

Preparation for Retirement, Financial Literacy and Cognitive Resources

Adeline Delavande, Susann Rohwedder and Robert Willis



Project #: UM08-07

Preparation for Retirement, Financial Literacy and Cognitive Resources

Adeline Delavande
RAND, Universidade Nova de Lisboa and CEP

Susann Rohwedder
RAND

Robert Willis
University of Michigan

September 2008

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
<http://www.mrrc.isr.umich.edu/>
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98362-5-04). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Regents of the University of Michigan

Julia Donovan Darrow, Ann Arbor; Laurence B. Deitch, Bingham Farms; Olivia P. Maynard, Goodrich; Rebecca McGowan, Ann Arbor; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; S. Martin Taylor, Grosse Pointe Farms; Katherine E. White, Ann Arbor; Mary Sue Coleman, ex officio

Preparation for Retirement, Financial Literacy and Cognitive Resources

Adeline Delavande, Susann Rohwedder and Robert Willis

Abstract

Traditional economic models assume that individuals have full information and act perfectly rationally. However, we show that there is considerable variation in financial literacy in the population and propose modeling the acquisition of financial knowledge in a human capital production framework. The model makes several predictions, notably with respect to portfolio choice. For example, it helps explain household non-participation in the stock market for some fraction of the population, and it provides guidance about the share of risky assets to hold for other types of households. Estimation of the human capital production function for financial knowledge on data from the Cognitive Economics Survey yields results that are consistent with important features of the model.

Authors' Acknowledgements

We gratefully acknowledge support for the research in this paper from the Michigan Retirement Research Center and the National Institute on Aging. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the Federal Government. Data used in the paper were collected with support from NIA P01 AG026571 program project grant, "Behavior on Surveys and in the Economy Using HRS," Robert J. Willis, PI; NIA R37-AG007137, "Assessing and Improving Cognitive Measurements in the HRS," John J. McArdle, PI. We also thank Cynthia Doniger for excellent research assistance.

1. Introduction

Successful planning for retirement requires forming expectations about a number of events, even far into the future, and integrating and translating these into economic decisions. Standard economic theory offers clear prescriptions on how individuals should make financial and retirement decisions, but many individuals have only a vague idea of these prescriptions or even have quite incorrect ideas. In this paper we develop a theoretical model that treats financial knowledge as a form of human capital and the acquisition of knowledge as an investment. We assume that investment in improved financial knowledge allows households to obtain a higher expected rate of return on their assets, holding risk constant. The cost of acquiring additional knowledge depends on cognitive ability, effort, the existing stock of knowledge, and payments for financial education or advice. The benefit of additional knowledge is equal to the increase in the expected rate of return multiplied by the amount of savings that will earn the higher return. The scale economy created by this multiplicative relationship helps to explain why the fraction of wealth held in stock tends to be an increasing function of total wealth, why low wealth households often hold no stock and, in addition, shows how cognitive ability, the opportunity cost of time and the availability of advice influence portfolio decisions.

We use this theory as a framework for our empirical work using new data from the American Life Panel and Cognitive Economics Survey. In particular, we study

how variations in financial literacy (knowledge) and effort vary in the population.

Acquiring additional financial information for a knowledgeable person will be low cost, whereas it will be high cost for a person largely lacking in financial knowledge.

However, if the person lacking in knowledge is willing to exert effort to learn or seek outside advice he may nevertheless make optimal or close to optimal portfolio decisions.

2. A Theory of the Accumulation of Financial Knowledge

2.1 Finite Financial Knowledge. In order to interpret our empirical work, it is useful to present a very simple human capital model of acquisition of financial sophistication, and its value to retirement planning. From a theoretical point of view, a person with more financial sophistication faces a lower cost of thinking constructively about retirement, and what actions to take. Conversely, a person with a strong concern about retirement may learn more about financial matters. These are general characteristics of any human capital decision in which the cost of acquiring knowledge is independent of its future rate of utilization.

The following is a simple formal model that embodies these ideas. As a point of departure, we consider a model of optimal portfolio choice due to Merton (1969). In the Merton model, it is assumed that there is a safe asset with a sure rate of return r and a risky asset that follows a continuous time random walk (Brownian motion) whose annual returns are normally distributed with mean μ and variance σ^2 . Assume that a household saves for retirement at some predetermined age T , that it cares only about the amount of wealth at retirement, W_T , and that its preferences are CRRA where γ is its coefficient of

relative risk aversion. Merton derives an elegant and very intuitive rule for optimal saving:

$$(1) \quad s_t = s^* = \frac{\mu - r}{\gamma \sigma^2}.$$

Assuming that $\mu - r > 0$, this model implies that it is optimal to hold a constant fraction, $s^* > 0$, of household wealth in the risky asset, which we shall think of as stock. It further implies that this fraction is higher the lower the degree of risk aversion, the higher the expected return on stock and the less risky they are. The Merton model assumes that individual households are able to buy and sell stock without any transaction cost, that they have well formed probability beliefs about the distribution of stock returns and that they have the knowledge to obtain these returns through their stock market transactions. Kezdi and Willis (2008) explore the relationship between stock market expectations and portfolio choice under the (implicit) assumption that people know how to construct an efficient portfolio in which (μ, σ) is a point on an efficient mean-variance frontier.

In this paper, we assume that people may not know how to construct such a portfolio. As we show later in this paper, our data indicate that there is substantial heterogeneity in objective measures of people's knowledge about the working of financial markets and, moreover, that only about a quarter of the people in our sample believe that they "understand the stock market reasonably well." In addition, Campbell, Laurent and Soldini (2007) show that there is heterogeneity in the efficiency of the

portfolios held by individual households and they provide estimates of the welfare loss due to inefficiencies. On the other hand, our data also suggest that people can improve their understanding of financial matters through investments in knowledge. For example, controlling for years of education, the number of economics courses taken in the past by people in the CogEcon sample of people over age 50 has a large effect on their score on a twenty-five item test of financial knowledge.

In the Merton model, the expected rate of return on the risky asset, μ , is treated as a known parameter and, in addition, in many applications of the model it is assumed that expectations are rational in the sense that μ and σ have common values across households which may be estimated from the historical pattern of stock market returns and that these parameters will continue to govern future returns. These assumptions imply that $\mu > r$ and, in turn, (1) implies that all households should hold at least some stocks in their portfolio and that the optimal share of stock, s^* , is independent of the level of wealth. As is well known, neither of these implications holds empirically. Many households fail to hold any stocks at all—a phenomenon dubbed “the stockholding paradox” by Halioussos and Bertaut (1995)—and the share of stocks in household portfolios increases dramatically with total wealth.

Kezdi and Willis (2008) instead view the Merton model in (1) as a simple microeconomic model of demand for stock holding. In their version, individuals may have different degrees of risk aversion and different beliefs about both expected stock market returns and risk (μ and σ) and different degrees of risk aversion (γ) without

any restrictions imposed on the distribution of these parameter across household by general equilibrium considerations. They use survey data on individual's probabilistic beliefs about stock returns and household stock holding from the HRS in order to jointly estimate a model of the determinants of risk aversion and the mean and variance of subjective beliefs about stock market returns. They find that the failure of many households to hold stock can be explained very simply. Contrary to historical data, many households believe that the expected return on stocks is lower than the interest rate and, in addition, believe that the variability of returns, or risk, is much higher than indicated by the historical data.

Although the model we develop in this section follows Kezdi and Willis by using the Merton model as the basis for a microeconomic model of household asset demand, we suppress considerations of variation in probabilistic beliefs about general stock market returns in order to focus on investment in the acquisition of financial knowledge within the simplest possible deterministic framework. Nonetheless, one aspect of the stock of financial knowledge that goes into people's investment decisions is their understanding of the probability distribution of asset returns. In future research, we hope to generalize the model presented in this paper to incorporate probabilistic beliefs of stock market returns and of future retirement income from pensions and Social Security as a component of an individual's human capital and to consider activities undertaken by an individual to learn about the distribution of these events as an investment in human capital.

For our purposes, we shall be thinking of the standard extension of the Merton model to a situation in which there are many different risky assets comprised of stocks, long term bonds and the like. In such a world, each possible portfolio of risky assets is characterized by an expected rate of return and degree of risk denoted by (μ_j, σ_j) for the j^{th} such portfolio. Given complete knowledge of the means, variances and covariances of the distribution of returns for all risky assets together with a mastery of finance theory and free access to appropriate financial institutions, it is possible to construct a set of portfolios that describe an efficient mean-variance frontier, $\mu^* = \phi(\sigma)$, that gives the maximum possible expected rate of return for any given level of risk.

Ordinary people lack the degree of theoretical and empirical knowledge and institutional resources necessary to construct such portfolios. Such people do have access to a wide range of professional financial advice from private or public sources, amateur advice, self-help books, newspapers and magazines from which they may obtain useful knowledge about the potential risks, returns and costs of various financial instruments available to them, their tax implications, and the institutional means to implement an efficient investment strategy. By pursuing these opportunities to increase their own financial knowledge, people may increase the expected return from their portfolio of risky assets without incurring additional risk.

On the other hand, the market is full of badly priced or excessively risky assets (or risks that are hidden from view) that a myriad of salesmen and television gurus with

bogus credentials are only too happy to sell to the gullible.¹ Those with low levels of financial knowledge who are unable to distinguish good deals from bad ones or good advice from bad advice may receive risk-adjusted returns on risky investments below the rate of return available on safe assets. If people recognize that they have a low level of financial knowledge—in other words, they “know what they don’t know”—it may be rational for them to avoid investment in risky assets. Those who “don’t know what they don’t know” are likely to be vulnerable to scams or more likely to become “stock jocks” that follow get rich quick trading strategies. (See Kimball and Shumway, 2008.)

In our theoretical model, we attempt to capture some of these ideas in a very simple way. For simplicity, we shall assume that people do “know what they know,” an assumption that we later show is reasonable for many people. Holding the level of risk constant, the maximum expected rate of return on the portfolio of risky assets that is available to a given person i depends on his financial knowledge, K_i , and advice received from others, A_i , as follows:

$$(2) \quad \mu_i = f(K_i, A_i, \sigma); \mu_i < \mu^* .$$

Equation (2) is illustrated diagrammatically in Figure 1. At any given level of risk, σ , we assume that increased knowledge allows the individual to obtain a better

¹ A number of financial dangers faced by older Americans face are described in a series of *New York Times* articles by Charles Duhigg (2006, 2007, 2008) dealing with a range financial products such as reverse mortgages, tax shelters, annuity products, identity theft, long term care insurance, and life insurance. In addition, of course, there many reputable books and magazines that discuss matters such as the importance for long term retirement saving of shopping for low-priced mutual funds, the value of diversification or the danger of eating up profits by excessive trading.

expected return ($f_k > 0$) with the expected return asymptotically approaching a theoretical maximum of μ^* on the mean-variance frontier as the level of knowledge increases indefinitely. Below some level of knowledge denoted by \underline{K} , however, we assume that the expected returns from a risky portfolio are lower than what is available from safe, unsophisticated assets because of perceived high costs and high risks of a portfolio whose components are chosen “blindly”.

Economic theory suggests that when people know that they do not understand a complex product but do understand the seller’s motivation to make a profit, they will distrust the information they receive and may pass up products that a more informed buyer would find to be beneficial (Milgrom, 2008). Advice from others and one’s own financial sophistication are likely to be complements ($f_{kA} > 0$). For people with sufficiently low levels of knowledge, the effect of advice on the expected return may be negative ($f_k < 0$). These financially unsophisticated people recognize that much of the advice they receive (whether from friends and relatives, professional advisors or salesmen) may be ill-informed or self-serving and they also recognize that they are unable to distinguish good from bad advice. Their best course of action may be to avoid investing in anything but riskless assets. Conversely, people with more sophisticated levels of knowledge may be better able to utilize advice and sophisticated financial products in order to achieve expected returns on portfolios of risky investments that exceed not only the safe return but also the returns that they could achieve by constructing their own portfolios from scratch. Possible effects of increased financial advice are illustrated diagrammatically in Figures 2a and 2b. In the first case, an increase

in the availability of advice would reduce the threshold value of financial knowledge that is needed in order to obtain a higher expected return from the financial market with the opposite effect in the second case.

Obviously, the value of financial knowledge depends not only on how much it increases the rate of return per dollar of investment but also on how many dollars of investment will benefit from this increased return. Letting $EV(K)$ denote the expected value of financial knowledge and noting that $EV(\underline{K})=0$, the expected value of a stock of financial knowledge above the threshold value is:

$$(3) \quad EV(K_t) - EV(\underline{K}) = EV(K_t) = (\mu_t - r)s^*(\mu_t)EW(\mu_t, \sigma); K_t \geq \underline{K}.$$

In words, this equation says that the value of a given level of knowledge is equal to the excess return of the risky asset multiplied by the amount of retirement wealth that will be held in the risky asset where $s^*(\mu_t, \sigma)$ is the optimal share of the risky asset determined by (1) and $W(\mu_t)$ is the level of retirement wealth.²

Because of the scale economy implied by the multiplicative relationship between excess returns and magnitude of retirement wealth, Equation (3) has important implications both for incentives to investment in forms of human capital that enhance financial knowledge and for decisions about the level and portfolio composition of

² The notation used for this equation probably needs further work to be valid for a general continuous time dynamic optimization problem under uncertainty. Later in this section, we will use this equation in a very simple two period model where the intuition underlying this equation and its implications for portfolio choice and human capital accumulation are transparent.

retirement savings. In the next subsection, we specify a human capital production function that describes the technology of the acquisition of financial knowledge and, in the following subsection we develop a simple two period model of life cycle decisions which incorporates this production function. These models form the basis for our empirical work in remainder of the paper.

2.2. Investment in Financial Knowledge. We assume that financial knowledge is accumulated according to a human capital production function inspired by the pioneering model of Ben-Porath (1967) and more recent work by Cunha and Heckman (2007). In the Ben-Porath model the rate of acquisition of new human capital depends on the allocation of own time and purchased inputs; the productivity of own time depends on the stock of knowledge the person has accumulated in the past; and total factor productivity in the investment process depends on general intelligence. Although focused on skill development during childhood and labor market careers, Cunha and Heckman emphasize features of human capital accumulation that are also relevant to the development of non-market skills at later stages of the life cycle such as skill in financial or health decision making. In particular, they assume that there are many forms of human capital and emphasize dynamic complementarities between current learning and past investments.

For simplicity, we follow Ben-Porath by assuming that the production function for investment in human capital has a Cobb-Douglas form:

$$(4) \quad Q_t = \frac{dK_t}{dt} = \alpha (e_t K_t)^{\beta_1} H_t^{\beta_2} M_t^{\beta_3} E_t^{\beta_4}; \quad \beta_1 + \beta_2 + \beta_3 + \beta_4 < 1$$

where Q_t is the amount of additional financial knowledge learned during period t . The actual amount of time or effort devoted to learning is e_t , but the number of efficiency units of time is $e_t K_t$, where K_t is a function whose elements include the stock of financial knowledge the person possesses at the beginning of period t . Thus, financial knowledge is self-productive in the sense that greater current knowledge enhances the efficiency with which future knowledge may be obtained.

In addition, the productivity of learning new financial knowledge depends on other inputs. These include the current stock of other forms of human capital, H_t ; purchased inputs such as books and magazines, educational courses, or financial advice, M_t ; and unpriced “environmental” inputs including the knowledge of family and friends, news provided by television and newspapers, E_t . Finally, α is an ability parameter that can be interpreted as a measure of fluid intelligence or IQ defined as the ability to reason in novel situations without regard to specific context (Kimball and Willis, in progress). The assumption that $\beta_1 + \beta_2 + \beta_3 + \beta_4 < 1$ implies that there are decreasing returns to increased effort within a given period causing the marginal opportunity cost of Q_t to be an increasing function of Q_t . This, in turn, suggests that it is efficient to spread the accumulation of knowledge over time, creating a rising life cycle trajectory of the stock of knowledge.

3. The Implications of Financial Knowledge for Retirement Savings and Portfolio Choice

To model the acquisition of financial knowledge and its impact on retirement saving, we shall consider a two period life cycle model where the first period corresponds to the life cycle phase, typically beginning sometime in the forties, in which a household begins saving for retirement and the second period corresponds to the age of retirement. The household's utility during the rest of their lifetime is

$$(5) \quad U = U(C_1, C_2).$$

where consumption during the two periods is denoted by C_1 and C_2 .

At the beginning of period 1, the household's current budget constraint is given by

$$(6) \quad W_1 + Y_1 = C_1 + S_1 - SST_1 - DBC_1 - I_1,$$

where W_1 is the value of assets at the beginning of period 1; Y_1 , C_1 and S_1 are, respectively household income, consumption and saving during period 1; SST_1 is Social Security tax paid, DBC_1 and I_1 is the cost of investment in financial knowledge during the period. The level of the household's retirement consumption, C_2 , is equal to its

accumulated wealth, W_2 plus the Social Security and DB pension benefits it receives, $SSB_2 + DB_2$;

$$(7) \quad C_2 = W_2 + SSB_2 + DB_2.$$

Retirement wealth accumulation depends on both the amount saved during the first period and on the realized rate of return on the portfolio of financial assets into which the household places its savings. That is,

$$(8) \quad W_2 = W_1 + (1 + \rho)S_1$$

The portfolio rate of return, ρ , depends on the rates of return and portfolio shares of the risky and safe assets as defined in (1). That is,

$$(9) \quad \rho = s\mu + (1 - s)r.$$

We shall use a two-period version of the theoretical framework presented in Sections 2.1 and 2.2, to illustrate the effect of financial knowledge on portfolio choice. For simplicity, we ignore financial advice and other forms of human capital and treat the return from the risky asset as a certainty-equivalent. The household's rate of return to investing in the risky asset is

$$(10) \quad \mu = f(K_1)$$

where K_1 is the stock of financial knowledge the household possesses when it makes its savings and portfolio decisions at the end of period 1. This stock of knowledge, in turn, is equal to the sum of its initial stock of knowledge at the beginning of period 1 plus the amount of investment in knowledge during period 1:

$$(11) \quad K_1 = K_0 + Q_1.$$

As discussed above in Section 2.1, there is a threshold value of financial knowledge, \underline{K} , such that $\mu = f(K_1) \leq r$ for $K_1 \leq \underline{K}$ and $\mu = f(K_1) > r$ for $K_1 > \underline{K}$. That is, the attainable rate of return on the risky asset does not exceed the rate of return on the safe asset unless the household's level of financial knowledge exceeds some threshold level.

The dollar cost of investment in financial knowledge during period, I_1 , is equal to the opportunity cost of the time and effort expended to acquire additional financial knowledge given by the cost function,

$$(12) \quad I_1 = I(Q_1); \quad I' > 0, I'' > 0.$$

Given diminishing returns to effort, the marginal cost curve of investment is upward sloping.

The simple indifference curve diagram in Figure 3 puts together equations (5)-(12).³ The household's endowment is at point e . At this point, the household's resources at the beginning of period 1, which consist of income and initial assets (given by the distance $0a$) and the present value of its Social Security benefits and any other defined benefit pensions that will be received at the beginning of period 2 (given by the distance $0b$). Assuming that the household only has access to the riskless asset and that it can borrow or lend at the riskless interest rate r , it may achieve any intertemporal consumption path, (C_1, C_2) , along the wealth constraint given by the solid black line passing through point e with an absolute slope of $1 + r$. The present discounted value of its lifetime wealth is given by the distance $0w$. Its optimal consumption path is indicated by the tangency between its indifference curve, U^* and the wealth constraint at point c . Since this point is to the northwest of the endowment point e , the household must engage in retirement saving.

The household may also choose to invest in risky assets. As drawn, the figure assumes that the household's initial level of financial knowledge is considerably below the threshold, \underline{K} . If it wishes to obtain a higher rate of return than r , it will need to augment its knowledge through human capital investment during period 1. The concave dotted line in the figure indicates how the household life cycle consumption path is influenced by increased investment in human capital and associated changes in the amount of savings, the allocation of retirement savings between the safe asset and risky assets and the expected rate of return to the portfolio. Initially, the dotted line lies below

³ The formal derivation of this diagram, which is straightforward but tedious, is relegated to Appendix 1 (not yet written).

the solid line because the cost of investment is not offset by additional returns from the purchases of the risky asset. Indeed, beginning in a northwest direction from point e , the absolute slope of the dotted line is initially smaller than $1 + r$, indicating the stock of financial knowledge remains smaller than \underline{K} .

As additional investment causes the level of financial knowledge to exceed \underline{K} , it pays to begin including risky assets in the portfolio. As both the rate of return rises and the volume invested in risky assets increases, the dotted line eventually crosses the solid line, implying that the household's life time welfare may be increased by making a sufficiently large investment in financial knowledge and applying this knowledge to a sufficiently large volume of the risky asset. The optimal levels of investment and savings occur at the tangency between the indifference curve U^* and the dotted frontier at point d . Note that the present value of lifetime consumption, discounted at the riskless rate of interest, this point is given by the distance $0w'$ where the dotted straight line passing through point d intersects the horizontal axis.

Although the investment in financial knowledge raises lifetime wealth in Figure 3, note that lifetime utility is the same at points c and d since both points lie on the same indifference curve. The possibility of multiple optima arises because the multiplicative relationship in equation (3) between the rate of return on the risky asset, the share of risky assets in the portfolio and the amount invested in the risky asset causes the opportunity set to be non-convex.

It is easy to use this indifference curve diagram to generate a number of comparative static implications about the decision to invest in risky assets. For example, if the initial stock of financial knowledge, K_0 , is increased, less investment, Q_1 , will be required. This will cause the dotted line beginning at point e to become steeper, allowing the household to attain a higher indifference curve if it chooses to invest in risky assets. An increase in α , the level of fluid intelligence, would have a similar effect by increasing the efficiency of investment, thereby reducing its cost. An increase in Social Security wealth would reduce the incentive to invest in risky assets if the increase crowds out retirement saving. Variations in time preference will also influence the choice of whether to invest in risky assets. Households that are more patient than the one depicted in Figure 3, will wish to choose to increase the ratio, C_2/C_1 , which they can accomplish with increased saving. The increase in saving makes more likely that an investment in financial knowledge and purchase of risky assets will improve welfare.

4. Data

The data for our empirical investigations of financial knowledge, its determinants and how it relates to retirement saving and portfolio choice come from two sources: the Cognitive Economics Survey and the American Life Panel. Both employ innovative design elements and interview all or a majority of the sample over the Internet. We will use the data from these two surveys to investigate the empirical support for the theory that we have developed above.

4.1 Cognitive Economics Survey

The Cognitive Economics Survey (CogEcon) is an innovative new survey administered by mail and internet to a national sample of 1,222 persons, age 51 and older

and their spouses regardless of age. It was designed by a team of economists to help understand the cognitive bases of economic decision making.⁴ The CogEcon questionnaire, which has a median length of 53 minutes on the internet version, includes a battery of twenty-five questions on financial sophistication, detailed measures of income, wealth and portfolio allocation plus measures of risk tolerance, self-assessed financial knowledge, use of records other sources of information and several questions on decision making. Respondents were drawn from the pool of participants in the “NCGS+HRS” cognition study led by John J. McArdle. That project conducted an extremely detailed, three hour cognitive assessment of sample members, measuring many components of fluid and crystallized intelligence. In this paper we will combine the information collected in the CogEcon Survey with the cognitive measures from the “NCGS+HRS” cognition study.

4.2 American Life Panel (ALP)

The ALP is an Internet panel of respondents 18 and over with a sample of about 1,200 participants. Respondents in the panel either use their own computer to log on to the Internet or a Web TV, which allows them to access the Internet, using their television and a telephone line. The technology allows respondents who did not have previous Internet access to participate in the panel and furthermore use the Web TVs for browsing the Internet or use email. About once a month, respondents receive an email with a request to visit the ALP URL and fill out questionnaires on the Internet. Typically an

⁴ In addition to Willis, the design team includes Daniel Benjamin, Andrew Caplin, Miles Kimball, Kathleen McGarry, Claudia Sahn, Matthew Shapiro, and Tyler Shumway. Gwen Fisher, Brooke Helppie, Joanne Hsu oversaw the internet and mail data collection and also provided valuable help on the survey design. We also wish to extend our thanks to many other scholars from outside the project made valuable suggestions.

interview will not take more than 30 minutes. Respondents are paid an incentive of about \$20 per thirty minutes of interviewing (and proportionately less if an interview is shorter). The respondents in the ALP are recruited from among individuals age 18 and older who are respondents to the Monthly Survey (MS) of the University of Michigan's Survey Research Center (SRC).⁵ The interviews cover a wide range of topics, including information on income, wealth, health, retirement benefits, expectations, and attitudes. The 12th wave of the ALP administered a large battery of questions on financial literacy to half the sample. 25 of these questions are identical to questions fielded in the CogEcon Survey. The resulting variables along with variables on demographics, income, wealth, and expectations about future Social Security benefits form the core of our empirical analysis. We have observations on financial literacy for 574 respondents.

4.3 Financial literacy questions

Both the ALP and the CogEcon Survey used the same format for administering the financial literacy questions, asking respondents to express how strongly they believe a statement to be true or false. The text box below shows the exact instructions from the CogEcon Survey.

⁵ The MS is the leading consumer sentiments survey that incorporates the long-standing Survey of Consumer Attitudes (SCA) and produces, among others, the widely used Index of Consumer Expectations. Each month, the MS interviews approximately 500 households, of which 300 households are a random-digit-dial (RDD) sample and 200 are re-interviewed from the RDD sample surveyed six months previously. SRC screens MS respondents. It asks MS-respondents age 18 or older if they have Internet access and, if yes, whether they would be willing to participate in Internet surveys (with approximate response categories “no, certainly not,” “probably not,” “maybe,” “probably,” “yes, definitely”). If the response category is not “no, certainly not,” respondents are told that the University of Michigan is undertaking a joint project with RAND. They are asked if they would object to SRC sharing their information about them with RAND so that they could be contacted later and asked if they would be willing to actually participate in an Internet survey. Many MS-respondents are interviewed twice. At the end of the second interview, an attempt is made to convert respondents who refused in the first round. This attempt includes the mention of the fact that participation in follow-up research carries a reward of \$20 for each half-hour interview.

Next we would like to ask you a series of statements about financial matters. We would like to know whether, in your opinion, the statement is generally “**True**” or generally “**False**” and how strongly you believe this to be the case.

An example of a true-false statement is the following:

Example Question: A savings bank never offers a checking account.											
Most Likely False						Most Likely True					
Surely False				Guess False		Guess True		Surely True			
100%	90%	80%	70%	60%	50%	50%	60%	70%	80%	90%	100%

←Please Circle One Number→

If you think that this statement is most likely to be **true**, please choose a number in the right half of the box above. If you think that the statement is surely true, circle “**100%**.” If you think it is only 60% likely to be true, please circle “**60%**.”

Similarly, if you think that this statement is most likely to be **false**, please choose a number in the left half of the box above. If you think that the statement is surely false, circle “**100%**.” If you think it is only 70% likely to be false, please circle “**70%**.” If you are completely unsure and have “no idea” whether the statement is true or false, please make your best possible guess and circle whether you would like to guess true with 50% confidence or guess false with 50% confidence.

The ALP administered a total of 70 questions following this format, 25 of which are identical to the ones administered in the CogEcon Survey. The questions query concepts such as portfolio diversification, compound interest, and institutional knowledge, for example, aspects of how annuities work and so on.

In our analysis we focus on the 25 questions that are identical in the two surveys. Both surveys employed randomizations of the way in which the questions were presented to

respondents to allow checking for and correcting potential biases.⁶ We recode answers on a 12-point scale so that a score of 1 means that the answer to a particular question is false and a score of 12 means that the answer is perfectly correct. We compute respondents' average score and interpret it as a proxy of their financial knowledge. In addition we use respondents' scores on single questions to study the acquisition of financial knowledge in the context of a human capital framework.

5. Empirical Investigation of Financial Knowledge and its

Determinants

The point of departure of this paper is the large heterogeneity in the population with respect to financial knowledge. We start out our empirical investigations by showing evidence of this heterogeneity both in terms of accuracy of respondents' knowledge and in terms of how long it took respondents to answer financial literacy questions. We also provide empirical support for the assumption that people know what they know which we discussed in the model presented above.

Our study of determinants of financial knowledge follows very closely the structure implied by the production function for investment in human capital (equation (4)).

5.1 Heterogeneity in Financial Knowledge in the Population

Using the 25 questions on financial literacy which were administered in both surveys (ALP and the CogEcon Survey) we determine a score of financial sophistication

⁶ Both surveys randomized respondents into several different design scenarios. One randomized whether the scale displayed on the screen went from true to false or from false to true when reading from left to right. In addition, each question that was phrased so that "true" was the correct answer had a second version with wording as closely comparable as possible, but with "false" as the true answer. Which version of a particular question a respondent received was again determined at random. In addition, the ALP randomized the question order.

which is computed as the average score of all answers a respondent gave to the 25 questions.

Table 1 shows summary statistics of this score by individual characteristics in the ALP. We find that women have a lower score than men; high education and high income are also associated with higher scores. Older respondents achieve higher scores than younger ones suggesting that there is accumulation of knowledge over time.

There is also substantial variation in the amount of time that respondents use to answer the financial literacy questions. Computing the average time an individual took per question we find a population average of 17.4 seconds in the ALP.⁷ However, Table 2 shows that the person at the 25th percentile took about half as long as the one at the 75th percentile (12.5 compared to 21.0 seconds). When comparing the timings by individual characteristics the largest differences emerge by age with older respondents taking longer: the 60-90 year olds take 21.1 seconds on average whereas those age 45 or younger take 14.3 seconds on average. Note that older respondents – while taking longer – also achieve greater accuracy on average.

Individuals who are less financially literate could compensate for their lack of knowledge by seeking advice from others. A prerequisite for this to happen is that these individuals have a realistic assessment of their own knowledge so that they are in the position to realize *when* to get advice and whom to ask. The CogEcon Survey asked respondents to rate themselves along three dimensions related to financial sophistication:

⁷ Studying response times in an internet survey needs to take account of the fact that respondents may interrupt their session and even leave their desks. We therefore top coded the response times for each question at 100 seconds before computing the average time per question for each respondent. The results are very similar if we choose higher cut-offs.

Q11. I am pretty good at math.

Q10. I understand the stock market reasonably well.

Q9. I am good at understanding day-to-day financial matters, such as checking accounts, credit-cards, mortgages, installment payments and budgeting.

Respondents were provided a 6-point scale ranging from “strongly agree” to “strongly disagree.” Table 3 shows the cumulative distribution of answers for each of these items.

84 percent of the sample believe (“strongly agree” & “agree”) that they have a good understanding of day-to-day financial matters whereas only 25 percent would say this about their stock market knowledge; 70 percent believe that they are good at math.

To find how these self-assessments relate to an objective measure of financial sophistication we regress respondents’ mean_score attained in the 25 financial literacy questions on demographics (age, education, female, number series and number of economics courses) and on the self-assessments. Understanding the stock market and being good at day-to-day financial matters are strongly significant and are associated with higher financial literacy even when controlling for education and fluid intelligence (Table 4).

Note also in Table 4 the positive coefficient on economics courses indicating that investments in financial knowledge lead to increased understanding of financial matters.

This effect is significant and robust to controlling for education and fluid intelligence (measured by the score on the numbers series).

5.2 Determinants of Financial Knowledge

Our study of determinants of financial knowledge follows closely the structure implied by the production function for investment in human capital in equation (4), discussed above:

$$(13) \quad Q_t = \frac{dK_t}{dt} = \alpha(e_t K_t)^{\beta_1} H_t^{\beta_2} M_t^{\beta_3} E_t^{\beta_4} D^{-\beta_5}; \quad \beta_1 + \beta_2 + \beta_3 + \beta_4 - \beta_5 < 1$$

where

Q_t : amount of additional financial knowledge learned during period t

α : ability parameter (fluid intelligence)

K_t : stock of financial knowledge

e_t : effort

H_t : other forms of human capital

M_t : purchase inputs

E_t : unpriced “environmental” inputs

D_t : difficulty of the question

We proxy the amount of additional financial knowledge learned during period t , Q_t , by the score a person achieves on a particular financial literacy question j . We argue that this is an informative approximation because when decomposing the process that a person goes through in working out the answer to one of the questions it becomes clear

that this process involves many of the same elements that are involved in the acquisition of financial knowledge: A person who answers the j^{th} question of the financial literacy battery combines his or her knowledge, effort and intelligence to find the answer to the question. The score is an indicator of how productive their effort is, given the difficulty of the question measured by D_t . Therefore we propose that a respondent's answer score to the j^{th} question can be interpreted as a measure of Q_t . The time (in seconds) that a respondent spends on answering any particular question will be our proxy-measure for effort e . In addition, the CogEcon Survey assesses individuals' fluid intelligence by means of a number series test.⁸ We will use the score on this test as our measure of ability. Building on our earlier result that people largely appear to “know what they know”, individuals' self-rated knowledge about the stock market and about day-to-day financial matters will be our measure of the *stock* of financial knowledge K_t .

The individual solves the utility maximization problem

$$(14) \quad \max U_i = \sum_{j=1}^{25} [Q_t - \gamma_t e_{ij} - wH(a_t)e_{ij}]$$

trading off the utility of additional financial knowledge (measured as the score on the financial literacy questions) against the disutility of effort e and the opportunity cost of time w . Deriving the optimality conditions one can show that an individual should

⁸ The Number Series test in the W-JIII battery is a measure of fluid intelligence that measures quantitative reasoning (Woodcock & Mather, 2001). This ability involves reasoning with concepts that depend upon mathematical relationships. The task required the respondent to look at a series of numbers with a number missing from the series. The respondent needed to determine the numerical pattern, and then provide the missing number in the series. Answers were scored correct or incorrect for each item, and a standardized score (called a W-score) was computed based on WJ-III standard scoring (Woodcock & Mather, 2001).

allocate effort such that the marginal score per unit effort is equated across items. Further

one can show that $\frac{\partial e_{ij}}{\partial D_j} < 0$, implying that the respondent should spend less time on

difficult items; and $\frac{\partial e_{ij}}{\partial K_i} > 0$ which indicates that the respondents should spend more

time on items where existing knowledge is greater. The effect of fluid intelligence on effort is ambiguous.

To provide some further intuition of how our empirical approach relates to the theoretical model Figure 4 depicts Q_j , the accuracy of the respondent's answer to the j^{th} question, as a function of the time (or effort e) in seconds to respond to question j . Higher effort, that is, spending more time on working out the answer to question j , leads to more accurate answers and with that to a higher score on that question. The production function Q^{low} pertains to an individual for whom the marginal return of spending additional effort or time on the j^{th} question is smaller than for the person with the production function Q^{high} . Differences in the marginal return to effort can be due to differences in individuals' ability a , measured by their number series score. If both individuals have the same preferences for effort (parallel indifference curves in Figure 4) the person with Q^{high} will choose to spend more time (e^{high}) than the person with Q^{low} (e^{low}).

Descriptive analysis of the scores measuring the accuracy of respondents' answers as a function of the time spent on the respective question provides support for this framework. Stacking respondents' scores on all 25 financial literacy questions and the

associated timings Figure 5 (left panel) plots the average score by time to answer. We find a relatively steep positive return of additional time spent up to about 40 seconds. There appears to be no additional return for spending more than 40 seconds, in fact the return becomes negative for time spent in excess of 60 seconds. Dividing the sample into quartiles by number series score, which is our proxy for ability (see right panel in Figure 5), we find very similar slopes for the highest two quartiles even though the line for the 4th quartile reflects a higher level of accuracy of respondents' answers. The curve for the 2nd quartile is steeper for times between 5 and 30 seconds, but largely flat at higher values. The line for the lowest quartile stands out from all others, because it is completely flat indicating no return to additional time spent at any level. This finding is consistent with equation (13) which implies that the productivity curve corresponding to a more difficult question has a flatter slope so that additional time has a lower marginal product measured by gain in test score.

In CogEcon, there is no time limit on the financial literacy questions. However, if satisfaction from getting a better score on a give question trades off against the disutility of additional effort, our model predicts a negative correlation between effort and difficulty, holding ability and knowledge constant.

So far we have discussed the case of identical preferences for effort. However, it is likely that for some people – even conditional on ability – learning about financial matters is more burdensome or painful than for others leading them to have different tastes for how much effort to place on improving their financial knowledge. For

example, a person who already knows a lot about finance (i.e., K_t is high) learning more on the same topic tends to be easier because this person has a lot of background information to relate the new knowledge to. Similarly, somebody who has a personality suited to and a facility of learning by asking others is also likely to find it easier to acquire additional financial knowledge than a person who is afraid of asking and has to work out everything on his own. Figure 6 illustrates how variation in individuals' preferences for effort would lead to different choices of effort or time spent working out a question and therewith to different qualities of answers. Holding ability constant we focus on the production function for Q^{high} . $U(high\ cost\ effort)$ represents the indifference curve for a person who finds it "painful" to apply himself to financial problems, so this person's effort is high cost; $U(low\ cost\ effort)$ represents the indifference curve for a person who does not mind thinking about financial topics. $U(low\ cost\ effort)$ has a flatter slope; its tangential point with the production function Q^{high} indicates that this person would choose to exert more effort than the person with high cost effort.

Over the life course, our model of human capital accumulation will tend to generate correlation of personality characteristics and knowledge in much the way that the Cunha and Heckman's (2007) model leads to correlations of earnings with cognitive and non-cognitive skills during working life. Imagine, for example, two individuals, A and B, who are initially endowed with identical stocks of financial knowledge, ability and all other forms of human capital. Assume that A has a personality trait such that he gets less disutility from effort cost than B and, therefore, chooses to spend more time learning. At the end of the period, financial knowledge and this personality trait will be correlated.

In addition, because human capital is self-productive, the productivity of A's time in learning new financial knowledge will exceed B's in future periods despite the fact that the personality trait has no causal effect on either financial knowledge or learning ability.

We will take advantage of this correlation in estimating the production function of investment in financial knowledge from financial literacy questions and use personality traits as valid instruments for endogenously chosen effort under the assumption that these traits do not directly enter the production function. That is, controlling for ability and the stock of human capital and other variations in equation (13), variations in personality shift the slope of indifference curves, tracing out the production function, as is illustrated in Figure 6. The estimation approach used in the next section follows this approach.

5.2.1 Estimating the production function

We estimate the production function using respondents' answers to each of the 25 financial literacy questions. To that end we stack the data and obtain about 11,400 person-question observations. More specifically we estimate the equation

$$\begin{aligned}
 (15) \quad score_{ij} = & \beta_0 + \beta_1 \ln timetoanswer_{ij} + \beta_2 question\ difficulty_j \\
 & + \beta_3 age_i + \beta_4 female_i + \beta_5 education_i + \beta_6 Number\ Series\ Score_i \\
 & + \beta_7 SRstockmarket_i + \beta_8 SRfinance_i + \varepsilon_{ij}
 \end{aligned}$$

for the i = respondents and $j= 1, \dots, 25$ questions and taking into account in the computation of the standard errors that each respondent contributes up to 25 observations and the potential correlation in the errors that might results from this.⁹

The left hand variable is the score in the CogEcon survey of respondent i on question j from the financial literacy battery. The score depends on an individual's choice of effort, measured by the log of the length of time to answer, which is treated as an endogenous variable and identified by a first stage effort supply function described below. In addition to effort, an individual's score depends on his or her cognitive ability and stock of financial knowledge. Fluid reasoning ability is measured by the "*Number Series Score*," the respondent's score on the number series test. "*SRstockmarket*" and "*SRfinance*" are the self-rated measures of knowledge of the stock market and of day-to-day financial matters that we discussed in section 5.1 which we use to measure the stock of financial knowledge K_i . We also control for the *difficulty* of question j which is defined as 12 (maximum attainable score on question j) minus the population average score on question j obtained by an independent sample of respondents in the ALP. Finally, we include age, sex and years of education in order to control for aspects of ability and knowledge not captured by the direct measures.

The first stage equation in (16) estimates the effort supply function derived from the hypothesis that individuals choose their effort level, balancing the utility they get

⁹ We estimate by 2SLS in Stata accounting for arbitrary correlation among observations pertaining to the same respondent.

from performing well on the test against the disutility of effort, measured by the length of time spent answering each question, and using measures of personality and survey fatigue to identify an effort supply function.

$$\begin{aligned}
 \ln timetoanswer_{ij} = & \gamma_0 + \gamma_1 question\ difficulty_j + \gamma_2 age_i + \gamma_3 female_i + \gamma_4 education_i \\
 (16) \quad & + \gamma_5 NSscore_i + \gamma_6 SRfs_i + \gamma_7 SRsm_i \\
 & + \gamma_8 qorder_j + \gamma_9 extroversion_i + \gamma_{10} no\ records_i + \eta_{ij}
 \end{aligned}$$

Note that the theory implies that the optimal strategy for CogEcon respondents, like that for high school students taking the SAT, is to choose to work on easy questions and spend little time on long questions so that we expect γ_1 to be negative.¹⁰ The Cobb-Douglas form of the production function in (13) implies that more able or more knowledgeable individuals will devote more time to answering questions because these resources raise the marginal product of effort. However, this effect might be offset by a higher opportunity cost of time for such people or because the technology in (13) is misspecified. For example, a person with more education and higher fluid intelligence may be able to determine more quickly whether or not a question is too difficult to be worth thinking about.

¹⁰ On the SAT, students are penalized for guessing and, hence, are advised to skip questions unless they are fairly sure that they know the answers. CogEcon simply presents a series of statements, asking respondents to indicate whether the statement is true or false and how confident they are of their answer by clicking on one of 12 radio buttons. The respondent does not receive feedback on their score and there is no penalty for wrong answers; indeed two of the radio buttons are explicitly labeled “just guessing true” or “just guessing false.” Moreover, the format of the test discourages skipping a question because a respondent who clicks the “next” button without clicking on a radio button is presented with a reminder screen asking them to give an answer. In practice, almost all of the questions are answered by everyone.

The coefficient of log effort in the human capital production function in (15) is identified by the exclusion restriction that three variables—extroversion, no records and question order—do not affect an individual’s capacity to answer a question but do affect the disutility of effort in (16). Extroversion is the “Big 5” personality characteristic that is most correlated with effort.¹¹ Another major component of the CogEcon survey collects detailed data on wealth and portfolio composition. At the beginning of this section, respondents are invited to use information such as income tax records, 401k plan statements or other records to help them give accurate answers. At the end of the wealth section, they are asked whether they used any records. We hypothesize that those who did not use any records revealed a high disutility of effort in answering survey questions. Finally, question order indicates when a given question occurs in the 25-item test and, thus, serves as an indicator of how fatigued (or bored) an individual might be in answering questions of this type.

Table 5 presents the summary statistics of the variables and Table 6 has the estimation results of the human capital production function. We find that log effort, fluid intelligence and self-rated knowledge are all significant and increase the accuracy of the responses. Note that the coefficient on difficulty of the question is negative in line with our discussion above. Including controls for the “Big 5” personality traits does not affect

¹¹ Personality refers to relatively stable characteristics of thought, affect, and behavior: conscientiousness (being goal-directed, organized, and detail-oriented), agreeableness (having a tendency to get along easily with others), extroversion (enjoys social engagement and interacting with others), openness to experience (willing to try new things), and neuroticism (having a tendency to worry a lot). These five characteristics were measured via self report with the 44-item Big Five Inventory (BFI; John & Srivastava, 1999). Participants indicated the extent to which he/she agreed with a series of statements that describe him/herself using a 5-point Likert-type response scale ranging from Strongly Disagree to Strongly Agree.

We have estimated models using all five factors, but this appears to lead to a “weak instrument” problem.

the other estimates much. Extroversion, and openness are significant at the 5%-level; conscience is significant at the 10%-level.

Table 7 shows the results from the first stage regression with effort supply as left-hand variable. Several ones of the instruments are strongly significant. Note also the the estimates associated with difficulty, ability and education: they are all significant at the 5-percent level and exhibit the signs implied by the theory derived above. Higher degree of difficulty of the questions, higher ability of the respondents and more years of education are all estimated to be associated with the respondent spending less time on a particular question.

6. Conclusions

Traditionally economic models assume perfect information. However, we document in that there is important variation in financial knowledge in the population. For example, women score lower than men on a measure of financial sophistication while high education, high income and being older are associated with higher scores. In this paper we develop a theoretical model that treats financial knowledge as a form of human capital and the acquisition of knowledge as an investment. There is an important scale economy in this investment process: while increased knowledge raises the feasible expected return per dollar, the total value of the investment depends on the number of dollars to which the improved return is applied. This creates a non-convexity that helps explain household non-participation in the stock market and also implies a set of comparative static and dynamic implications; for example, that higher levels of fluid intelligence will be related to higher learning rates on financial sophistication items and

to the choice of riskier portfolios. We use this theory as a framework for our empirical work using new data from the internet version of the Cognitive Economics Survey. In particular, we estimate a model in which individuals combine their knowledge, fluid reasoning ability and effort in order to answer each question in a 25-item battery of questions testing their financial sophistication. While their knowledge and ability are predetermined, individuals choose their effort level, balancing the utility they get from performing well on the test against the disutility of effort, measured by the length of time spent answering each question, and using measures of personality and fatigue to identify an effort supply function. We find that the data and the estimation of the human capital production function of financial knowledge lend empirical support to many implications derived from the theoretical model. For example, respondents should spend less time on difficult problems of financial literacy and that is indeed what the estimated effort supply function indicates. Similarly, a larger stock of knowledge should make it easier to learn about new related issues and the estimates once again support this hypothesis.

We illustrate that the model also makes important predictions about portfolio choice and retirement saving. In future research we plan to investigate whether these predictions are supported by the data. For example, we will study to what extent the variation in financial knowledge, effort and cognitive resources explains variation in outcomes that capture preparation for retirement such as retirement savings in excess of Social Security and pensions, portfolio diversification and uncertainty about retirement resources.

References

Bayer, Patrick J, Bernheim, B. Douglas, and Scholz, John Karl. 2008. "The Effects of Financial Education in the Workplace: Evidence from a Survey of Employers." *Review of Economics and Statistics*.

Ben-Porath, Yoram. (1967) The Production of Human Capital and the Life Cycle of Earnings. *The Journal of Political Economy*, Vol. 75, No. 4, Part 1, pp. 352-365

Campbell, John Y., Calvet, Laurent and Sodini, Paolo (2007)"Down or Out: Assessing the Welfare Costs of Household Investment Mistakes", with Laurent Calvet and Paolo Sodini, *Journal of Political Economy* 115:707-747.

Cattell, R.B. (1941). Some theoretical issues in adult intelligence testing. *Psychological Bulletin*, 38, 592.

Cattell, R.B. (1971, 1987). *Intelligence: Its structure, growth and action*. Amsterdam: Elsevier.

Cunha, F. and Heckman, James. , (2007) "The Technology of Skill Formation," *American Economic Review*, 97(2):31-47.

Haliassos, Michalis and Bertaut, Carol (1995), "Why Do So Few Hold Stocks?" *The Economic Journal*. 195, 1110-29.

Horn, J.L. (1991). Measurement of intellectual capabilities: A review of theory. Chapter 7 in McGrew, K. S., Werder, J. K. & Woodcock, R. W. *Woodcock-Johnson Technical Manual* (pp. 197-246). Allen, TX: DLM Teaching Resources.

Horn, J.L. (1998). A basis for research on age differences in cognitive abilities. In J. J. McArdle & R. Woodcock (Eds.), *Human Cognitive Abilities in Theory and Practice* (pp. 25-28). Chicago: Riverside.

Horn, John L. and McArdle, John J. (2007) "Understanding Human Intelligence Since Spearman." In Robert Cudeck and Robert C. MacCallum, eds., *Factor Analysis at 100: Historical Developments and Future Directions* (Mahway, NJ and London, Lawrence Erlbaum Associates, Publishers): 205-247.

Milgrom, Paul (2008). "What the Seller Won't Tell You: Persuasion and Disclosure in Markets." *Journal of Economic Perspectives* 22, pp. 115-131.

Table 1: Financial literacy score by characteristics

	Score			
	Mean	P25	Median	P75
Gender				
Male	9.04	8.28	9.16	9.76
Female	8.33	7.52	8.20	9.12
Age				
<=45	8.30	7.40	8.20	9.24
46-59	8.67	7.88	8.70	9.44
60-90	8.93	8.16	8.94	9.68
Education				
High school	7.94	7.06	7.84	8.78
College	8.69	7.88	8.68	9.48
Grad School	9.05	8.16	9.16	9.84
Income				
Lowest tercile	8.27	7.44	8.16	9.08
2nd	8.77	8.00	8.76	9.64
Highest tercile	9.05	8.32	9.16	9.76

Table 2: Time to answer financial literacy questions by characteristics

	Average Time in seconds			
	Mean	P25	Median	P75
All	17.3	12.5	16.2	21.0
Gender				
Male	17.9	13.2	17.0	21.8
Female	16.8	12.3	15.4	20.3
Age				
<=45	14.3	10.4	13.4	16.8
46-59	16.7	12.3	15.6	19.8
60-90	21.1	16.0	19.7	24.7
Education				
High school	18.3	12.0	16.5	21.7
College	17.2	12.5	16.2	21.0
Grad School	16.7	12.7	15.6	20.3
Income				
Lowest tercile	18.3	13.0	16.8	22.2
2nd	16.8	12.4	16.4	20.5
Highest tercile	16.1	12.3	15.3	18.3

Table 3: Cumulative distribution of responses to self-ratings

	Good at Math	Good at Household Finance	Know Stock Market
Strongly agree	15.49	38.75	3.91
Agree	55.10	84.15	24.66
Somewhat agree	76.08	95.11	53.03
Somewhat disagree	86.86	98.04	65.56
Disagree	97.06	99.61	88.06
Strongly disagree	100.00	100.00	100.00

Table 4: Explanatory power of self-ratings for financial literacy

	Basic specification		Add fluid intelligence		Add self-assessments	
	Coeff.	t-value	Coeff.	t-value	Coeff.	t-value
Age	0.002	0.53	0.009	2.18	0.005	1.20
Education	0.113	5.83	0.072	3.59	0.054	2.95
Economics courses	0.087	3.32	0.076	3.00	0.005	0.19
Female	-0.345	-4.43	-0.257	-3.34	-0.127	-1.76
Number series score			0.012	5.72	0.011	5.48
Good at math					-0.014	-0.44
Self-rated stock market knowledge						0.23
Self-rated day-to-day fin. knowledge						0.15
Constant	-1.714	-4.45	-7.980		-8.459	-7.77
<i>R-squared</i>	0.163		0.215		0.360	
<i>N</i>	479		479		479	

Table 5: Summary statistics

Variable	N	Mean	Standard Deviation
Score on question j	11403	8.730	3.461
time to answer	11475	20.809	15.380
(time to answer) ²	11475	669.536	1204.864
self-rated stock market knowledge	18475	3.215	1.475
self-rated day-to-day fin. knowledge	18700	5.086	1.006
number series score	11450	527.032	19.850
difficulty of question	23075	3.489	0.772
used records = no	23025	0.441	0.496
Personality traits			
Extroversion	22900	3.345	0.823
Agreeable	22900	4.258	0.538
Conscience	22900	4.104	0.613
Neuroticism	22900	2.639	0.634
Openness	22900	3.725	0.711

Table 6: Estimates of the Production Function by 2SLS,
regressing the score measuring rightness on covariates.

	Coefficient	P-value
Log of time to answer	1.858	0.000
Difficulty of question j *	-1.362	0.000
Age	-0.020	0.006
Female	-0.008	0.947
Years of education	0.111	0.000
Number Series score (fluid intelligence)	0.015	0.000
Good at Household finance (self-rated)	0.236	0.000
Good stock-market knowledge (self-rated)	0.280	0.000
Constant	-2.306	0.276

* Difficulty of question j is defined as
(12 – average score for the same question in the ALP survey). A higher value
of the variable “difficulty” reflects higher difficulty of question j .

Note: Standard errors are adjusted for clustering at the person level to account for the
fact that each person contributes up to 25 observations in the question / person data set
for estimation.

Instrumented: Log_e

Instruments: difficulty, age, female, years of edu, NSscore, good HH fin, good SM,
question order, extroversion, no records.

Table 7: First Stage Regression which has the direct interpretation as the Effort Supply Function (left-hand variable: log of time to answer); additional included instruments: question order, extroversion and not using any records to answer questions.

	Coefficient	P-value
Difficulty of question j *	-0.082	0.000
Age	0.013	0.000
Female	-0.057	0.083
Years of education	-0.029	0.001
Number Series score (fluid intelligence)	-0.002	0.049
Good at Household finance (self-rated)	-0.009	0.596
Good stock-market knowledge (self-rated)	0.014	0.215
Question order	-0.011	0.000
Extroversion	-0.065	0.000
No records used	-0.108	0.000
Constant	4.113	0.000

* Difficulty of question j is defined as (12 – average score for the same question in the ALP survey). A higher value of the variable “difficulty” reflects higher difficulty of question j .

Note: Standard errors are adjusted for clustering at the person level to account for the fact that each person contributes up to 25 observations in the question / person data set for estimation.

Test of overidentifying restrictions: Score chi2(2) = .662524 (p-value = 0.7180)

Figure 1. Effect of Financial Knowledge on Portfolio Rate of Return

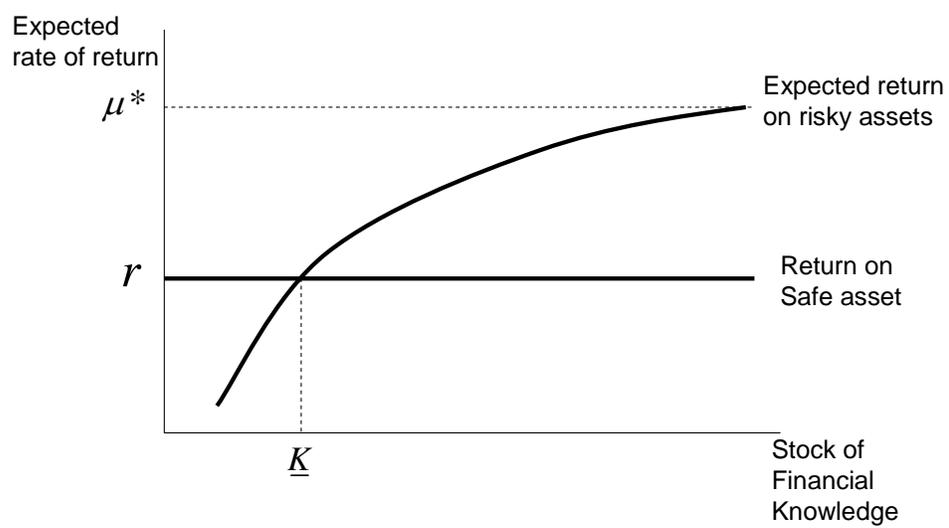


Figure 2a. Effect of Financial of Advice on Portfolio Rate of Return

(Case 1: Increased Advice Reduces Knowledge Threshold for Investment in Risky Asset)

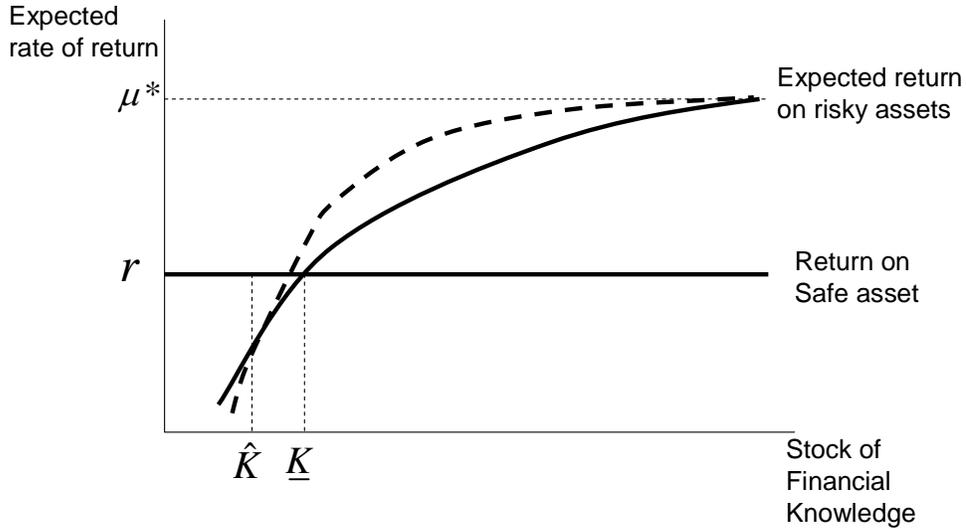


Figure 2b. Effect of Financial of Advice on Portfolio Rate of Return

(Case 2: Increased Advice Raises Knowledge Threshold for Investment in Risky Asset)

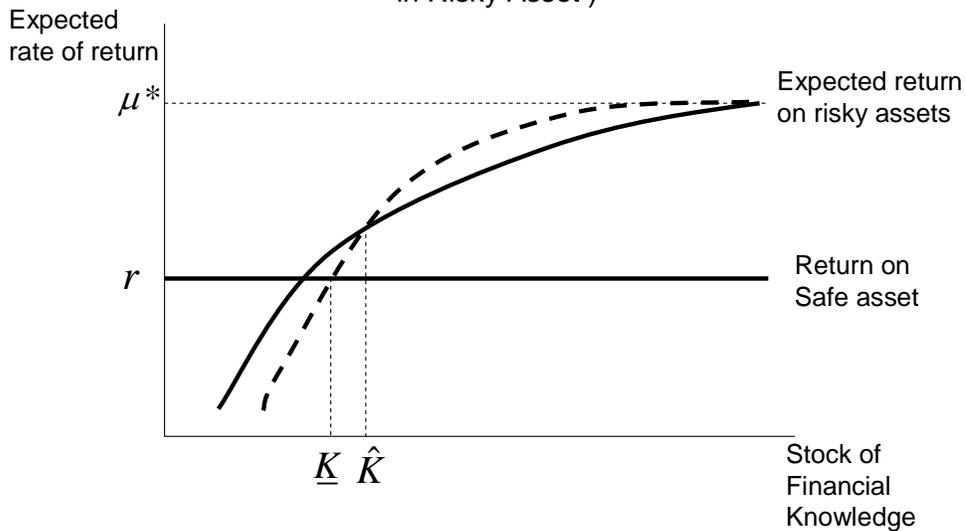


Figure 3. Joint Choice of Financial Knowledge, Savings and Rate of Return

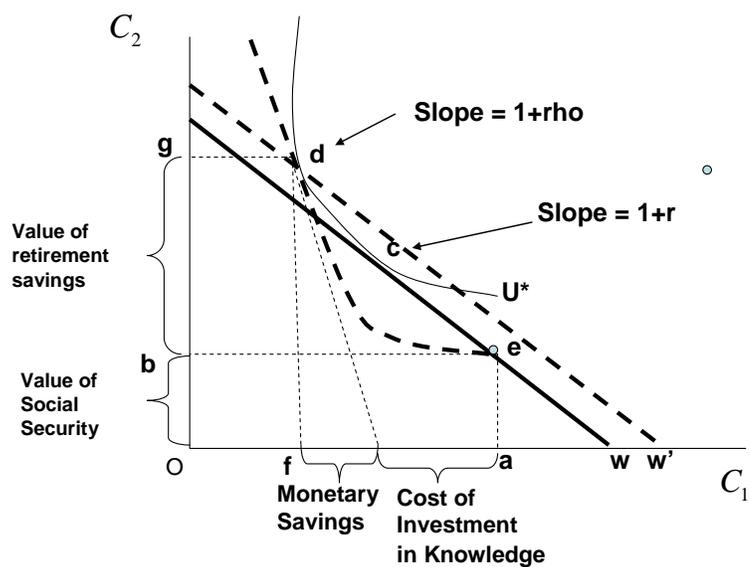


Figure 4:

Production Function for jth Question

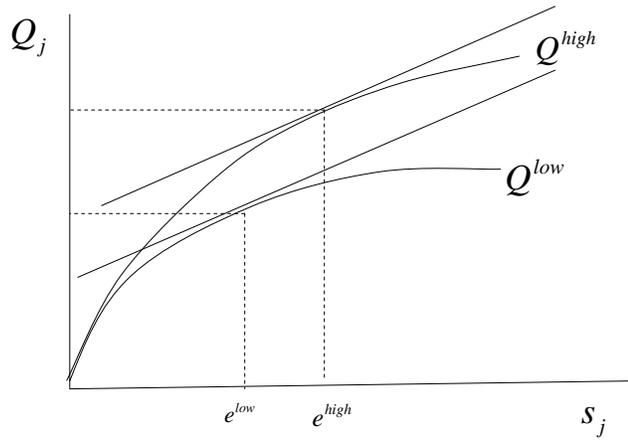


Figure 5: Relationship between time to answer question j and the accuracy of the answer

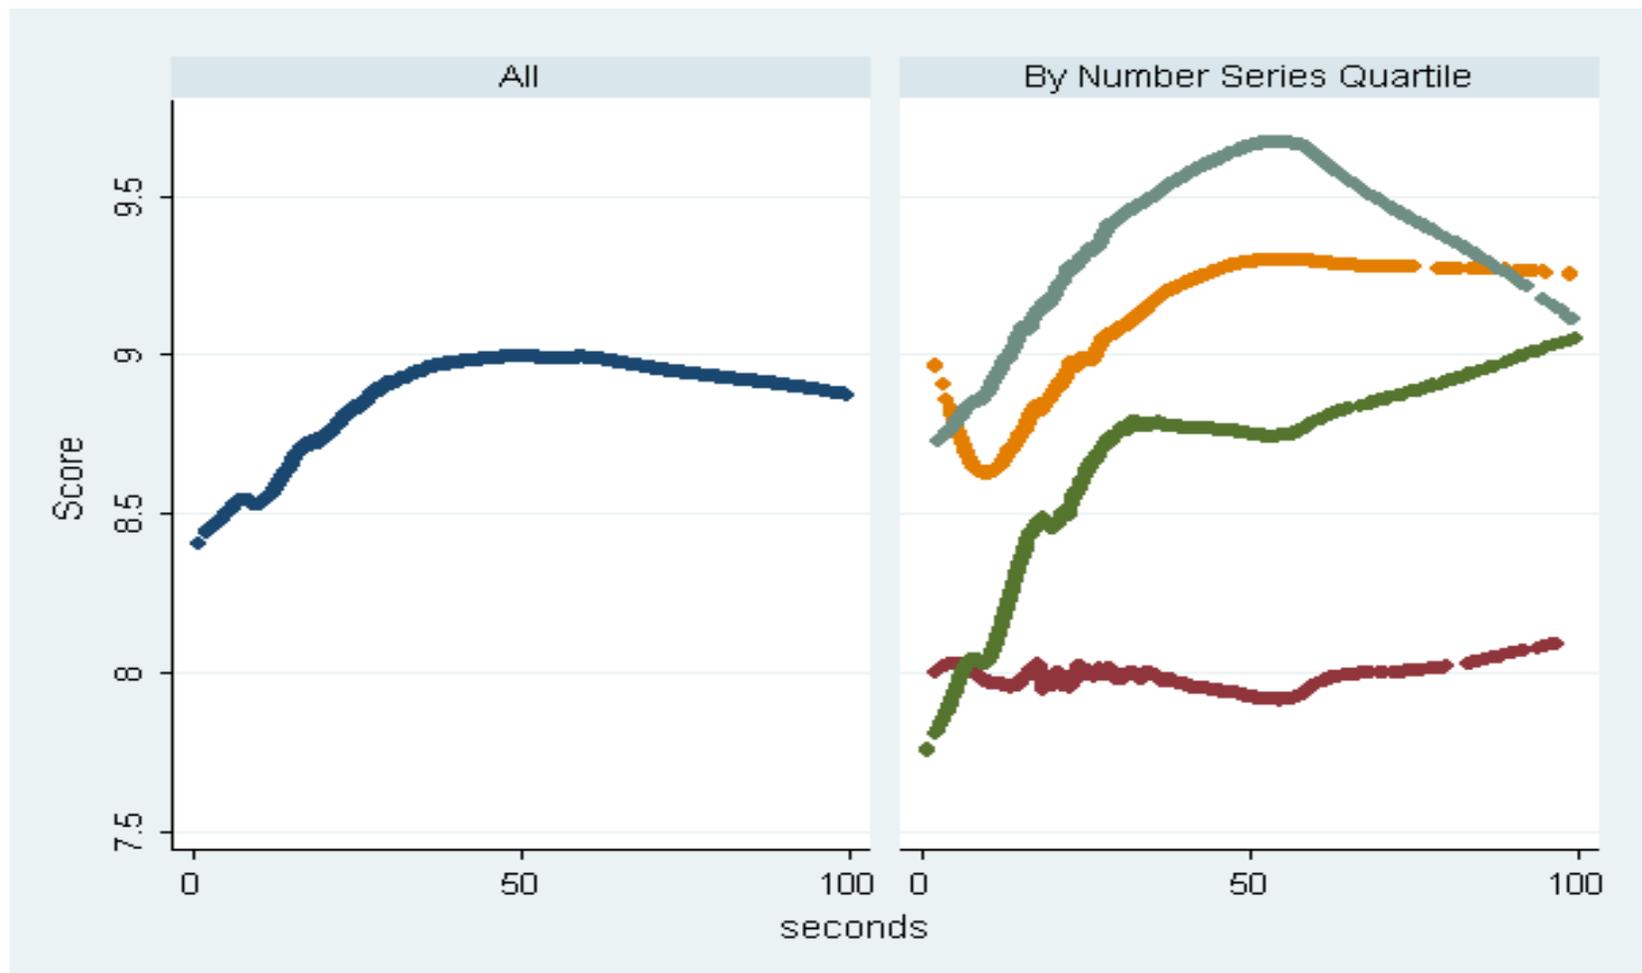


Figure 6. Identification: Variation in Disutility of Effort Due to Personality & Fatigue Trace Out Production Function

