

GMM Estimation of an Asset Pricing Model with Habit Persistence

José L. Fillat and Hugo Garduño*

Department of Economics
The University of Chicago

Abstract

The asset pricing literature has calibrated models with external habits and documented that these models are successful at generating a large set of stylized facts about asset prices. In this paper, we re-consider this evidence by estimating the preference specification using GMM under three different market settings. First we assume complete insurance among all the individuals and estimate the model using aggregate consumption data. Second, we acknowledge that not all the households are actually trading and holding stocks, thus we use household-level data. The third market setting is estimated assuming market incompleteness among stockholders along with limited participation. We find evidence that a complete markets model is better able to explain average returns, whereas a model that includes limited participation of agents in the stock market and incomplete consumption insurance among individuals is better able to explain the equity premium and does so with a lower value of the RRA coefficient than a model with complete markets.

Keywords. Habits, Returns Behavior, Stochastic Discount Factor, Limited Participation, Incomplete Markets.

JEL classification. C12, C13, C52, D52, E21, E44, G12.

1 Introduction

It is well known in modern finance that the original specification of the consumption-based model has not been very successful in explaining the risks that drive prices

*Department of Economics, The University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637; jlfilat@uchicago.edu and hgarduno@uchicago.edu. We thank John Cochrane, George Constantinides, Eugene Fama, James Ellis, Monika Piazzesi for useful comments. We have benefited from comments of the participants at the seminars in the University of Chicago and the ALACDE conference at the University of California, Berkeley. Jose L. Fillat thanks Fundación Rafael del Pino and the Division of Social Sciences of the University of Chicago for financial support. Hugo Garduño is grateful for the financial support of the Fulbright Commission and Conacyt. Usual disclaimers apply.

and expected returns. A popular shortcoming of the model is the Equity Premium Puzzle documented in Mehra and Prescott (1985). This puzzle conveys an intriguing shortcoming of the standard theory: for a class of competitive pure exchange economies in which returns and consumption growth follow stationary processes, sensible values of risk aversion cannot account for the large equity premium (the difference of the return of the market portfolio and the risk-free rate) found in data. In order to match the equity premium, risk aversion should take unusually large values relative to those traditionally used. Even if we concede that investors are very risk averse, problems would not end here, since allowing for large values for risk aversion induces a very large risk-free rate, and data shows that the return on T-bills (a proxy for the risk free rate) is small and has low volatility.

The literature that followed this puzzling fact looked for theoretical innovations to the standard model. These other theories include alternative specification for preferences, incomplete markets, borrowing constraints, limited stock market participation, market imperfections, and modified probability distributions, among others.¹

Our objective in this paper is to estimate and test a consumption-based model with external habit formation following Campbell and Cochrane (1999) in three different circumstances: complete markets, limited stock market participation, and incomplete markets. There are a number of reasons for doing this. On the one hand there are not enough empirical analysis of habits models based on estimation and testing. Most of the empirical work in this literature are calibration exercises, and without discrediting whatsoever their legitimacy, formal econometric techniques would help to better evaluate these models. The classic reference of estimation and testing of the consumption-based model is Hansen and Singleton (1982), but to our knowledge, there are a few papers that have performed estimation of habit-based models.² On the other hand there is evidence of market incompleteness and limited participation of agents in the stock market; hence, adding these features to the habits model might give further insights into the behavior of stock returns.

The standard theory assumes a time and state separable utility function with a constant relative risk aversion coefficient. A natural response for the poor performance of the model is to look for alternative representation for preferences in which this assumption is relaxed. One approach that has also become common is to use a utility function that takes into account the fact that rational individuals behave according to habits. Previous studies have documented the importance of habits in consumption. Deaton (1994) provides a recent overview of evidence of habitual consumption. In this monograph some evidence is shown of positive autocorrelation in consumption series. Heaton (1995) studies the case in an asset pricing framework and finds evidence of habit persistence only if consumption is allowed to be locally substitutable. Henceforth the idea of

¹For a recent summary of these approaches see Mehra and Prescott (2003).

²A noticeable example is Chen and Ludvigson (2004), this paper utilizes a semiparametric approach to estimate the habit, which is a function of current and lagged consumption, so in principle accounts for an internal habit.

including habits in the utility function is appealing both theoretically, to relax the separability assumption, and empirically, given the evidence that favors the presence of habits in consumption.

The influence of past consumption on current behavior is obvious. Past consumption affects marginal utility for present consumption and thus optimal behavior of individuals. As Becker (1992) points out, past consumption can produce both positive and negative effects on current consumption, allowing for intertemporal complementarities or substitution. He also stresses that habit formation can be understood by recognizing that the utility of many goods depends on how present consumption of these goods compare with the amount consumed in the past. That is known as internal habit formation.

Asset pricing has incorporated this idea into a number of papers. A paper worth mentioning is Abel (1990), which introduces a utility function that nests three classes of utility functions: time-separable utility, internal habits, and what he calls “catching up with the Joneses”, which basically corresponds to external habit formation where the utility depends on how far the individual is from past aggregate consumption. When calibrating the model, the values obtained for the equity premium are very small for low values of the preference parameter.

Probably the most notorious example of habit formation in asset pricing is the model developed in Campbell and Cochrane (1999). The aim of this paper is to show how a habit-based model can account for explaining some features observed in the data.³ Habit formation is external, which means that the consumer cares about consumption relative to a habit that does not depend on lagged consumption. In this setup a state variable is included through the definition of the “surplus consumption ratio,” the percentage by which aggregate consumption is above the subsistence level. This surplus consumption ratio adds volatility to the stochastic discount factor beyond that one already induced by consumption growth. This leads to the conclusion that consumers fear stocks because they are likely to do poorly in recessions rather than the resulting risk in consumption or wealth.

The finance literature views Campbell and Cochrane (1999) as a successful model that proposes a pricing kernel that delivers predictability and has been used increasingly in subsequent papers. Campbell and Cochrane (2000) use again their habits specification to explain why the CAPM is a better approximate asset pricing model than the standard consumption based model. Otrok, Ravikumar, and Whiteman (2002) explore how habits model (this one included) solve the equity premium and risk-free rate premium. Matsen (2003) utilized the habit model to analyze the link between international trade in financial assets, economic growth and welfare, and shows that the welfare gain from asset trade is lower with habit persistence.

Chen, Cosimano, and Himonas (2003) show that the integral equation for the price-dividend function yields a unique, bounded, continuous and infinitely

³Time-varying risk premia, the puzzles surrounding the equity premium and the low volatility of the risk-free rate.

differentiable solution. Later Chen, Cosimano, and Himonas (2004) develop an explicit formula for the distribution of U.S. returns that explains many of its observed characteristics.

Menzly, Santos, and Veronesi (2004) use this habit model along with multiple technologies and generate cross-sectional differences in expected returns and long-horizon return predictability. Santos and Veronesi (2005) use the model in order to account for the value premia of expected returns.

Li and Zhong (2005) do an empirical investigation of the predictability and cross-section of returns from the international equity markets using this habit specification. Chue (2005) studies how variations in investor risk aversion can be a potent source of international market co-movements. Pastor and Veronesi (2005) uses this habit model to argue that the number of firm going public changes over time in response to time variation in market conditions.

The increasing use of Campbell and Cochrane (1999) specification gives another reason to estimate and test their model and constitutes one of the contributions of this paper.

The estimations performed in this paper also take into account some other features of the stock market, namely that not all agents in the economy hold stocks. Limited participation is an important feature of the stock market, Mankiw and Zeldes (1991) document this fact, showing some evidence that consumption by stockholders is more volatile and more highly correlated with the market excess return, and this differences might help to explain the size of the equity premium. They also find that the implied coefficient of risk aversion decreases for stockholders.

Some other papers have pointed out the relevance of limited participation. Vissing-Jørgensen (2002b) points out that there are some facts about portfolio choice heterogeneity that may account for returns behavior: a large proportion of households in the U.S. do not hold stocks; large fluctuation of exits and entries into the market; large changes in portfolio shares for equity over time; and heterogeneity in the share of financial wealth invested in stocks.

Vissing-Jørgensen (2002a) finds that accounting for limited participation is important for the estimation of the elasticity intertemporal substitution (EIS). In particular, differences of EIS estimates between stockholders and non-stockholders is high and statistically significant, and this might have effects on the equity premium. Attanasio, Banks, and Tanner (2002) use UK data and show that sensible values for the EIS can be obtained when estimating the Euler equations for the group of likely stockholders, whereas estimates for the whole sample (which includes non-likely stockholders) parameters are unappealing.

Finally we go beyond limited participation. We question the usual assumption of complete consumption insurance; i.e., consumers have a set of assets that ensures them to smooth consumption perfectly across states of the world. That simplifying assumption allows the use of the representative consumer. If we drop that, then the consumption of the representative agent is no longer valid to compute the marginal rate of substitution. On the contrary, we have to compute them one by one for each consumer in the economy (in this case, for each stockholder). Only then we will be able to average the marginal rate

of substitution to obtain the pricing equation.

Complete consumption insurance has also been relaxed in a number of papers although Constantinides and Duffie (1996) revived incomplete markets after a series of empirical difficulties. This paper demonstrates that there exists a specification of the income process that is capable of explain any premium. This model does so without introducing borrowing constraints, short-sale restrictions, transaction costs restrictions on the net supply of bonds, all of which were used in previous non successful papers.

Brav, Constantinides, and Geczy (2002) examine the pricing implications of limited participation and of relaxing the assumption of complete consumption insurance, reinforcing the rationale of using individual marginal rates of substitution for testing purposes in Euler equations instead of the marginal rate of substitution of the representative aggregate consumer. They test Euler equations *a là* Hansen and Singleton (1982) with household data and find evidence supporting the hypothesis of incomplete consumption insurance. When it comes to taking into account limited participation, they find some evidence that a stochastic discount factor driven by a per capita consumption growth can explain the equity premium with a high coefficient of relative risk aversion.

In summary: the original specification of the consumption-based model has shown poor success in explaining the equity premium; there are not enough empirical evaluations of habits models based on estimation and testing, and there is strong evidence that a large fraction of households do not hold stocks. We acknowledge this facts and estimate and test the model of habits in Campbell and Cochrane (1999), taking into account complete markets, incomplete markets, and limited participation.

We proceed as follows: Section 2 explicates the theoretical model. Section 3 presents a description of the data used for the estimations. The fourth section reports the results of the estimations and tests. Section 5 concludes the paper.

2 Theory

This section describes the theoretical framework of this paper. The main objective of an empirical evaluation of an Asset Pricing model is to correctly explain the behavior of stock market returns and ideally have some predictive power. There are different approaches and models explaining returns. One of them relies on factors as explanatory variables for asset returns and uses the beta representation for returns, whereas another approach is the stochastic discount factor (SDF) method, which has become common in the recent literature for econometric evaluation purposes. For a deeper analysis of these two branches and the relationship between them, we refer the reader to Cochrane (2001) and Jagannathan and Wang (2002).

A SDF has the property that the value of a financial asset p_t is equal to the expected value of the payoff (the value of the asset in the next period plus dividends) of such asset times the SDF:

$$p_t = E_t [m_{t+1}(p_{t+1} + d_{t+1})]. \quad (1)$$

This relationship can be transformed in terms of returns and reduced to the expression

$$E_t [m_{t+1}R_{t+1}] = 1, \quad (2)$$

where R_{t+1} is the return of an asset or portfolio. Any random variable m_{t+1} that satisfies (2) is a valid *stochastic discount factor*.

The testable implications of a consumption-based model, like the one used in this paper, are the Euler equations, which can be expressed as in (2). In this case the SDF is the intertemporal marginal rate of substitution. In the standard representative agent model with a CRRA utility function and complete markets, the SDF becomes $m_{t+1} = \delta(C_{t+1}/C_t)^{-\gamma}$. However, as Mehra and Prescott (1985) pointed out, this model does not do a good job in replicating the value of equity premium found in the data.

As mentioned above, this paper tests a model that includes habits in the utility function. Now we describe this model in more detail, both in a complete and incomplete markets framework.

Throughout this paper we use the habits model of Campbell and Cochrane (1999). The basic idea is to include slow-moving external habit to the standard utility function. The habitual level of consumption depends on the history of aggregate consumption rather than on the history of individual past consumption. The utility function for this economy is given by

$$E_0 \left[\sum_{t=0}^{\infty} \delta^t \frac{(C_{i,t} - X_t)^{1-\gamma} - 1}{1-\gamma} \right], \quad (3)$$

where $C_{i,t}$ is the level of consumption for each individual $i = 1, \dots, I$ and X_t is the habit, which can also be interpreted as a time-varying subsistence level. The consumer will maximize utility subject to his or her budget constraint. The relevant relation that is obtained from the optimization problem is as usual the Euler equation. However, the presence of the habit changes the marginal utility of consumption and the intertemporal marginal rate of substitution. This marginal rate of substitution determines the SDF that prices assets, so denoting R_{t+1}^j to the return of asset j , the corresponding conditional Euler equation is given by

$$E_t \left[\delta \left(\frac{C_{i,t+1} - X_{t+1}}{C_{i,t} - X_t} \right)^{-\gamma} R_{t+1}^n \right] = 1. \quad (4)$$

This is true for all individuals i and all assets n . Notice that if the habit were not external, it would have an effect on the future marginal utility and thus in the pricing equation. In a complete markets environment with identical preferences, all individuals choose the same level of consumption, thus $C_{i,t} = C_t$. Campbell and Cochrane (1999) include a state variable dubbed *surplus*

consumption ratio S_t , which is the percentage by which consumption is above the habitual consumption:

$$S_t \equiv \frac{C_t - X_t}{C_t}. \quad (5)$$

With this definition we can rewrite marginal utility in terms of S_t , and thus (4) becomes

$$E_t \left[\delta \left(\frac{C_{t+1} S_{t+1}}{C_t S_t} \right)^{-\gamma} R_{t+1}^n \right] = 1. \quad (6)$$

Notice also that the coefficient of relative risk aversion will be in terms of this variable:

$$\eta_t \equiv -C_t \frac{u''(C_t)}{u'(C_t)} = \frac{\gamma}{S_t}. \quad (7)$$

Hence this is a model that distinguishes risk aversion from intertemporal substitution through non time-separable preferences. The surplus consumption ratio will determine the risk aversion and the curvature of the utility function; in particular, risk aversion is no longer constant and increases when consumption is close to the subsistence level X_t . Periods in which consumption is close to the habit (i.e., when S_t is low) are identified with periods of recession, since the agent's consumption is closer to the subsistence level. When this is the case, risk aversion rises; hence, periods of recession are related to periods with high risk aversion.

In this model, consumption growth is assumed to follow an i.i.d. lognormal process

$$c_{t+1} - c_t = g + v_{t+1}, \quad (8)$$

where $c_t \equiv \log C_t$ and $v_t \sim N(0, \sigma^2)$. There should also be an underlying evolution of the surplus consumption ratio. In this case Campbell and Cochrane (1999) assume that $s_t \equiv \log(S_t)$ follows an auto-regressive heteroscedastic process in the following fashion:

$$s_{t+1} = (1 - \phi)\bar{s} + \phi s_t + \lambda(s_t)(c_{t+1} - c_t - g). \quad (9)$$

The sole inclusion of the state variable is not enough to account for risk premia. A high surplus consumption ratio decreases the marginal utility of consumption between periods. This will encourage the agent to increase consumption today by borrowing more. This would lead to an increase in the risk-free interest rate. Since they want to keep this return constant, they introduce a sensitivity function $\lambda(s_t)$, which controls the response of the surplus consumption ratio to random changes in consumption growth. Given this, the next step is to choose an appropriate sensitivity function. This was done according to three conditions: first, a constant risk free-rate. This is an important feature of the model. As was mentioned previously, former models predict volatile risk-free rates, and data show that these returns are relatively constant. Second,

the habit is predetermined near the steady state \bar{s} . And third, habit moves non-negatively with consumption everywhere.

A sensitivity function that satisfies these conditions is as follows:

$$\lambda(s_t) = \begin{cases} \frac{1}{\bar{S}} \sqrt{1 - 2(s_t - \bar{s})} - 1 & s_t < s_{max}, \\ 0 & s_t \geq s_{max}, \end{cases} \quad (10)$$

where s_{max} is the value for s_t that makes the first part of (10) equal to zero:

$$s_{max} = \bar{s} + \frac{1}{2}(1 - \bar{S}^2), \quad (11)$$

and the steady-state value for the surplus consumption ratio is given by

$$\bar{S} = \sigma \sqrt{\frac{\gamma}{1 - \phi}}. \quad (12)$$

In a nutshell, the standard model was modified to include external habit formation. A state variable was included as the driving force that prices assets. This variable can be interpreted as a recession variable. People fear stocks because they do poorly on recessions. As S_t decreases, consumption decreases towards the habit or subsistence level. This recession also has effects on risk aversion: the lower S_t , the higher risk aversion is. Hence, the SDF that we use in the estimations for the complete markets case is

$$M_{t+1}^{CM} = \delta \left[\frac{C_{t+1}}{C_t} \frac{S_{t+1}}{S_t} \right]^{-\gamma}. \quad (13)$$

In Section 4.2 we explore the results of the first modification of the benchmark model. We acknowledge that only a fraction of consumers actually hold stocks. Therefore, we may assume that the behavior of those with no stocks does not affect the price, since they are not marginal in the decision of devoting \$1 to invest in a risky asset or consume it immediately. Those who are not in the stock market (for any reason, which we will leave aside for the moment) do not face this decision. So they are not relevant in the pricing process of the risky assets.

For this purpose, the model remains the same. The only consideration is that now we are aggregating consumption only among stockholders. Thus, the model explained above is still valid. The only change is that now we will have a representative agent of those individuals who actually hold assets. The Euler equation to be used for the estimation is still (6), where the consumption is the average consumption of the agents participating in the stock market, according to the different definitions of stockholders: those who declared more than \$1 in assets, more than \$10,000, \$20,000, \$30,000, and \$40,000. In this case, we use data from the Consumer Expenditure Survey to be able to aggregate consumption only among those who hold stocks.⁴

⁴Technically, who hold a variety of assets. The question addressed in Section 3.

The formation of the sample moments out of the population moment in (6), including the limited participation feature, is discussed here. As mentioned in the introduction, there is evidence that not all consumers possess stocks, either directly or indirectly, and that the consumption pattern differs between those who hold stocks and those who do not. It is important to note that the assumption of complete markets among stock holders is still valid. The existence of a complete set of markets allows households to insure against idiosyncratic shocks affecting their income. Then the different households are able to equalize their marginal rates of substitution in each state of the world. The pricing implications of using an average of individual marginal rates of substitution must be the same as using the marginal rate of substitution of the representative household. The agents that do not have access to stock market are partialled out of the model exogenously, and their market incompleteness does not affect the pricing equation of the agents who are marginal. In this section of the paper, we use the aggregation result among stock holders, since they can still completely insure their consumption stream with their access to complete financial markets. We compute the consumption of the representative-consumer economy. The SDF that we use in the estimation of the model with limited participation is

$$M_{t+1}^{LP} = \delta \left[g_{t+1} \frac{S_{t+1}}{S_t} \right]^{-\gamma}, \quad (14)$$

where $g_{t+1} = \frac{C_{t+1}}{C_t} = \frac{\frac{1}{I_{t+1}} \sum_{i=1}^{I_{t+1}} C_{i,t+1}}{\frac{1}{I_t} \sum_{i=1}^{I_t} C_{i,t}}$ and the consumption is aggregated among the households belonging to the group of “asset holders”.

Finally, the third part of the paper, presented in section 4.3, comprises the study of market incompleteness among stockholders. We keep the assumption of limited participation and then question whether we can aggregate among stockholders so that we can use the “representative stockholder” or not. Each household’s marginal rate of substitution is a valid stochastic discount factor as well. But since we have short time series of each household, we can mitigate some observation error by using as a valid SDF an equally weighted average of the households’ marginal rates of substitution, i.e.,

$$M_{t+1}^{IM} = \delta \frac{1}{I_t} \sum_{i=1}^{I_t} \left[\frac{C_{i,t+1}}{C_{i,t}} \frac{S_{t+1}}{S_t} \right]^{-\gamma}. \quad (15)$$

Tests of the SDF given by (14) against the SDF given by (15) are tests of the hypothesis of complete consumption insurance among stockholders against the alternative hypothesis of incomplete consumption insurance among stockholders.

3 Data

3.1 Returns data

This section briefly describes the data used in this paper. All the observations are in monthly frequency. As previously mentioned above, the goal is to estimate four preference parameters (γ , ϕ , S_0 , and δ) and test the fit of the habits model in three different setups. Identification requires us to have at least the same number of moments and parameters. In this case we decided to have an overidentified problem and include nine returns and proceed with the GMM estimation as described in Hansen (1982) in order to perform estimation and testing jointly. The returns used in the estimations are the excess return on the market, which is the value-weighted return on all NYSE, AMEX, and NASDAQ stocks minus the risk free rate; six portfolios sorted by size (market-equity) and the ratio of book-equity to market-equity; and two bonds with different maturities (one and five years). The returns on the bonds were obtained from the Federal Reserve Bank of St. Louis (FRED II), whereas the rest of the returns were downloaded from a website maintained by Kenneth French. All returns were deflated using a seasonally adjusted composed Consumer Price Index for nondurable goods and services.

The rationale behind the election of these portfolios is the following. On one hand, an important question is whether this model does a good job in predicting the equity premium; hence, a natural candidate for the estimation is the return on the market. On the other hand, both size and book-to-market factors have shown to be good explanatory variables of stock return behavior (see Fama and French (1993)). Finally, we wanted to include returns on bonds to complement the stocks returns and include different maturities, in case we wanted to analyze the term premium.⁵

3.2 Consumption Data

First we perform estimations of the benchmark model with complete markets. The consumption per capita series was constructed by aggregation of seasonally adjusted real personal consumption expenditures of nondurables and services as well as the population series, all of which were obtained from records of the Bureau of Economic Analysis and the Census Bureau of the U.S. Department of Commerce. The size of the sample in this case is the biggest, such that data were available for consumption, population, returns, and inflation ranging from January 1959 to December 2003.

In order to acknowledge the limited participation of stock holders and incomplete markets, we turn to household-level data. The database we naturally chose for this goal is the Consumer Expenditure Survey (CEX). This survey is a rotating sample provided by the Bureau of Labor Statistics (BLS). Each

⁵We carried out other estimations and tests with different assets like the 25 Fama French portfolios or the Fama French factors. However, the qualitative results are better with the returns reported in this paper and are not significantly different.

quarter, approximately 5,000 U.S. households⁶ are surveyed. Each household is interviewed for five consecutive quarters, although the first interview is not in the published data, because it is considered a training interview. Then it is replaced randomly by another suitable household. The available sample spans from the first quarter of 1980 to 2002. The main purpose of this survey is to compute weights of the CPI, rather than providing data to analyze consumption. Many problems affect this survey. Among others, there is substantial evidence that by aggregating CEX data, the resulting figures are far from close to those provided by the National Income and Product Accounts, Personal Consumption Expenditure data, resulting in declining consumption and some other rare behavior (see McCarthy et al. (2002)). These differences can be partly explained by differences in definition and coverage, but it is hard to explain the huge underestimation that the CEX has for national aggregates (around 35%). Furthermore, Battistin (2003) shows that the distance between aggregated CEX and the National Income and Product Accounts (NIPA) has increased considerably during the second part of the 1990s. That is of particular concern for this paper, because we are investigating the model with both aggregate data from the NIPA and from the CEX, and so it is an issue to keep in mind when drawing conclusions regarding the two parts of the estimation. To fulfill the main objective of computing weights for the CPI, the survey is divided in two separate surveys; one is based on retrospective interviews (Interview Sample, IS) and one based on weekly diaries (Diary Sample, DS). We use the IS.

The Panel Survey of Income Dynamics (PSID) is also a source of household-level data. We use the CEX instead of the PSID, because the latter only has data on food consumption. Besides, we need data on the asset position of each household, questions not asked in the PSID.

Consumption is defined in this paper as expenditures of non-durables and services. We have excluded from the definition of consumption durable goods, health, education, and housing rent. Durable goods are commonly disregarded in the literature when computing consumption, since they are not directly linked to consumption in the households. Health and education are excluded due to the investment component in this type of expenditure. And health expenditures consider only the payments that the consumer makes “out of the pocket.” Housing rents are also excluded, because we do not have an alternative method to impute rents to home owners.

As mentioned above, we used monthly per capita consumption of non-durables and services for each household. First we aggregated across detailed UCC⁷ data for non-durables and services. We excluded durables, health expenses, education, and housing as explained above. We transform nominal consumption into real consumption by deflating with the CPI for non-durables

⁶Households are defined as “all members of a particular housing unit who are related by blood, marriage, adoption, or some other legal arrangement,” Bureau of Labor Statistics (2003).

⁷An extensive list of the detailed Universal Classification Codes used for the aggregation is available upon request. We preferred not to include the list in the paper, due to its size.

and services⁸. To get the per capita consumption, we divided by the number of household members.

The sample that we have used goes from 1981:11 to 2001:11. We dropped the households surveyed in year 1980 and 1981, due to the questionable quality of the interview during those years (see Attanasio and G.Weber (1995)). In total, we have 241 observations ($T = 241$). We also exclude non-responders to the last interview, provided that data about asset holding is answered in that last interview. They are excluded for all the estimations, even those involving no asset holdings. Only data about urban households have been taken into account. This is to minimize the observation error, since the non-marketed items corresponding to household production of rural households may bias downward consumption, and also because they were not interviewed for several years during the first issues of the survey.

Also, to avoid measurement error, we dropped the households who did not answer at least three interviews; i.e. with less than 9 months of reported consumption. If at any month the household reported zero consumption, all the observations corresponding to that household were dropped as well, to avoid observation error from the suspicious reports.

Next a series of filters to consumption was applied in order to minimize the effects of reporting errors and observation errors. These filters target abnormal behavior of consumption growth and consumption. They follow closely those filters applied in Brav, Constantinides, and Geczy (2002).

First we deleted observations if the household consumption growth is bigger than $5^{1/3}$ or smaller than $5^{-1/3}$. This is slightly more restrictive than the one in Brav, Constantinides, and Geczy (2002), because it is two-sided and it disregards both too much growth and too much decline. The observation is also removed if two consecutive household growth rates are too uneven; i.e. if $c_{i,t}/c_{i,t-1} < 2^{-1/3}$ and $c_{i,t+1}/c_{i,t} > 2^{1/3}$ or similarly $c_{i,t}/c_{i,t-1} > 2^{1/3}$ and $c_{i,t+1}/c_{i,t} < 2^{-1/3}$. As they mention in their work, the resulting sample is substantially smaller than the original raw sample of the survey. The initial total number of usable households in the whole sample was 153,575. After the filters are applied, the resulting sample is of 77,688.

Aggregate consumption growth is seasonally adjusted following a simple procedure. We compute the per capita consumption growth of the economy as the average of all individual consumption growth rates, then regress this per capita consumption growth on dummy monthly variables. Then we subtracted the seasonal component given by the adjusted values from the actual values. A similar procedure is used for the aggregate consumption level. On the other hand, the seasonal adjustment of individual consumption is trickier and even questionable. We do not discuss the profound philosophy for adjusting for seasonal variation a time series. But we acknowledge that it is even more controversial to do so if we do not have a long time series of individual observation. Since we are using individual data to obtain individual marginal rates of substitution, we need

⁸Available at the DRI-Basic Economics database available at Wharton Research Data Services through the Center of Research in Security Prices at the University of Chicago.

individual data that is seasonally adjusted. For this, we compute a seasonal component for each group of agents, defining each group according to the level of asset holdings. Once this is done, we subtract the seasonal component to each household, according to the group they belong.

Something worth mentioning about difficulties when using the CEX is the change of IDs between the last interview of 1985 and the first one of 1986. IDs were reset to 0, making it practically impossible to match households who were interviewed at the end of 1985 and also at the beginning of 1986. Therefore, there is a substantial loss of observations in the last part of the 1985 data, since all households who did not answer the last interview, supposedly done in 1986, were dropped. The filter that removes all the households which reported for less than three consecutive quarters contributes to the loss of observations. Something similar occurred in the last month of 1996 with another considerable loss of observations, as is depicted in Figure 7.

3.3 Asset Holdings Data

Once per-capita consumption growth is constructed, we turn to the asset holding data. We used the difference between the variables SECESTX and COMPSECX from the CEX data set. SECESTX is defined as the “estimated value of securities, such as stocks, mutual funds, private bonds, government bonds or Treasury notes owned by the household head (or any members of his CU) on the last day of (last month),” and COMPSECX is defined as “the difference in the estimated market value of all stocks, bonds, mutual fund, private bonds, government bonds or Treasury notes on the last day of the last month compared with a year ago last month.” The new variable is the financial wealth of the household; if missing, we treated them as 0. We also deflate this variable using the same CPI index used for the consumption series, determining wealth in 1983 U.S. dollars.

We used several thresholds for the definition of asset holders and have repeated the estimation for each of the groups. The thresholds are households with more than \$1, more than \$10,000, \$20,000, \$30,000, and \$40,000 invested in securities, bonds, mutual funds, private bonds, government bonds, or Treasury notes.

4 Estimation and Results

We begin the estimation by testing the hypothesis that the marginal rate of substitution of the representative consumer in equation (13) is a valid SDF. Basically, we estimate the model using the Generalized Method of Moments presented in Hansen (1982). We constructed nine moments; one per each testing return presented in Section 3:

$$g_T^n(\beta)^{cm} = \frac{1}{T} \sum_{t=1}^T \delta \left(\frac{S_t}{S_{t-1}} \frac{C_t}{C_{t-1}} \right)^{-\gamma} R_t^n - 1 \quad \text{for } n=1,2,\dots,9. \quad (16)$$

Here, R_t^n is the return at time t of asset n . Therefore, g_T will be a 9×1 vector; that is, it will have one sample moment per each asset return used in the estimation. The GMM estimation will minimize the quadratic form with the parameters as arguments, so defining $\beta = [\phi, \delta, \gamma, s_0]$, we have the following program:

$$\min_{\beta} g_T(\beta)' W g_T(\beta). \quad (17)$$

We programmed a function that computes the moment, with the parameters as variables. Inside the function is generated the process for the surplus consumption ratio. This is the case because this variable is not part of the data; it is a process generated from the data on consumption. This process depends on the four parameters to be estimated. A consequence of this fact is the unevenness of the objective function in (17). Linearizing the function in order to facilitate the estimation would be a step backward, since the power of the model relies on the fact that the habit is adjusting to the history of consumption non-linearly. The main objective in the literature of the equity premium puzzle is basically to increase the volatility of the stochastic discount factor. Including the surplus consumption ratio as a linear transformation of the consumption would not help further that goal, as we see below in plots regarding the volatility of the SDF. For this reason we keep the model as in the original paper and estimate it directly by GMM without any linear-estimation procedure.

The original referenced paper, Campbell and Cochrane (1999), is an exercise of in calibration. The set of parameters that they used is presented in Table 1. They assume that there is complete insurance among consumers and therefore simulate the model with the representative consumer, making no difference between participants or non-participants in the market. Conversely, we study three different situations (complete markets, limited participation, and limited participation with incomplete markets) with actual data instead of generated data, so that we can estimate the model instead of calibrate it.

Table 1: Parameter Choices in Campbell and Cochrane (1999)

Parameter	Variable	Value
Persistence Coefficient	ϕ	0.988
Subjective Discount Factor	δ	0.99
Utility Curvature	γ	2.00
Steady-State Surplus Consumption Ratio	\bar{S}	0.057

The big issue of the estimation is which weighting matrix to use. We do not extensively discuss GMM, but some comments should be useful concerning present results with the identity matrix and the efficient spectral matrix. The objective of using the efficient matrix, given by the spectral density of the sample moments computed in the first stage of GMM, is to maximize the asymptotic information in the sample about a model. The “danger” of using such a matrix is that it may blow up standard errors rather than improve pricing errors.

The efficient matrix will focus on linear combinations of returns that have low variance. Therefore, it may ignore the value premium and the size effect if they are hard to price in terms of variability. With a pre-specified weighting matrix, we are giving up asymptotic efficiency but still obtaining consistent and more robust estimations. Nevertheless, we compute the spectral density matrix in order to have the *correct* variances of the estimates and the moments according to (18) and (19).

Hansen and Jagannathan (1997) propose the use of the second-moment matrix of the payoffs $W = E(RR')^{-1}$ as opposed to the spectral density matrix S . They point out that this matrix may be of interest because the minimized GMM loss function can be interpreted as the distance between the estimated SDF and the SDF that prices all assets (the true one). Another good characteristic of this weighting matrix is that it is invariant to the initial choice of assets, as the spectral density matrix. The problem is that it shares some other characteristics of it, like also being nearly singular, and that it might overweight the moments corresponding to portfolios with smaller variance.

The correct variances of the parameters and the moments when using a pre-specified weighting matrix are

$$\text{var}(\hat{\beta}) = \frac{1}{T}(d'Wd)^{-1}d'WSWd(d'Wd)^{-1} \quad (18)$$

$$\text{var}(g_T) = \frac{1}{T}(I - d(d'Wd)^{-1}d'W)S(I - Wd(d'Wd)^{-1}d') \quad (19)$$

The spectral density matrix S has been computed by the Newey-West estimator,

$$S = \sum_{j=-\infty}^{\infty} E[g_{T_t}(\beta)g_{T_{t-j}}(\beta)'], \quad (20)$$

with one lag, so the actual matrix used for the computations of the variance, both of the parameters and the moments, was

$$\hat{S} = 2E(g_t g_t') + E(g_t g_{t-1}') + E(g_{t-1} g_t'). \quad (21)$$

When taking into account the limited-participation issue, we use the same procedure explained above. What changes slightly is the SDF used in the estimation. While for the benchmark model we used the marginal rate of substitution of the representative consumer for the whole economy, in this case we use the marginal rate of substitution of the representative consumer for those who actually hold assets. For that matter and as detailed in Section 3, we aggregate the consumption data for the households who have a positive asset position according to different thresholds. Therefore, the SDF we are testing is (14). The set of moment conditions according to the use of this SDF is given by (16), with the difference being that the aggregation of the consumption is computed only among those who enter in the definition of stockholders using household-level data from the CEX. The moment conditions are

$$g_T^n(\beta)^{lp} = \frac{1}{T} \sum_{t=1}^T \delta \left(\frac{S_t}{S_{t-1}} \frac{C_t^*}{C_{t-1}^*} \right)^{-\gamma} R_t^n - 1 \quad \text{for } n=1,2,\dots,9, \quad (22)$$

where the asterisk stands for aggregate consumption only among stockholders.

Finally, to test this last assumption of limited participation against the market completeness among stockholders, we estimate the model under the incomplete markets assumption, as detailed above. We assume that the SDF constructed by the marginal rate of substitution of a representative agent is not valid, so we use the marginal rate of substitution of every individual and average them for each period of time. Thus, we use the SDF in (15) to construct the nine moment conditions:

$$g_T^n(\beta)^{im} = \frac{1}{T} \sum_{t=1}^T \delta \frac{1}{I_t} \sum_{i=1}^{I_t} \left[\frac{C_{i,t+1}}{C_{i,t}} \frac{S_{t+1}}{S_t} \right]^{-\gamma} R_t^n - 1 \quad \text{for } n=1,2,\dots,9. \quad (23)$$

4.1 Complete Markets

This section reports the estimations and tests for the habits model in a complete markets setup. The corresponding SDF in this section is the one in equation (13). Two different results are presented. The first uses a pre-specified weighting matrix in the GMM minimization program, namely the identity. The second uses the efficient matrix to weight the moments that produces minimum variance estimates.

Why use a pre-specified matrix? Consider the efficient matrix used to weight the moments, and also consider its Choleski decomposition $S^{-1} = C'C$. Hence, the GMM problem can be re-written as follows

$$\min_{\beta} (Cg_T(\beta))' (Cg_T(\beta)). \quad (24)$$

This illustrates a few points. First, $Cg_T(\theta)$ is a linear combination of the moments, and GMM is trying to minimize this linear combination. As a result GMM, will price according to this rule, and we can be concerned about economically interesting moments rather than portfolios formed out of some other portfolios. Second, C weights each moment and combinations of them. With large matrices, these weights can be extreme, so it might be sensible to use a pre-specified matrix when a large number of assets is used. Third, efficient GMM will pay attention to assets with small return variance; then the model would be evaluated in its ability to price the minimum variance portfolio.

Table 2 summarizes the estimation results when we are weighting the moments equally. Notice that the first important outcome is that there is statistical evidence that the Euler equation does not hold. The J_T statistic takes a value of more than 44. This generates a negligible p-value, which in turn leads to the conclusion that the null hypothesis (that the moments equal zero, or equivalently, zero pricing errors) cannot be accepted.

Table 2: First-Stage GMM

γ	ϕ	S_0	δ	J_T	p-value
4.8669	0.9166	0.0020	0.8190	44.8	0.0000
(7.7644)	(0.1635)	(0.0277)	(0.2011)		

* Standard errors in parentheses.

With this in mind, we can also observe that the estimates of the parameters show interesting results. The autoregressive coefficient takes a value slightly bigger than 0.91 which is an estimate consistent with the prediction of the model. Notice that the subjective discount factor is relatively low and the initial condition of the surplus consumption ratio is 0.002 (so its log is about -6.21). On the other hand, the estimate of the power of the utility function is near 4.9; remember that this is not the relative risk-aversion coefficient. The surplus consumption ratio drives the level of risk aversion and its variation over time. With exception of γ , the standard errors show that the other parameters are significant. In the case of the power of the utility function, the standard error is near 8.

One of the purposes of the model is to account for risk premia. Looking at the predicted excess returns as a function of γ will then give information about whether this model is able to explain this fact. Table 2 shows that the standard errors for γ are just below 8. Figure 1 plots excess returns for the market portfolio against γ . If the model is predicting the equity premium correctly, then the range of γ around its standard errors ($\gamma \pm \text{s.e.}(\gamma)$) would match those of the expected predicted returns by its standard errors $E(R^e) \pm \sigma(R^e)/\sqrt{T}$. The figure shows the optimal value of γ and the upper part of the interval bounded by its standard error of about 12.6. The results show, however, that the predicted return for γ is lower than the actual one. Furthermore, the interval of predicted returns corresponding to $\gamma \pm \text{s.e.}(\gamma)$ does not match the actual excess returns around its standard errors.

In general it is desirable to have efficient estimators of the parameters of a given model. Econometrically, using a consistent matrix is sufficient to use the distribution theory; of course the standard errors of the parameters will not be efficient. There are some reasons in terms of robustness or economic sense to use a pre-specified matrix. There is trade-off when using first-stage estimates. Even if they may give up some efficiency, they are still consistent and can be more robust for economic interpretation.

For the rest of the paper, we focus on efficient estimations of the parameters. The reason is that the estimations do not vary too much; thus, we do have some gains in efficiency and we do not have to pay attention to the possibility of model mis-specification. Table 3 shows the results for the estimation of the parameters.




Figure 1: Expected Returns and γ

Table 3: Second-Stage GMM

γ	ϕ	S_0	δ	J_T	p-value
4.7524	0.9370	0.0011	0.8000	71.3	0.0000
(7.7583)	(0.1201)	(0.0128)	(0.1239)		

* Standard errors in parentheses.

As expected, we have a gain in the accuracy of the estimation with lower standard errors. The power of the utility function γ is estimated at 4.7524. This and the surplus consumption ratio will determine the time-varying relative risk-aversion coefficient. This process is generated with an initial condition of 0.0011 and an autoregressive coefficient ϕ above 0.93. The model predicts that the surplus consumption ratio should be a very persistent process in order to get little variation in the risk-free rate. The estimations yielded this result. The subjective discount factor is 0.8, a small number compared to the values used in the literature. The table also shows how all the parameters besides γ have small standard errors that make them significant. This is an annoying result, because risk aversion is an important matter when analyzing the equity premium.

Notice also that the model is still rejected. The J -statistic is above 70, which leads to the conclusion that there is no statistical evidence that the pricing errors are equal to zero, or that the stochastic discount factor prices assets. Subsequent sections confirm these results, regardless of the assumptions of market completeness of limited participation.

This fact deserves further comments. For many years, empirical work in finance has devoted efforts to develop models, test them, and usually reject them. Rather than exclusively focus on the rejection of the model and wonder why it does not work, we have to analyze in what dimensions the model *do* work. Some other research has also gone beyond rejection and characterize the properties that a SDF should have.⁹ For example, we now know that the SDF should be very volatile, and the risk free rate should be very steady. Remember that in the benchmark model the volatility of the SDF was given by aggregate consumption growth, which turns out to be not very volatile, and the predicted risk free rate is too high and not very steady. Therefore, the model should be analyzed more deeply and go beyond the rejection test and analyze its SDF and the pricing errors that it yields; i.e., observe whether the predicted returns indeed line up with actual average returns.

For this reason, we now turn to analyze these other dimensions of the model. Table 4 shows the average actual returns and the predicted returns, which are given by

$$E(R_t^n) = \frac{1}{E(M_t)} - \frac{\text{cov}(M_t, R_t^n)}{E(M_t)}. \quad (25)$$

We use Figure 2 to evaluate the estimation at a first glance. This kind of figure is common and emphasizes the fact in which we are interested, showing how close are we to the actual returns with our predictions. The root mean square error will give us the precise answer to the question of how far we are from the actual returns.

As can be observed, the model predicts a market return slightly smaller than the observed one, although this difference is not very large, so it does not do a bad job of predicting the equity premium. On the other hand, the predicted returns of value and size portfolios are more spread out. Thus, the model does not account for this value premia as well as it does for the equity premium. Although the model is rejected, we can observe small pricing errors, especially for the market return (around 0.17%), which is a desirable feature of an asset pricing model. For each of the estimation we also computed the root mean square errors (RMSE) to evaluate the fit of the model. For the case of complete markets the RMSE is 0.0049.

As we mention above, the standard model poses all the volatility of the stochastic discount factor in consumption growth and does not suffice to account for the equity premium. In this model the volatility of the stochastic discount factor is also given by the growth rate of the surplus consumption ratio (equation (13)). The estimations show how this variable adds an important effect to the volatility of M_{t+1} . Figure 3 shows the two components of the SDF: consumption growth and the surplus consumption ratio growth.¹⁰ It can be seen that this variable adds more volatility to the SDF than the one induced solely by consumption growth.

⁹See, for example, Hansen and Jagannathan (1991)

¹⁰The SDF is a geometric mean of both of them.

Table 4: Actual Expected Returns and Predicted Returns with Complete Markets

Portfolios	Actual	Predicted
MKT	1.0057 (0.0442)	1.0047
BL	1.0053 (0.0476)	1.0064
BM	1.0062 (0.0434)	1.0061
BH	1.0075 (0.0459)	1.0057
SL	1.0061 (0.0693)	1.0090
SM	1.0090 (0.0534)	1.0083
SH	1.0109 (0.0566)	1.0080
B1	1.0014 (0.0029)	1.0025
B5	1.0019 (0.0031)	1.0025

*Standard Errors in parenthesis




Figure 2: Actual vs. Predicted Returns




Figure 3: Components of the Stochastic Discount Factor

The surplus consumption ratio also determines the (time varying) coefficient of relative risk aversion, as equation (7) shows. People become more risk averse when S_t is smaller; i.e., when consumption gets closer to the subsistence level. Figure 4 shows the evolution of the CRRA coefficient. These peaks and troughs show the periods when consumers become more risk averse or less risk averse. Remember that the standard model cannot account for the large premium observed in data unless high values for risk aversion are allowed. In this case, the model cannot produce small values for the CRRA coefficient. For the case of complete markets and thus aggregate-consumption data, the CRRA coefficient is very high, taking values above 100.

There are some other features worth mentioning. The model has some characteristics that are replicated by the estimation. First, the process for the surplus consumption ratio seems to be stationary. Second, the model incorporates a habit or subsistence level that is time-varying and slowly adjusts to consumption. This result is also obtained and the behavior of the habit is induced from the generation of the S_t , as shown in Figure 5.

This figure shows the evolution of the habit X_t . Including habits in the model captures the idea that people get used to an accustomed standard of living, so a decrease in consumption may have an effect after some years of good times (see Cochrane (2005)). In this model, times of recession are identified with periods in which consumption gets closer to the habit, and risk aversion increases. Figure 5 shows these facts. Even if the habits adjust slowly to consumption, notice that there are some periods in which the gap between the two variables widens.




Figure 4: Relative Risk Aversion

These periods coincide with those when the risk-aversion coefficient reaches a peak.

In a nutshell: the estimation and tests of the habits model under complete markets show that we cannot accept the null hypothesis that the moments are zero; i.e., there is no statistical evidence supporting the idea that the SDF does price assets. The estimations produce high values for relative risk aversion, persistent and stationary surplus consumption ratio, and it does a good job in predicting the market return. It is worth noting that this does not mean that we can categorically reject the habits model in asset pricing. This is the case for this particular group of assets, but the rejection of the model as a whole requires a much deeper work of estimation for all kinds of returns (assets, bonds, exchange rates, etc.).




Figure 5: Consumption and Habit Evolution

4.2 Limited Participation

In this section we analyze the same model, but here we take into consideration the fact that not all the consumers hold stocks, and therefore not all the consumers should count towards the pricing implications of the portfolios used in the estimation. We move from NIPA accounts to CEX data on consumption. We filter the consumers according to asset position, which is one of the answers in the questionnaire. This process is explained in Section 3. Whereas in the previous section we used aggregate consumption as a variable indicating the consumption of the representative agent, this agent now represents only the stockholders, also subclassified in groups, depending the amount of assets they declared.

Table 5 reports some descriptive statistics for asset holders and non-asset holders using CEX data. We can observe the consumption information for people who do not hold assets at all, as well as people who hold some wealth in stocks; among these people, some sub-groups were formed to analyze consumption for wealthier households. The table shows a significant difference in the average consumption level for non-stock holders and stock holders; for the latter we also observe more variance for all the groups relative to the non-stock holders. Notice also that consumption growth also differs between stock holders and people who do not hold any assets. Keep in mind that consumption

growth is the process that prices assets in the benchmark model, so if we consider that the returns are determined by those individuals who participate in the market, then these differences in consumption growth play an important role in predicting the equity premium.

The third column of Table 5 computes the correlation between consumption growth and the market return; it is this correlation that matters for the equity premium. In the benchmark model, consumption growth is the force that drives asset pricing, and the correlation between this and the market return is what determines the equity premium. We do need a positive correlation between these variables to get a positive excess return, as we can observe this is the case for the stock holders. We observe that as the definition of stockholder is tightened, i.e. the larger is amount of assets, this correlation decreases.

Table 5: Average Consumption of Asset Holders

Asset Holders	C_t	C_{t+1}/C_t	$\text{corr}(\Delta c_t, R_t^{mkt})$
No	328.1778 (15.7388)	1.0003 (0.0187)	0.0172
> 0	471.8377 (27.4271)	1.0031 (0.0439)	0.1134
> 10,000	488.2128 (30.1617)	1.0029 (0.0453)	0.0734
> 20,000	512.9074 (40.9138)	1.0037 (0.0579)	0.0472
> 30,000	531.2672 (48.7482)	1.0063 (0.0733)	0.0320
> 40,000	548.1291 (53.7001)	1.0063 (0.0756)	0.0229
> 50,000	561.5608 (60.6146)	1.0076 (0.0869)	0.0169
All	349.5341 (15.1012)	1.0006 (0.0181)	0.0471

* Standard deviations are shown in parentheses.

Figure 6 confirms the previous information. We can observe the evolution of consumption of asset holders and non-asset holders. The former group shows a higher average consumption than the latter. It is worth noting, as was mentioned above, that there is a large proportion of the population that does not participate in the market. Hence if the participants are the ones that count towards pricing assets, their consumption is the one that has to account for the equity premium.

As Vissing-Jørgensen (2002b) points out, one of the major facts in portfolio choice is the substantial turnover in the set of stock market participants. This can be observed in Figure 7: the plot shows the number of households participating in the asset market in each period. The plot also shows two important decreases in the number of households, which correspond to changes in the way the survey was done and the details will be covered in Section 3.




Figure 6: Consumption of Asset Holders




Figure 7: Evolution of Asset Holders

The facts presented in the table and figures above also motivate our investigation to consider limited participation models. Although there is literature focusing in limited participation issues, there are not many efforts to find a synergy between this model and others including habitual consumption.

The main results of this section are summarized in Table 6 and Figure 9. We present the estimated parameters of the model and their corresponding standard errors for all the subsamples of consumers. We aggregated consumption among stockholders and used that consumption for the Euler equation estimation, as in (22). As explained in Section 3, we have divided the stockholders in different groups, depending on the amount of assets they are holding. For each group of consumers, we have estimated the model, and the results are presented in different rows for each group. Throughout this section, we assume complete markets among stockholders, because the marginal rate of substitution that we are using is that of the representative consumer. Standard errors are also presented in the tables. They were computed numerically using two-sided finite differences.

Table 6: GMM Estimation under Limited Participation

Threshold	γ	ϕ	s_0	δ	J_T	RMSE
> 0	0.9018 (0.5185)	0.9702 (0.0726)	-1.2342 (7.3467)	0.9897 (18.8895)	105.6	0.0117
$\geq 10,000$	1.6925 (0.2001)	0.9675 (0.0965)	-1.0836 (5.3748)	0.9780 (11.1656)	106.1	0.0133
$\geq 20,000$	1.6372 (0.3297)	0.9679 (0.1613)	0.5164 (5.3376)	0.9790 (28.6672)	106.0	0.0211
$\geq 30,000$	0.9817 (0.2307)	0.9810 (0.1433)	0.2348 (5.3330)	0.9909 (32.9697)	107.4	0.0169
$\geq 40,000$	0.9900 (0.2901)	0.9833 (0.1487)	0.1812 (7.8322)	0.9924 (36.8399)	107.4	0.0189
$\geq 50,000$	0.9923 (0.2492)	0.9875 (0.1747)	0.2309 (5.7093)	0.9937 (38.0734)	107.2	0.0177

* Standard errors in parentheses.

The most important parameter to be estimated is γ . It represents the curvature of the utility function. We observe in its estimations a sharp change among those who are considerably wealthy. Note that, as mentioned above, this is not the relative risk-aversion coefficient, although is directly related. As explained in Section 2, the risk-aversion coefficient is given by (7) and depends on how far the household is from the habits level, so it should be time varying. Table 8 shows the average overtime of the risk-aversion coefficient. There we observe the same pattern as for the curvature parameter. It is a fact that when we narrow the definition of stockholder, the risk aversion is considerably lower. There is a high measurement error given by the dramatic decrease in number of observations for the groups with more assets. This causes non-monotonous changes in the relative risk aversion. Nevertheless it can be inferred that wealthier agents

are less risk averse than those with a lower level of assets.

We are also showing some moments of the returns next to the returns predicted by the model. Recall from the data description that the returns are monthly. Below, Figure 9 depicts the expected actual returns versus the predicted returns. Table 7 shows the numbers for that figure. Both of them are presented for the estimation that corresponds to the asset holder definition of “having more than 1\$ in assets”. The root mean square error gives us the global sense of the goodness of the estimation, more than the difference between each of the returns, although we can already observe that we were not able to generate much variability across portfolios. Graphically is more clear in Figure9.

Table 7: Actual Expected Returns and Predicted Returns with Limited Participation

Portfolio	Actual	Predicted
MKT	1.0092 (0.0445)	1.0016
BL	1.0097 (0.0493)	1.0068
BM	1.0092 (0.0441)	1.0069
BH	1.0099 (0.0421)	1.0069
SL	1.0067 (0.0711)	1.0068
SM	1.0099 (0.0486)	1.0070
SH	1.0105 (0.0502)	1.0070
B1	1.0026 (0.0026)	1.0071
B5	1.0033 (0.0027)	1.0071

*Standard Errors in parenthesis

Table 8 shows the average risk aversion with its corresponding standard deviation and inverse of the expected value of the SDF. As a valid discount factor, it follows that the risk free return must be $R_{t+1}^f = 1/E(M_{t+1})$. It is reasonable that as we filter the consumers more tightly and we require them to be financially wealthier, they will be further from the habit’s lower bound of consumption (higher surplus consumption ratio) and the better they will be able to insure themselves against risk. Therefore, the surplus consumption ratio will be less volatile, making the RRA coefficient less volatile. In bad times, consumers turn out to be more risk averse, so instead of borrowing to keep a normal level of consumption, they save to make sure they are not going to be even worse off the following period. This feature is given by the non-linearity

specification of the habits process. In Figure 8 we observe the evolution of the coefficient of risk aversion. We observe how the peaks of the process correspond to dips in consumption, therefore to periods where the consumption growth is considerably smaller than one.

Table 8: Average Risk Aversion, Predicted Risk Free Rate under Limited Participation

Threshold	$E\left(\frac{\gamma}{S_t}\right)$	$R^f = \frac{1}{E(M_{t+1})}$
Assets > 0	3.1119 (0.4886)	1.0071
Assets > 10,000	4.5659 (0.5799)	1.0102
Assets > 20,000	3.2779 (0.7942)	1.0016
Assets > 30,000	1.5326 (0.3020)	1.0036
Assets > 40,000	1.4099 (0.2333)	1.0027
Assets > 50,000	2.6896 (0.1446)	1.0227

* Standard Errors in parentheses

The estimates for the subjective discount factor δ are relatively low, even though we are using monthly frequencies. But the standard errors are big, resulting in less reliable estimates.

The parameter ϕ represents the coefficient at which the habit adjusts geometrically to the consumption. It is the autoregressive coefficient of the process for the log surplus consumption ratio in (9). Estimates are close to 1, but that does not blow up the process, in the sense that it remains stationary, away from a unit root process. That ensures that the lower bound is never attained.

The estimates for s_0 do not have a direct explanation; it is just a parameter that is needed to initialize the process for the surplus consumption ratio. However, we do not have any intuition to explain whether it is or it is not a reasonable result. What we can say is that the entire process behaves nicely; it never attains the upper bound represented by s^{\max} and it is stationary regardless of the fact that ϕ is close to 1.

The most illustrative result is given by Figure 9, however, where we show again how the predicted returns differ from the actual returns. The predicted returns have been computed from the expression given by (25). Comparing it to Figure 2, we observe that there is no improvement by adding the assumption of limited participation to the model. The predicted return of the market given by the benchmark model of complete markets is more accurate than that given by the model including only those who have stocks.

Taking a closer look at that figure, we observe than the market portfolio




Figure 8: Relative Risk Aversion

is the one that causes more difficulties in being predicted. The benchmark portfolios are closer to the 45 degree line. In this case we also computed the RMSE which amount 0.0117. As we can observe the fit of this model is worse than the case of complete markets. Later we will compare all the models and find that limited participation has the worst fit among the models tested.

Figure 10 shows one of the most important features of the model. What previous attempts to nail down the Equity Premium Puzzle lacked was the ability of generating a sufficiently volatile SDF. In this figure, we see how the model succeeds in creating this volatility. While the usual SDF is composed of consumption growth only, here it is also composed of the growth in the surplus consumption ratio, which is shown in the lower plot. It can be seen how this component adds volatility to the whole SDF, an improvement of the model as explained above.

The central task of financial economics is to figure out what are the real risks that drive asset prices and expected returns. Therefore, this non-linear specification of habits injects a high volatility in the SDF, and also permits precautionary savings to enter into the task of pricing the assets accordingly.

In Figure 11 we observe the evolution of the habit and how it adjusts to consumption. When there is a big decline in consumption, habit does not react immediately, and that causes panic in the consumer. He or she does not like being close to the habit, so he/she saves and therefore fears the stocks because of the bad behavior of returns in bad times.




Figure 9: Actual vs. Predicted Returns




Figure 10: Components of the Stochastic Discount Factor




Figure 11: Consumption and Habit Evolution

4.3 Incomplete Markets

In this section we repeat the exercise relaxing the assumption of complete markets. We estimate the parameters using all the subsamples of consumers. That will complete the possibilities, in the sense that we are testing the four possible combinations. Now we test incomplete markets with and without limited participation. In the previous section we tested complete markets with limited participation and in the first case we tested complete markets without limited participation. Thus, we maintain throughout this section the assumption of neither stockholders or the rest of the individuals can hedge themselves in all the possible states of the world. For that purpose, we do not compute anymore the aggregate consumption of any group for the stochastic discount factor. Alternatively, we compute the marginal rate of substitution for every single consumer and then average across individuals. We do these because the Euler Equation should hold for every individual, then if the individual SDF is a valid one an average of all the SDFs is also a valid SDF.

Under the null hypothesis of the existence of incomplete markets, that object should be a proper stochastic discount factor. The results of the second-stage estimations are summarized in Table 9. We can observe the estimates of the curvature parameter in column one. The model is also rejected following the J_T criterion for the GMM estimation; we do not report p -values because they are all zero. But as Fama and French (1996) emphasized, we focus on the spread

in average returns and the size of the pricing errors. That result is presented in the last column of Table 9. According to this criterion, the best fitting case is the estimation corresponding to the group of asset holders Table 11 displays more interesting information. The risk aversion coefficient is roughly declining as we tighten the definition of stockholders. It is an expected and reasonable result, since we claim that wealthier agents are farther away from the habit level, therefore they are expected to be less risk averse, given the functional form of the utility.

Table 9: GMM estimation under Limited Participation and Incomplete Markets

Threshold	γ	ϕ	s_0	δ	χ^2	RMSE
All	1.0012 (0.2169)	0.9904 (0.3948)	-1.3836 (2.0224)	0.8932 (28.2914)	108.9	0.0116
> 0	1.8200 (0.5115)	0.9304 (0.2878)	-0.3400 (0.2997)	0.6118 (12.9693)	100.1	0.0083
$\geq 10,000$	2.0754 (0.2080)	0.9239 (0.1082)	-0.9288 (0.1445)	0.5003 (12.8200)	101.6	0.0173
$\geq 20,000$	1.6930 (0.1688)	0.9902 (0.1317)	-0.3125 (0.2662)	0.6335 (12.4518)	107.5	0.0144
$\geq 30,000$	1.8912 (0.1371)	0.9649 (0.0926)	-0.2753 (0.1112)	0.5393 (9.4449)	97.5	0.0151
$\geq 40,000$	1.8239 (0.2052)	0.9524 (0.1174)	-0.5155 (0.1339)	0.5724 (10.4896)	95.9	0.0281
$\geq 50,000$	2.0357 (0.1810)	0.9529 (0.1419)	-0.3964 (0.2322)	0.5001 (10.1653)	100.2	0.0667

* Standard errors in parentheses.

We also compare the average actual returns with our predicted returns under this model specification. To match the first moment of the actual returns is basically the goal of the paper. We can see that the return on the market portfolio is matched accurately and it is evident that the most problematic portfolio is the small and low book-to-market. That can be seen in Table 10. By itself, the prediction for the market portfolio expected returns could be a desirable result. But once we take into account the rest of the moments to match, we distance from the results of the benchmark setup. This is true once we observe the pricing errors, about which we talk a few lines below. Again, a clear picture of the estimation is Figure 12. They correspond to the estimation including the group of consumers who own at least one dollar in the stock market, to make it comparable to the estimation presented in the previous section. In this case, we do generate variability across portfolios, obtaining better results than in the previous section. But not better than in the section where we assumed complete markets. In this particular case, the root mean square error of the estimation is 0.0083. They are lower than the ones found in the previous section for the first three estimations (i.e. asset holders defined as having at least \$1, \$20,000, or \$30,000 in assets) where we were taking into account limited par-

ticipation with complete markets among stockholders. However they are still lower than the complete markets case. To explore also the fourth case, which would be incomplete markets without limited participation, it is reported a root mean square error of 0.0116. It can be inferred from this RMSE that taking into account limited participation issues is not as important as we might think because it does not improve the quality of the predictions. As Cochrane (2005) cited, *it is not clear that the stockholder/nonstockholder distinction is vital*. He claims that people who is not investing in the stock market, choose not to do so optimally, in the face of trivial fixed costs. This assertion may be arguable but the data show that there is no improvement in the quality of pricing by excluding from the sample those who do not own stocks. In any case, non-stockholders are somehow “in” the market because they usually hold pension or benefit plans linked to stock market investments. So basically we support empirically Cochrane’s intuition by finding that using aggregates is not only not terribly wrong, but using micro-data to rule out non-stock holders does not help to predict better expected returns.

Table 10: Actual Expected Returns and Predicted Returns with Limited Participation and Incomplete Markets

Portfolios	Actual	Predicted
MKT	1.0092 (0.0445)	1.0092
BL	1.0097 (0.0493)	1.0097
BM	1.0092 (0.0441)	1.0065
BH	1.0099 (0.0421)	1.0082
SL	1.0067 (0.0711)	1.0132
SM	1.0099 (0.0486)	1.0072
SH	1.0105 (0.0502)	1.0079
B1	1.0026 (0.0026)	1.0018
B5	1.0033 (0.0027)	1.0018

*Standard Errors in parenthesis

We also report figures 13, 14, and 15 as we did in the previous sections. It is observed a much more volatile risk aversion coefficient since the consumption is also more volatile than in the benchmark case. The much higher volatility of the micro data provided in the CEX is accountable for this fact.

Table 11: Average Risk Aversion, Predicted Risk Free Rate under Limited Participation and Incomplete Markets

Threshold	$E\left(\frac{\gamma}{S_t}\right)$	$R^f = \frac{1}{E(M_{t+1})}$
All	4.0249 (0.4348)	1.0090
Assets > 0	6.7252 (1.2093)	1.0017
Assets > 10,000	7.7275 (1.1648)	1.0013
Assets > 20,000	2.0731 0.0577	1.0107
Assets > 30,000	2.9473 (0.2175)	1.0069
Assets > 40,000	3.2797 (0.2578)	1.0121
Assets > 50,000	3.1793 (0.1996)	1.0266

*Standard Errors in parentheses




Figure 12: Actual vs. Predicted Returns




Figure 13: Relative Risk Aversion




Figure 15: Consumption and Habit Evolution




Figure 14: Components of the Stochastic Discount Factor

5 Conclusions

The purpose of this paper is to estimate and test a consumption-based model in Asset Pricing in which the utility function includes an external habit. The tests were performed in three different stages: complete markets, limited stock market participation and incomplete markets within limited participation. Several reasons justify our work: there is not enough empirical work in habits models that perform estimation as opposed to calibration. There is also strong evidence that a large proportion of households do not participate in the stock market. And we can claim that there is evidence of the presence of market incompleteness, meaning that the individuals are not able to insure their consumption path against any state of the world in the future.

When it comes to statistically testing the model under any of the three setups, we found that the models are rejected, according to the J -test. This implies that for the group of assets considered, we cannot accept the null hypothesis that the Euler equation holds. But we do not stop there. Over the years, empirical work in financial economics has developed models, tested them and rejected them. Rejection, though, is not the only thing we should take into account when analyzing the performance of a model. We cannot categorically disregard the consumption-based model solely relying on a statistical rejection for a group of assets. As Cochrane (2005) points out, two of the most important papers in asset pricing, namely Hansen and Singleton (1982) and Fama and French (1996), are formally rejected, but the latter shows in what dimensions

the model works. Besides p-values and J -tests, it is also important to look at pricing errors and predicted returns.

We observe that a model with complete market best predicts average returns since shows the smallest RMSE. When we include either limited participation or incomplete markets, the fit of the models are very similar, the RMSE are considerably higher than the complete-markets case. Since the inclusion of only one of these features did not improve the fit of the model we estimated a SDF with limited participation and incomplete consumption insurance among individuals. This model is able to drastically reduce the pricing errors but still the RMSE is bigger than the complete markets case. However it is a better predictor of the equity premium. One reason for the success of the model is the inclusion of the habits in the utility function, resulting in countercyclical effect of the savings decision of the investors due to precautionary savings and a much higher volatility of the Stochastic Discount Factor needed to drive down the equity premium.

References

- Abel, Andrew B. 1990. "Asset Prices under Habit Formation and Catching Up with the Joneses." *A.E.R. Papers and Proceedings* 80 (May): 38–42.
- Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura. 2004. "What really happened to consumption inequality in the US?" *Working Paper*.
- Attanasio, Orazio, and G. Weber. 1995. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy* 103 (95): 1121–1157.
- Attanasio, Orazio P, James Banks, and Sarah Tanner. 2002. "Asset holding and consumption volatility." *The Journal of Political Economy* 110 (4): 771–92.
- Basak, Suleyman, and Domenico Cuoco. 1998. "An Equilibrium Model with Restricted Stock Market Participation." *The Review of Financial Studies* 11 (2): 309–341 (Summer).
- Battistin, E. 2003. "Errors in survey reports of consumption expenditures." *Working Paper W03/07, Institute for Fiscal Studies, London*.
- Becker, Gary S. 1992. "Habits, Addictions, and Traditions." *Kyklos* 45 (3): 327–45.
- Brav, Alon, George M. Constantinides, and Christopher Geczy. 2002. "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." *Journal of Political Economy* 110 (4): 793–824.
- Bureau of Labor Statistics. 2003. Chapter 16, Consumer Expenditures and Income of *Handbook of Methods*, edited by Bureau of Labor Statistics.
- Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–251.

- . 2000. “Explaining the poor performance of consumption-based asset pricing models.” *The Journal of Finance* LV (6): 2863–2878 (December).
- Chen, Xiaohong, and Sydney C. Ludvigson. 2004, May. “Land of Addicts? An Empirical Investigation of Habit-Based Asset Pricing Models.” Unpublished manuscript, Columbia University.
- Chen, Yu, Thomas F. Cosimano, and Alex A. Himonas. 2003, November. “By Force of Habit: An Exploration of Asset Pricing Models using Analytic Methods.” unpublished manuscript, University of Notre Dame.
- . 2004, November. “Internal versus External Habits: Why does the Distribution of Stock Returns have ARCH Properties?” unpublished manuscript, University of Notre Dame.
- Chue, Timothy K. 2005. “Conditional market comovements, welfare, and contagions: The role of time-varying risk aversion.” *Journal of Business* 78 (3): 949–967.
- Cochrane, John H. 2001. *Asset Pricing*. Princeton: Princeton University Press.
- . 2005, February. “Financial Markets and the Real Economy.” forthcoming in *The International Library of Critical Writings in Financial Economics*.
- Constantinides, George M. 1990. “Habits Formation: A Resolution to the Equity Premium Puzzle.” *The Journal of Political Economy* 98 (3): 519–543 (June).
- Constantinides, George M., and Darrell Duffie. 1996. “Asset Pricing with Heterogeneous Consumers.” *Journal of Political Economy* 104 (2): 219–240.
- Deaton, Angus. 1994. *Understanding Consumption*. New York: Oxford University Press.
- Duffee, Gregory R. 2005. “Time variation in the covariance between stock returns and consumption growth.” *Journal of Finance* 60 (4): 1673–1712 (August).
- Duffie, Darrell, and Kenneth J. Singleton. 1993. “Simulated Moments Estimation of Markov Models of Asset Pricing.” *Econometrica* 61 (4): 929–952 (July).
- Epstein, Larry G., and Stanley E. Zin. 1991. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis.” *The Journal of Political Economy* 99 (2): 263–286 (April).
- Fama, Eugene F., and Kenneth R. French. 1992. “The Cross-Section of Expected Stock Returns.” *The Journal of Finance* 47 (2): 427–465 (June).
- . 1993. “Common Risk Factors in the Returns of Stocks and Bonds.” *Journal of Financial Economics* 33:3–56.

- . 1995. “Size and Book-to-Market Factors in Earnings and Returns.” *The Journal of Finance* 50 (1): 131–155 (March).
- . 1996. “Multifactor Explanations of Asset Pricing Anomalies.” *The Journal of Finance* 51 (1): 55–84 (March).
- Fillat, José L., and Hugo Garduño. 2004, june. “Beyond Habits: An Empirical Investigation.” The University of Chicago, unpublished manuscript.
- Hansen, Lars P. 1982. “Large Sample Properties of Generalized Method of Moments Estimators.” *Econometrica* 50:1029–1054.
- Hansen, Lars P., and Ravi Jagannathan. 1991. “Implications of Security Market Data for Models of Dynamics Economies.” *The Journal of Political Economy* 99 (2): 225–262 (April).
- . 1997. “Assessing Specification Errors in Stochastic Discount Factor Models.” *The Journal of Finance* 52 (2): 557–590 (June).
- Hansen, Lars P., and Kenneth J. Singleton. 1982. “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models.” *Econometrica* 50 (5): 1269–1286.
- Heaton, John. 1995. “An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specification.” *Econometrica* 63 (3): 681–717 (May).
- Huang, Chi-fu, and Robert H. Litzenberger. 1988. *Foundations of Financial Economics*. North-Holland.
- Jagannathan, Ravi, and Zhenyu Wang. 2002. “Empirical Evaluation of Asset-Pricing Models: A comparison of the SDF and Beta Methods.” *The Journal of Finance* LVII (5): 2337–2367 (October).
- Lee, Bong-Soo, and Beth Fisher Ingram. 1991. “Simulation Estimation of Time Series Models.” *Journal of Econometrics*, vol. 47.
- Li, Yuming, and Maosen Zhong. 2005. “Consumption habit and international stock returns.” *Journal of Banking and Finance* 29 (3): 579–601 (March).
- Mankiw, N.G., and S.P. Zeldes. 1991. “The Consumption of Stockholders and Non-Stockholders.” *Journal of Financial Economics* 29:97–112.
- Matsen, Egil. 2003. “Habit persistence and welfare gains from international asset trade.” *Journal of International Money and Finance* 22 (2): 239–260 (April).
- McCarthy, M., D. Johnson, T. Garner, and B. Passero. 2002. “Issues in Construction and Research Use of the Consumer Expenditure Survey.” *paper presented at the 2002 NBER Summer Institute*.
- Mehra, Rajnish, and Edward C. Prescott. 1985. “The Equity Premium: A Puzzle.” *Journal of Monetary Economics* 15 (March): 145–161.
- . 2003. Chapter 14 of *Handbook of the Economics of Finance*, edited by G.M. Constantinides, M. Harris, and R. Stulz, 887–936. Elsevier B.V.

- Menzly, Lior, Tano Santos, and Pietro Veronesi. 2004. “Understanding Predictability.” *Journal of Political Economy* 112 (1): 1–47.
- Otrok, Christopher, B. Ravikumar, and Charles H. Whiteman. 2002. “Habit formation: a resolution of the equity premium puzzle?” *Journal of Monetary Economics* 49 (6): 1261–1288 (September).
- Pastor, Ľuboš, and Pietro Veronesi. 2005. “Rational IPO waves.” *Journal of Finance* 60 (4): 1713–1757 (August).
- Santos, Tano, and Pietro Veronesi. 2005, October. “Cash-flow Risk, Discount Risk and the Value Premium.”
- Vissing-Jørgensen, Annete. 2002a. “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution.” *Journal of Political Economy* 110 (August): 825–853.
- . 2002b. “Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures.” *The University of Chicago, Working Paper*, March.
- . 2003. “Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion.” *AEA Papers and Proceedings* 93 (2): 383–391 (May).