widetilde{\pi}''_{vv}(v|\epsilon,\sigma)m(v) + 2\widetilde{\pi}'_v(v|\epsilon,\sigma)m'(v) + \widetilde{\pi}(v|\epsilon,\sigma)m''(v) < 0 \tag{2.30}$$

Now, notice that for $v = u(\epsilon)$, $F(u(\epsilon)) = H(\epsilon)$ in this setting. Then, from (2.4), we can write

$$m(u(\epsilon)) = \frac{1}{\left(1 + K_E\left[1 - H(\epsilon)\right]\right)^2}\frac{1+K_E}{1+K_U} \tag{2.31}$$

Now, equations (2.28) and (2.31) yield the following differential equation for $u(\epsilon)$:

$$u'(\epsilon) = -\frac{1}{\psi(u,\epsilon)}\frac{2K_E h(\epsilon)}{1+K_E(1-H(\epsilon))}$$

[11] This is because Proposition 1 requires $F(v)$ to be continuous and positive on its support, so if a single ϵ mapped to a v interval of positive measure, that would require there to be a positive mass of type ϵ firms, which we are assuming is not the case.

The boundary condition, as in the discrete type case, is $u(\underline{\epsilon}) = v^*$, assuming that $\widetilde{\pi}(v^*|\epsilon, \sigma) \geq 0$ for all ϵ in the support of $h(\epsilon)$.

It can then be shown, as in (Burdett and Mortensen, 1998), that the second order condition (2.30) is indeed satisfied on the relevant support. Note also that it can be shown formally that higher ϵ types will offer the higher utility jobs. If (2.30) holds, then by the Implicit Function Theorem, (2.28) allows us to define $u(\epsilon) = v$, that is, the optimal choice of v as a function of ϵ. Moreover, we have

$$u'(\epsilon) = -\frac{\widetilde{\pi}''_{v\epsilon}(v|\epsilon,\sigma)m(v) + \widetilde{\pi}'_{v}(v|\epsilon,\sigma)m'(v)}{\widetilde{\pi}''_{vv}(v|\epsilon,\sigma)m(v) + 2\widetilde{\pi}'_{v}(v|\epsilon,\sigma)m'(v) + \widetilde{\pi}(v|\epsilon,\sigma)m''(v)}$$

so we see $u'(\epsilon) > 0$ since we assume the denominator is less than zero, and we know the numerator is greater than zero by our assumption stated in (2.26).

Thus is the distribution of the utility dispersion determined. Given $u(\epsilon)$, the actual distribution of wage-amenity pairs in equilibrium in determined by $H(\sigma|\epsilon)$. A change of variables will make it possible to thus derive the joint distribution of w and x.

2.5.1 Example

Take $v(w, x) = wx^2$ (i.e. $\alpha = 2$) and firm technology given by by $\phi_j(x) = \epsilon\sqrt{\sigma} - \frac{\sigma}{2}x^2$ (i.e. $\epsilon = \frac{\rho}{\sqrt{\sigma}}$ in the notation of the discrete type examples). Then we have

$$\widetilde{\pi}(v|\epsilon,\sigma) = \sqrt{\sigma}(\epsilon - \sqrt{2v}). \tag{2.32}$$

Suppose ϵ and σ are independently distributed, with $H(\epsilon) \sim U([1,2])$ and $H(\sigma) \sim U([1.5, 2.5])$. In particular, notice $F(u(\epsilon)) = H(\epsilon) = \epsilon - 1$, for $\epsilon \in [1, 2]$. Then the first order condition (2.29), after substituting (2.31) and some cancellations, yields:

$$\frac{-\frac{\sqrt{2\sigma}}{2}u(\epsilon)^{-\frac{1}{2}}}{\epsilon\sqrt{\sigma} - \sqrt{2\sigma u(\epsilon)}} = -\frac{2K_E\frac{1}{u'(\epsilon)}}{1 + K_E(1 - (\epsilon - 1))}$$

or, after some rearranging:

$$u'(\epsilon) = \frac{2\sqrt{2}\epsilon u^{\frac{1}{2}} - 4u}{\frac{1}{K_E} + 2 - \epsilon}$$

Fortunately, this has the form of a Bernoulli Equation, which has an exact solution.

If we assume $v* = \frac{1}{2}$ and thus $u(1) = \frac{1}{2}$, then we get

$$u(\epsilon) = \frac{1}{2}\left(\frac{(\frac{1}{K_E}+2-\epsilon)^2}{(\frac{1}{K_E}+2)(\frac{1}{K_E}+1)} + \frac{\epsilon^2}{\frac{1}{K_E}+2}\right)$$

Using a change of variables, this expression along with the distributions of ϵ and σ could be used to find an expression for the joint distribution of w and x in equilibrium.

2.5.2 Discussion

Although the differential equation complicates things somewhat, the continuous type formulation appears to have the potential to be much more intuitive and clean to work with. Instead of discrete intervals scattered around the wage-amenity space, as in Figure 1, the equilibrium in the continuous type model is a continuous density of jobs over wage-amenity space, which is a more sensible model and potentially more amenible to estimation. This gives a still clearer picture of the full range of wage-amenity distributions that can result, depending primarily on the distributions of firm technologies that give rise to them.

As an aside, the smooth properties of this model allow us to consider more cleanly what happens in the model as frictions ($\frac{1}{K_E}$ and $\frac{1}{K_U}$) go to zero. The offer distribution $F(v)$ changes very little as frictions go to zero. What does change and, in fact, converges to a perfect information "one price" outcome is the observed distribution of jobs as expressed by $G(v)$. In the limit as frictions go to zero, only the very highest types (say, with equal ϵ but a range of σ values) will have a positive number of workers in equilibrium, as can be seen in (2.3). Thus the wage-amenity distribution as frictions go to zero converges to the indifference curve corresponding to zero profit for the highest types, which is the same as the equilibrium outcome in the traditional hedonic wage model.

2.6 Imposing Equal Profits on Firms: A Baseline

The discussion in Section 2.4 and briefly in the previous section of the various ways that utility dispersion in job offers can alter the observed wage-amenity relationship from that expected in the perfect information model begs the question of which configuration is most likely, if any.

As a baseline case, let us consider the implications of imposing equal profits among firms in equilibrium. Without explicitly modelling it, we can imagine a free entry explanation of why this would tend to be the case. For instance, suppose an economy is made up of various industries, each characterized by a different σ determining their marginal costs of providing the job amenity. The revenue productivity of each industry, and thus its ϵ in the notation of the previous section, depends negatively on the number of firms entering the industry. Now, firms would choose to enter in the industry, or σ type, that would yield the highest expected profit, lowering the profit of that type somewhat by their entry. This free entry process would continue until all σ types yield equal profits. Specifically, if we suppose firms have equal costs of creating a vacancy regardless of industry, the free entry equilibrium will be such that each firms' ex post profits are equal to that entry cost.

This entry argument is just a sketch, and there may even be reasons to believe, in some cases, that jobs that offer higher amenities will tend to yield more or less expected profit to firms.[12] However, it is still helpful to see the implications of imposing equal profits per vacancy created, as this will illuminate as well the implications of expected profits instead being correlated with amenity level.

Before presenting the general case, consider a simple example: two types of firm, with amenities exogenously given: $x_L = 1$ and $x_H = 2$, with worker preferences given by $v(w, x) = wx$. Then we have

$$\begin{aligned} \widetilde{\pi}_j(v) &= \rho_j - \frac{v}{x_j} \\ &= \frac{\epsilon_j - v}{x_j} \end{aligned} \tag{2.33}$$

where it can be shown that the type with the larger $\epsilon_j = \rho_j x_j$ will be the one offering higher utility jobs, according to Proposition 1. The question at hand is which type will have the higher ϵ_j if both received equal expected profits in equilibrium.

Recall from Section 2.4 that with discrete types, the lower type will offer jobs in a utility range from the reservation utility v^* up to some maximum utility $\overline{v}_1$, all of which will yield equal expected profit to the firms of that type in equilibrium. The higher type will offer jobs in an adjacent utility range starting at $\overline{v}_1$ and again earn equal expected profit in that range. Notice that both types offer jobs with utility $\overline{v}_1$ and thus face the same $m(\overline{v}_1)$ for those jobs. If their expected profits are equal then, it must be that $\widetilde{\pi}_L(\overline{v}_1)$

[12]For example, larger training costs or other fixed costs for certain types of jobs may necessitate higher expected profits from the worker to compensate for them. If higher fixed costs are associated with higher or lower amenity jobs, we may expect to see those jobs tend to yield higher expected (variable) profits.

and $\widetilde{\pi}_H(\overline{v}_1)$ are equal. But setting these equal yields:

$$\epsilon_L - \overline{v}_1 = \frac{\epsilon_H - \overline{v}_1}{2} \tag{2.34}$$

If both types are offering jobs in equilibrium, then we must have $\epsilon_j > \overline{v}_1$ for both. Thus we have $\epsilon_L < \epsilon_H$ and we see that imposing equal profits means that the higher amenity type offers higher utility jobs in equilibrium.

Briefly, we can see that if ρ_H were increased so that ϵ_H were larger than the equal profit level, the higher amenity jobs would continue to be the higher utility jobs, even higher in fact. On the other hand, if ϵ_H were below the equal profit level given by (2.34) but still higher than ϵ_L, the higher amenity jobs would continue to be higher utility although now the type H firms would be receiving less expected profit than the type L firms. If ϵ_H were decreased still further to less than ϵ_L, that is, if ρ_H were less than $\rho_L/2$, then the lower amenity jobs would be the higher utility jobs. In this case, by again comparing the expressions for $\widetilde{\pi}_L(\overline{v}_1)$ and $\widetilde{\pi}_H(\overline{v}_1)$, we can see that the expected profits of the higher amenity firms would need to be less than half that of the lower amenity firms. In summary, the high amenity jobs will be the higher utility jobs unless the higher amenity firms earn substantially less profit than the low amenity firms.

Now consider the general case to see that this somewhat surprising result is not limited to this example.

Proposition 2

Assume preferences such that the amenity is a normal good. Consider two types of firms such that at any utility level v, type L would offer less of the amenity x than type H due to higher marginal costs of providing it. These may be two discrete types amidst others or even points along a continuum of types as in Section 2.5. Assume that both types of firm earn the same expected profits in equilibrium at whichever utility levels they choose to offer (labeled as v_L and v_H accordingly):

$$\widetilde{\pi}_L(v_L)m(v_L) = \widetilde{\pi}_H(v_H)m(v_H) \tag{2.35}$$

Finally, assume the two types are ordered according to (2.10) so that one of the types necessarily offers jobs of higher utility than the other type.

Type L offering less of amenity must be the lower type, offering lower utility jobs than type H.

Proof. Recall the firm's problem (2.6) of choosing w and x to provide a given utility level v to the worker. We see from the envelope theorem that

$$\widetilde{\pi}_j'(v) \quad = \quad -1/v_w(w,x). \tag{2.36}$$

It is a property of x being a normal good that along an indifference curve of utility v, the marginal utility of w increases with increasing x.[13] (See the Appendix for a proof.) Since firm L would provide less amenity x than H if both were offering utility v, $v_w(w,x)$ is less for the utility v jobs offered by firm L, so that (2.36) implies

$$\widetilde{\pi}_L'(v) < \widetilde{\pi}_H'(v) \quad \forall\, v \tag{2.37}$$

Now, from profit maximization, we know $\widetilde{\pi}_L(v_L)m(v_L) \geq \widetilde{\pi}_L(v_H)m(v_H)$. But if we substitute for the LHS using the equal profits condition (2.35), we find $\widetilde{\pi}_L(v_H) \leq \widetilde{\pi}_H(v_H)$. But combine this with (2.37) for v_H and we have:

$$\frac{\widetilde{\pi}_L'(v_H)}{\widetilde{\pi}_L(v_H)} < \frac{\widetilde{\pi}_H'(v_H)}{\widetilde{\pi}_H(v_H)}$$

That is, by Proposition 1, it must be that type L is the lower type and offers lower utility jobs than type H. □

The result depends critically on the assumption that the amenity is a normal good for the workers. If the amenity is an inferior good, the opposite is true: if the firm types earn equal profits, the lower amenity jobs will be the higher utility jobs (i.e. case (c) in Figure 1). If the workers' preferences are quasilinear, so that the amenity is neither normal nor inferior, all firm types would earn equal profits by all offering the same utility.

Returning to the normal good case, what is the intuition behind the above result? The decision of which utility each firm offers depends on balancing the benefit of increasing $m(v)$ and the cost of decreasing $\widetilde{\pi}(v)$ by increasing v. All firms face the same

[13] It may seem that since we would typically assume that w (standing in for consumption) is also a normal good, there should be a symmetric argument showing the opposite result. Indeed, it is the case that $v_x(w,x)$ is decreasing along an indifference curve as x increases, so the marginal utility of x is greater for firm L when offering the same utility v as firm H. However, whereas all firms share the same unitary marginal cost of providing wages, firms differ in their costs of providing the amenity. Indeed, this is the reason firms offer different levels of the amenity at all. Recalling equation (2.7) , we see that $\phi'(x)/v_x(w,x) = -1/v_w(w,x)$ since optimizing firms equalize the marginal costs of providing utility by the two avenues of wages and amenities. So although low amenity firms could increase worker utility at a greater rate by increasing x, their marginal costs of increasing x are even higher and their overall marginal cost of increasing utility is therefore higher than that of firms offering more amenity at the same utility v.

marginal benefit $m'(v)$ but those offering less of the amenity face a larger marginal cost $\widetilde{\pi}'(v)$ at each utility level v. This means that firms with higher marginal costs of providing the amenity will optimize by choosing to offer lower utility jobs than the lower cost (higher amenity) firms.

2.7 Conclusion

I have illuminated the full range of possible wage-amenity distributions that can result from wage dispersion in a general equilibrium search framework with homogeneous workers, correcting some mistaken implicit assumptions in the existing literature. This is important in itself because search frictions and wage dispersion are increasingly understood as important in the labor market, and the compensating wage differential is a fundamental concept in labor economics. In addition, my conjecture that firms tend to earn equal expected profits is the first attempt to explain why higher amenity jobs would tend to be higher utility jobs in equilibrium, an idea that has so far been taken for granted in the literature without justification.

Existing empirical work using the job durations method is consistent with the idea that higher amenity jobs tend to be more desirable to workers, as it indicates that workers' true MWP for a job amenity is larger than the slope of the observed relationship between wages and amenities. However, the measures they report do not indicate whether or not my conjecture of equal expected profits is in fact the case, and further work is needed on the job durations method in general.

Bibliography

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.

ATTANASIO, O. AND M. BROWNING (1995): "Consumption over the Life Cycle and over the Business Cycle," *The American Economic Review*, 85, 1118–1137.

ATTANASIO, O. AND H. LOW (2004): "Estimating Euler equations," *Review of Economic Dynamics*, 7, 406–435.

ATTANASIO, O. P. AND G. WEBER (2010): "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy," *Journal of Economic Literature*, 48, 693–751.

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.

BANERJEE, A. AND S. MULLAINATHAN (2010): "The Shape of Temptation: Implications for the Economic Lives of the Poor," *NBER Working Paper*.

BESLEY, T. (1995): "Savings, credit and insurance," *Handbook of development economics*, 3, 2123–2207.

BLISS, C. (2004): "Some implications of a variable EIS," *mimeo*.

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.

BONHOMME, S. AND G. JOLIVET (2009): "The pervasive absence of compensating differentials," *Journal of Applied Econometrics*, 24, 763–795.

BROWNING, M. AND T. CROSSLEY (2000): "Luxuries are easier to postpone: A proof," *Journal of Political Economy*, 108.

BUERA, F., J. KABOSKI, AND Y. SHIN (2012): "The Macroeconomics of Microfinance," *mimeo*.

BURDETT, K. AND D. T. MORTENSEN (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–73.

DALE-OLSEN, H. (2006): "Estimating Workers' Marginal Willingness to Pay for Safety using Linked Employer-Employee Data," *Economica*, 73, 99–127.

DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2008): "Returns to Capital in Microenterprises: Evidence from a Field Experiment," *Quarterly Journal of Economics*.

——— (2009): "Measuring microenterprise profits: Must we ask how the sausage is made?" *Journal of Development Economics*, 88, 19–31.

——— (2012): "One-Time Transfers of Cash or Capital Have Long-Lasting Effects on Microenterprises in Sri Lanka," *Science*, 335, 962–966.

DUFLO, E., M. KREMER, AND J. ROBINSON (2008): "How High are Returns to Fertilizer? Evidence from Field Experiments in Kenya," *AER Papers and Proceedings*.

——— (2010): "Nudging farmers to use fertilizer: theory and experimental evidence from Kenya," *mimeo*.

DUPAS, P. AND J. ROBINSON (2012a): "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya," *mimeo*.

——— (2012b): "Why Don't the Poor Save More? Evidence from Health Savings Experiments," *mimeo*.

DYNAN, K. (1993): "How prudent are consumers?" *Journal of Political Economy*, 1104–1113.

EPSTEIN, L. AND S. ZIN (1989): "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica: Journal of the Econometric Society*, 57, 937–969.

GRONBERG, T. J. AND W. R. REED (1994): "Estimating Workers' Marginal Willingness to Pay for Job Attributes Using Duration Data," *Journal of Human Resources*, 29, 911–31.

GUVENEN, F. (2006): “Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective,” *Journal of Monetary Economics*, 53, 1451–1472.

HECKMAN, J., R. MATZKIN, AND L. NESHEIM (2010): “Nonparametric identification and estimation of nonadditive hedonic models,” *Econometrica*, 78, 1569–1591.

HWANG, H.-S., D. T. MORTENSEN, AND W. R. REED (1998): “Hedonic Wages and Labor Market Search,” *Journal of Labor Economics*, 16, 815–47.

KABOSKI, J. AND R. TOWNSEND (2011): “A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative,” *Econometrica*, 79, 1357–1406.

KARLAN, D., R. KNIGHT, AND C. UDRY (2012): “Hoping to win, expected to lose: Theory and lessons on micro enterprise development,” *mimeo, Yale University*.

KARLAN, D. AND J. MORDUCH (2009): “Access to Finance,” in *Handbook of Development Economics*, ed. by D. Rodrik and M. Rosenzweig, vol. 5.

LANG, K. AND S. MAJUMDAR (2004): “The Pricing of Job Characteristics When Markets Do Not Clear,” *International Economic Review*, 45, 1111–28.

MANNING, A. (2003): *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton University Press.

MORTENSEN, D. (2003): *Wage Dispersion: Why Are Similar Workers Paid Differently?*, Massachusetts Institute of Technology.

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.

ROSEN, S. (1974): “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy*, 82, 34–55.

——— (1986): “The Theory of Equalizing Differences,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and R. Layard, Elsevier Science Publishers, chap. 12, 641–90.

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

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

VAN OMMEREN, J., G. VAN DEN BERG, AND C. GORTER (2000): "Estimating the Marginal Willingness to Pay for Commuting," *Journal of Regional Science*, 40, 541–563.

Appendix A

Chapter 1 Appendices

A.1 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 utility[1]: $u(c) = \frac{1}{1-\gamma}(c - c_s)^{1-\gamma}$ 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:

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

Atkeson and Ogaki (1996) combine subsistence levels and additive multi-good prefer-

[1] This is also called Stone-Geary preferences in demand analysis, although then c_s is restricted to be positive unlike for the general HARA form.

ences in what they call extended addilog utility, adding additional flexibility. In any case, these specifications 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 to avoid the econometric difficulties with estimating non-linear equations. However, the resulting expressions for how the EIS depends on the consumption level are less clean than the specification I use in Section 1.7.1.

A.2 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 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).

A.3 Tables and Figures

A.3.1 Tables

Table A.1: Descriptive Statistics

	Mean					
	Control	Treated	Std. Dev.	Min	Median	Max
Age of Business	9.89	11.49	10.43	1	7	55
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
Digitspan Score	5.7	5.8	1.3	3	6	11
Initial Capital	272.2	280.5	252.5	1.4	200	952
Initial Profits	37.2	41.1	34.9	2	30	250
Init. Household Income	88.7	82.4	52.8	12	80	450
Init. Consumption/person	23.6	27.1	18.2	3.6	21.8	129.5
Init. Household Size	4.9	4.9	1.7	1	5	10
Init. Food Share	0.611	0.630	0.164	0.075	0.637	0.963
Init. Rec. / Elec. Share	0.027	0.031	0.064	0.000	0.006	0.456
2 year Savings Rate	0.015	0.036	0.120	-0.174	0.006	1.220
2 yr % Change Capital	0.554	0.789	0.917	-1.303	0.575	4.868
2 yr % Change Hhld Inc.	0.442	0.578	0.670	-1.191	0.483	2.603
2 yr % Change Cons.	0.169	0.120	0.557	-1.401	0.163	2.180

Table A.2: Euler Equation Estimates

	Total Income		Profits Only	
Parameter	IV	OLS	IV	OLS
EIS β_1	0.028*	0.004	-0.06	0.001
	(0.014)	(0.004)	(0.07)	(0.009)
Constant term β_0	-0.06	0.079	0.24*	0.10
	(0.10)	(0.053)	(0.15)	(0.050)
N	129	129	127	127

IV estimation is by GMM, and all standard errors reported in parentheses are robust. * means statistically different from zero at the 10% level. Negative returns estimates are left out (28 and 23 for total income and only profits, respectively). $\hat{r}_i$ 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.

Table A.3: OLS Regression of Returns on Instruments

Parameter	Log Returns to Total Income	Log Returns to Profits
Age of business	0.006	-0.004
	(0.09)	(0.052)
Manufacturing	4.9*	-1.4
	(2.1)	(1.3)
Services	3.9	-0.59
	(3.0)	(1.9)
Male owner	-2.9	-1.0
	(1.8)	(1.1)
Owner education	-0.61*	0.03
	(0.33)	(0.15)
Digitspan score	0.47	-0.02
	(0.74)	(0.58)
Constant	8.3	3.1
	(5.8)	(3.5)
F test p value	0.06	0.60
N	129	127

Estimation is OLS with robust standard errors reported in parentheses. * means statistically different from zero at the 10% level. Negative returns estimates are left out (28 and 23 for total income and only profits, respectively). The independent variable is the natural log of annualized returns. The default is a female in a retail industry.

Appendix B

Chapter 2 Appendices

B.1 Property of Normal Goods

Assume the worker has wealth M and can purchase the amenity at price p: $px+w = M$. The worker's problem can be written as

$$\max_x \; v(M - px, x)$$

with F.O.C.

$$v'_w(M - px, x) \cdot (-p) + v'_x(M - px, x) = 0$$

By the implicit function theorem, we have

$$\frac{dx}{dM} = -\frac{-pv''_{ww} + v''_{wx}}{p^2 v''_{ww} - 2pv''_{wx} + v''_{xx}}$$

where the denominator must be negative due to the second order condition. Thus, if x is a normal good and $\frac{dx}{dM} > 0$, it must be that

$$(-\frac{v'_x}{v'_w})v''_{ww} + v''_{wx} > 0.$$

Finally, define $w(x)$ as the wage that gives utility v for any amenity level x. We want to show that $v_w(w(x), x)$ increases as x increases. But the total derivative with respect to x is just the above expression that we said must be positive if the good x is normal.

(Notice $w'(x) = -v'_x(w(x), x)/v'_w(w(x), x)$, the marginal rate of substitution.)

B.2 Tables and Figures

B.2.1 Figures

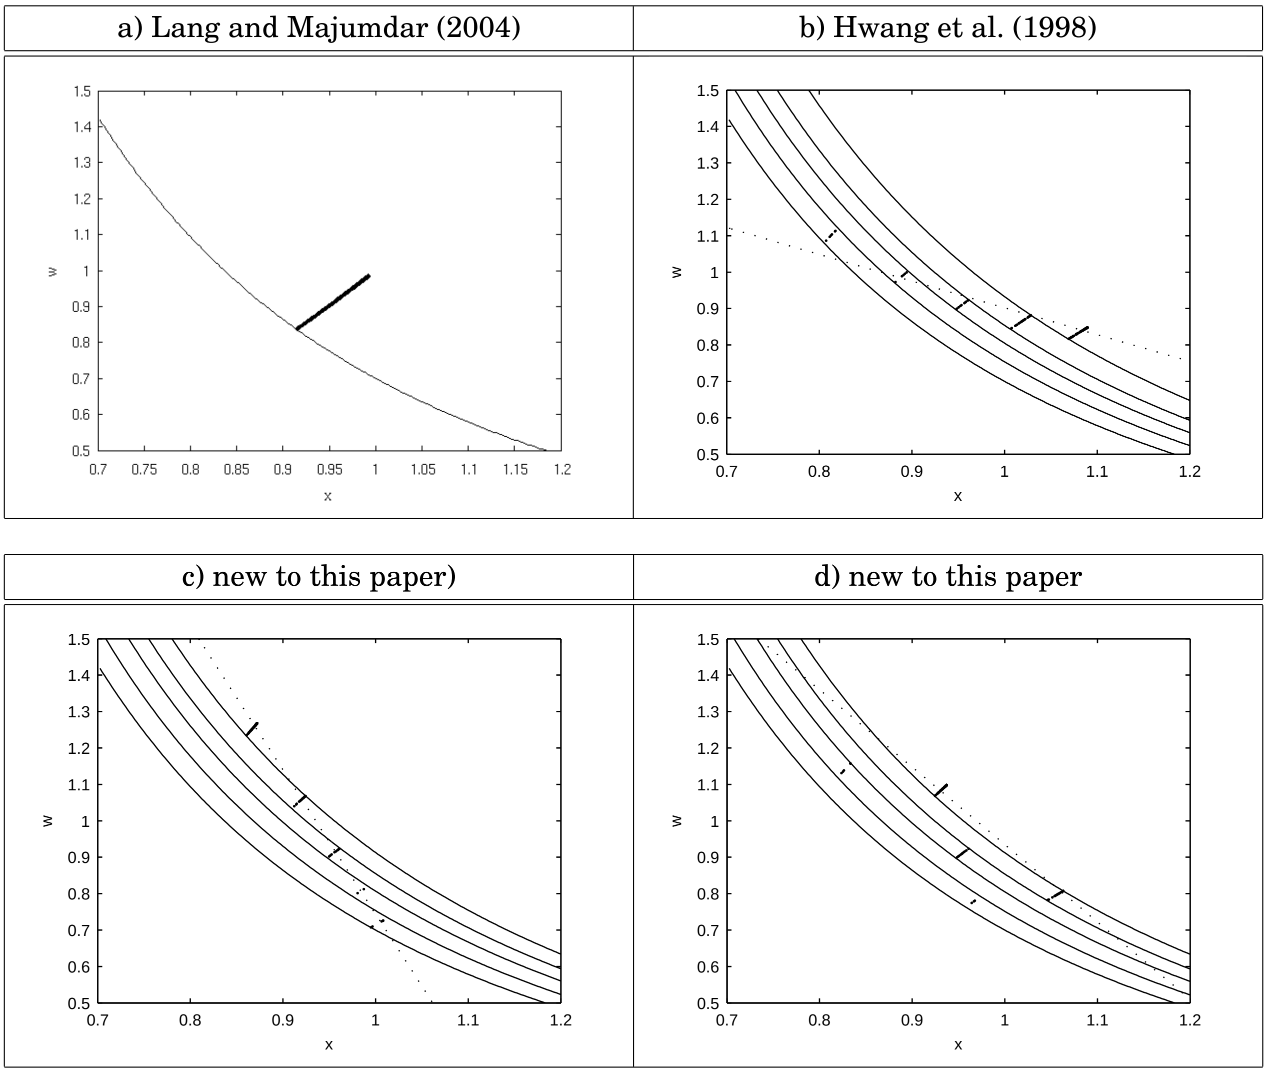

Figure 1: Wage-amenity offers for different distributions of firm heterogeneity. The curved lines are indifference curves. The other marks are wage-amenity pairs characterizing a sample of 200 jobs in steady-state equilibrium. The dotted line in (b-d) is a best-fit line through the wage-amenity pairs, such as would be used to estimate a linear specification of a hedonic wage model.

a) Homogeneous firms with $\rho = 2\sqrt{2}$ and $\sigma = 2$.
b) $\{(\rho_j, \sigma_j)\} = \{(2.4, 10/3), (2.1, 10/4), (1.9, 2), (1.8, 10/6), (1.7, 10/7)\}$.
c) $\{(\rho_j, \sigma_j)\} = \{(1.57, 10/7), (1.72, 10/6), (1.9, 2), (2.17, 10/4), (2.54, 10/3)\}$
d) $\{(\rho_j, \sigma_j)\} = \{(1.69, 10/6), (2.43, 10/3), (1.9, 2), (1.864, 10/7), (2.20, 10/4)\}$