SCHOOL OF ECONOMICS AND MANAGEMENT
FACULTY OF SOCIAL SCIENCES
AARHUS UNIVERSITY

CREATES
Center for Research in Econometric
Analysis of Time Series

CREATES Research Paper 2010-46

Habit-based Asset Pricing with Limited Participation Consumption

Christian Bach and Stig Vinther

School of Economics and Management
Aarhus University
Bartholins Allé 10, Building 1322, DK-8000 Aarhus C
Denmark

Habit-based Asset Pricing with Limited Participation Consumption*

Christian Bach[†] Stig Vinther Møller[‡]

June 16, 2010

Abstract

We calibrate and estimate a consumption-based asset pricing model with habit formation using limited participation consumption data. Based on survey data of a representative sample of American households, we distinguish between assetholder and non-assetholder consumption, as well as the standard aggregate consumption series commonly used in the CCAPM literature. We show that assetholder consumption outperforms non-assetholder and aggregate consumption data in explaining bond returns, bond yields, and the volatility of bond yields. We further show that the high volatility of assetholder consumption enables the model to explain the equity premium puzzle and the risk-free rate puzzle simultaneously for a reasonable value of relative risk aversion.

JEL Classifications: C32, G12.

Keywords: CCAPM, Limited participation consumption data, Habit formation, Real term structure, Risk premium, GMM estimation.

*We would like to thank Peter O. Christensen, Tom Engsted, Claus Munk, Thomas Q. Pedersen, and seminar participants at Aarhus University for useful comments. We are grateful to CREATES (funded by the Danish National Research Foundation), D-CAF, and the Danish Social Science Research Council for research support.

[†]School of Economics and Management, Aarhus University, Bartholins Allé 10, 8000 Aarhus C, Denmark; email: cbach@creates.au.dk.

[‡]Finance Research Group, Aarhus School of Business, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark; email: svm@asb.dk.

1 Introduction

The habit-based asset pricing model developed by Campbell and Cochrane (1999) has become one of the most prominent models within the consumption-based framework. The model explains countercyclical variation in equity premia and other stylized facts of the aggregate stock market. Campbell and Cochrane balance intertemporal substitution with precautionary savings to obtain a constant risk-free rate, while Wachter (2006) extends the model such that the risk-free rate is time-varying. This also leads to countercyclical variation in bond risk premia, and Wachter shows that the model has the ability to produce a reasonable fit of the means and volatilities of bond yields.¹

In this paper we extend the work of previous studies by taking into account that not all households trade and own assets. Our objective is to analyze the ability of the habit model to price assets using the consumption of households who actually do invest in assets. In order to perform such an analysis, we use interview data from the U.S. Consumer Expenditure Survey (CEX) and distinguish between the consumption of households that do invest in assets and the consumption of those that do not. For comparison reasons, we also consider aggregate consumption, which is typically used when working with consumption-based models.²

While previous literature primarily uses household-level data to explain stock market behavior (see, for instance, Mankiw and Zeldes (1991), Parker (2001), Brav et al. (2002), and Malloy et al. (2009)), we also focus the attention on the bond market. As we include the bond market in our analysis, we follow Wachter (2006) and allow the risk-free rate to be time-varying. Furthermore, as in Campbell and Cochrane (1999) and Wachter (2006), we use calibration and simulation techniques to evaluate the performance of the model, but additionally we also consider formal econometric estimation of the model using Generalized Method of Moments (GMM). We find it useful to apply different methodologies as this allows us to test the model in different dimensions and check the robustness of our results.

Our main contribution is to show that the assetholder/non-assetholder distinction is im-

¹The NBER working paper version of Campbell and Cochrane (1999) also has an analysis with a time-varying risk-free rate.

²For instance, Campbell and Cochrane (1999) and Wachter (2006) both use aggregate consumption.

portant when pricing assets in the habit formation framework. With calibrated parameter values we simulate the model and demonstrate that it works substantially better with as-setholder consumption than with non-assetholder and aggregate consumption. The model fits the means and standard deviations of bond yields with much higher accuracy using as-setholder consumption in comparison to the two other consumption measures. Thus, when we use the consumption of households who face bond market risk, the ability of the model to price bonds improves substantially.

When it comes to equity data, all three consumption measures work well in explaining the mean and standard deviation of the excess return on the aggregate stock market. Interestingly, however, there are substantial differences in the implied steady state risk aversion necessary to explain the large equity premia. When using assetholder consumption, the model only needs a steady state risk aversion of around 8 to explain the large equity premia, while the steady state risk aversion is around 11 with non-assetholder consumption and as high as 88 with aggregate consumption. In the original study of Campbell and Cochrane (1999), the model needs a high steady state risk aversion of 35 to explain the large equity premia.³ By contrast, in the limited participation setting that we apply in this paper, the model explains the large equity premia using an economically plausible level of risk aversion, and at the same time the model accounts for many salient features on the bond market. The result of a more plausible level of relative risk aversion is economically intuitive as those households who invest in asset markets are exposed to more risk than those who do not.

As a supplement to the calibration and simulation analysis, we also use GMM to estimate the model in a cross-sectional setup based on maturity-sorted bond returns. The GMM estimation confirms that the model's ability to price bonds improves when using the consumption of households that do invest in bonds. With non-assetholder and aggregate consumption the model has difficulties in explaining the cross-sectional variation in returns on maturity-sorted bond returns, but with assetholder consumption the model works reasonably well.

The rest of the paper is organized as follows. Section 2 briefly outlines the model, Section

³In the Wachter (2006) specification, the model also implies a high risk aversion of 52 at steady state.

3 describes the data and gives summary statistics, Section 4 demonstrates how we calibrate and simulate the model and presents results, Section 5 demonstrates how we estimate the model using GMM and presents results, and, finally, Section 6 concludes.

2 The Model

The habit model implies that each household maximizes:

$$E \sum_{t=0}^{\infty} \delta^t \frac{(C_t - X_t)^{1-\gamma} - 1}{1-\gamma}, \quad (1)$$

where δ , γ , C_t , and X_t are, respectively, the subjective discount factor, the utility curvature parameter, consumption, and an external habit level. From the specification of the utility function, it follows that the relative risk aversion is given by γ/S_t where S_t is the surplus consumption ratio defined as:

$$S_t \equiv \frac{C_t - X_t}{C_t}.$$

The relative risk aversion moves countercyclically over time: when C_t is well above X_t in cyclical upswings, the relative risk aversion decreases and when C_t gets close to X_t in cyclical downturns, the relative risk aversion increases. Following Campbell and Cochrane (1999) and Wachter (2006), we specify the log of the surplus consumption ratio, $s_t = \log(S_t)$, as an autoregressive process (throughout, lowercase letters indicate logs):

$$s_{t+1} = (1 - \phi) \bar{s} + \phi s_t + \lambda(s_t) v_{t+1}, \quad (2)$$

where ϕ is the persistence parameter and $\lambda(s_t)$ is a sensitivity function that determines how innovations in consumption growth v_{t+1} influence s_{t+1} . The consumption growth process is given by:

$$c_{t+1} - c_t = g + v_{t+1}, \quad v_{t+1} \sim \text{niid}(0, \sigma^2), \quad (3)$$

where g and σ are the mean and volatility of the log consumption growth. The sensitivity function $\lambda(s_t)$ is specified as follows:

$$\lambda(s_t) = \frac{1}{\bar{S}} \sqrt{1 - 2(s_t - \bar{s})} - 1, \quad \text{if } s_t \leq s_{\max}, 0 \text{ otherwise,} \quad (4)$$

where

$$\bar{S} = \sigma \sqrt{\frac{\gamma}{1 - \phi - b/\gamma}} \quad (5)$$

is the steady state level of S_t and

$$s_{\max} = \bar{s} + \frac{1}{2}(1 - \bar{S}^2) \quad (6)$$

is the value of s_t at which the expression in Eq. (4) becomes zero. With this model specification, the stochastic discount factor equals:

$$M_{t+1} = \delta \left(\frac{S_{t+1}}{S_t} \frac{C_{t+1}}{C_t} \right)^{-\gamma} = \delta e^{-\gamma\{g + (\phi-1)(s_t - \bar{s}) + [1 + \lambda(s_t)]v_{t+1}\}}, \quad (7)$$

and the log real risk-free rate is a linear function of s_t :

$$\begin{aligned} r_{f,t+1} &= \log \left[\frac{1}{E_t[M_{t+1}]} \right] \\ &= -\log(\delta) + \gamma g - \frac{\gamma(1 - \phi) - b}{2} - b(s_t - \bar{s}). \end{aligned} \quad (8)$$

Campbell and Cochrane (1999) set $b = 0$, which implies a constant real risk-free rate and a flat yield curve, as they have their main focus on explaining stylized facts on the equity market. We follow Wachter (2006) and allow for a time-varying real risk-free rate (i.e. $b \neq 0$). Campbell and Cochrane (1999) and Wachter (2006) both use aggregate consumption when they analyze the habit model's ability to explain asset prices. We extend their work by distinguishing between the consumption of households that actually own and trade assets and the consumption of those that do not. Thus, we take into account that a large group of households do not invest in assets. In the next section, we explain the data that we use in our application of the habit model.

3 Data and Summary Statistics

Calibrating and estimating the model, as well as comparing the implications of the model to the empirical data, require data on consumption, zero coupon bonds, inflation and prices and dividends on the market index. In this section we present the data as well as summary statistics of the consumption data.

The focus of this paper is to compare assetholder consumption data with non-assetholder and aggregate consumption data. We use quarterly aggregate real per capita consumption of non-durables and services from the NIPA tables available from the Bureau of Economic Analysis (BEA). Assetholder and non-assetholder consumption data is provided by Annette Vissing-Jørgensen⁴ and is explained in detail in Malloy et al. (2009).⁵ Therefore, we will only shortly explain the aspects of the data which are the most relevant for the application in this paper. The assetholder and non-assetholder data set is on a monthly frequency and contains series of consumption growth over the past 1, 2, 4, 8, 12, 16, 20, and 24 quarters. We use the series with quarterly growth and to avoid observations with overlapping information, we limit ourselves to a quarterly frequency of the data.⁶ Assetholder and non-assetholder data is available from the second quarter of 1982 to the fourth quarter of 2004, which limits our study to 91 observations on quarterly consumption.

Assetholder and non-assetholder consumption data is calculated based on household data from the CEX. The data is based on interviews with 4,500-7,500 households per quarter, and each household is asked to report its consumption for the past three months. Each household is interviewed five times and the first interview is discarded.⁷ The data is corrected for changes in family size and seasonalities. The assetholders are broadly defined as households answering positively to holding "stocks, bonds, mutual funds and other such securities".⁸ Non-assetholders are broadly defined as those not "holding stocks, bonds, mutual funds

⁴The data is available from <http://www.kellogg.northwestern.edu/faculty/vissing/htm/research1.htm>.

⁵Also see Battistin (2004) for further details on the survey.

⁶The use of overlapping observations would not be consistent with the random walk model of consumption in Eq. (3).

⁷A household is also discarded from the sample if the 2nd, 3rd, or 4th interview is missing.

⁸Malloy et al. (2009) attempt to verify this through a probit estimation (see their appendix B). If the result of the probit analysis is not in accordance with the answer given in the survey, then the observation is discarded from the sample.

and other such securities". The data set also contains a consumption growth series for top assetholders (wealthiest one third of assetholders). However, considering the time series properties of the quarterly consumption growth indicates that this series has significant measurement noise and we discard this series from our analysis.⁹

We do not have a time series of consumption growth for each household that spans the entire sample period. As in Malloy et al. (2009), we use a time series of average consumption growth calculated across the cross-section of households. For assetholders (AH) the average log consumption growth rate is given by:

$$\frac{1}{H_t^{AH}} \sum_{h=1}^{H_t^{AH}} \left(c_{t+1}^{h,AH} - c_t^{h,AH} \right),$$

where $c_t^{h,AH}$ is household h 's consumption at time t , and H_t^{AH} is the number of assetholders at time t . In a similar way, the average log consumption growth rate for non-assetholders (NAH) is given by:

$$\frac{1}{H_t^{NAH}} \sum_{h=1}^{H_t^{NAH}} \left(c_{t+1}^{h,NAH} - c_t^{h,NAH} \right).$$

Table 1 shows correlations as well as first and second moments of the consumption series. There is a positive correlation between all consumption series. It is also seen that assetholders have higher consumption growth than non-assetholders, but much lower than for the aggregate consumption series. One possible explanation is that the CEX and NIPA consumption series are based on different procedures and are therefore not directly comparable. The standard deviation of assetholder consumption data is higher than that of non-assetholder and aggregate consumption data, which is important because previous CCAPM studies point out that the volatility of consumption is too low to explain the volatile asset markets. We note that the additional volatility could potentially be due to measurement error in the CEX interview data. However, as we use the average consumption across the cross-section of households rather than individual consumption, we may be able minimize

⁹One could smooth the series using filter techniques, however, the presence of noise seems too prominent for this approach to yield trustworthy results.

the issue of measurement error.¹⁰ Of course, working with interview data makes measurement error inescapable, but as Parker (2001) points out the CEX interview data is "the best household-level data on consumption over time in the United States".

Monthly observations on returns and prices on the value-weighted NYSE/AMEX index are from CRSP. The monthly price of the index, as well as returns including and excluding dividends, is used to calculate monthly dividend payments. A time series at quarterly frequency of the price-dividend ratio is calculated by dividing the price of the index with the sum of dividends over the previous year.

We use the quarterly consumer price index (CPI) from CRSP as a measure of inflation, as this series is the most commonly used in the literature.¹¹ We use the nominal three month yield on a quarterly frequency from the Fama Risk Free Rates database in CRSP as the one period nominal risk-free rate. This rate is adjusted by inflation to obtain the real three month risk-free rate, corresponding to the risk-free rate calculated in Eq. (8). We use one, two, three, four, and five year zero coupon bond prices both to calibrate the model and to compare model implications to moments observed in the market. The bond data used has a quarterly frequency and is from the Fama-Bliss Discount Bond section in the CRSP database.¹² The inflation series is used to calculate real yields. The bond and inflation data is also used to calculate series of one year real bond holding returns with quarterly frequency. We follow Cochrane and Piazzesi (2005) and calculate one year holding returns as

$$R_{n,t} = \ln(P_{n,t}) - \ln(P_{n+4,t-4}), \quad (9)$$

where $R_{n,t}$ is the period t one year holding return from buying a bond with $n + 4$ periods to maturity one year (four periods) ago, and $P_{n,t}$ is the price at time t of a bond with unit pay-off in n periods. The inflation during the year is used to calculate real holding returns.¹³

Finally, in the GMM estimation we use quarterly real returns on government bond port-

¹⁰See Brav et al. (2002) for thorough analysis on measurement error in household consumption data.

¹¹Using the personal consumption expenditures (PCE) deflator from the NIPA tables to measure inflation yields similar results.

¹²For more information, see Fama and Bliss (1987).

¹³When we use expected inflation calculated based on, for instance, an AR(1) process, the results are nearly identical as the results that we obtain using ex post realized inflation. Results available on request.

folios with maturities of 1, 2, 5, 7, and 10 years, also available from CRSP.

4 Calibration and Simulation Methods

This section first describes the calibration methods and results. Then we show how the calibration relies on a simultaneous simulation of the model to calculate long-term bond and stock prices. The simulation method, in general, follows Wachter (2005) and the calibration method, in general, follows methods described in Campbell and Cochrane (1995), Campbell and Cochrane (1999), and Wachter (2006). In the previous literature the model has been calibrated to the aggregate market and bond yields, while we calibrate to the aggregate market and one year real bond holding returns. We do this to be consistent with the GMM estimation in Section 5.

4.1 Calibration

As described in Section 2, the model has six parameters to be calibrated: g , σ , γ , ϕ , δ , and b . The mean consumption growth and standard deviation of consumption growth were calculated in Section 3 to their corresponding moments in the data. The calibration results for the remaining parameters are shown in Table 2 for each of the consumption growth series. The habit persistence parameter, ϕ , can be calibrated directly from the data without having calibrated γ , δ , and b . We follow the previous literature and calibrate ϕ to match the first-order autocorrelation of the price-dividend ratio of the AMEX/NYSE index and we obtain $\phi = 0.9689$.¹⁴ The three remaining parameters must be calibrated jointly by e.g. minimizing an error function over choices of γ and b . Since the GMM estimation relies on moments based on real bond returns, we calibrate the model to match a cross section of real bond holding returns. We define a vector with one year holding returns of bonds calculated in the

¹⁴This value is similar to results found in previous calibrations. However, since the model is relatively sensitive to this parameter value, we have also calculated it for the AMEX/NYSE/NASDAQ index and find that $\phi = 0.9677$. The choice between the two values has only little impact on our results.

model¹⁵ with one to five years to maturity as:

$$\hat{\mathbf{R}}_{n,t} = \left(\hat{R}_{1,t}, \hat{R}_{2,t}, \hat{R}_{3,t}, \hat{R}_{4,t}, \hat{R}_{5,t} \right)'.$$

A similar collection of one year holding returns calculated from the data is denoted as $\mathbf{R}_{n,t}$. We then minimize the criterion function:

$$F(\gamma, b) = \mathbf{e}' \mathbf{I}_5 \mathbf{e},$$

where $\mathbf{e} = \frac{1}{T} \sum_{t=1}^T \left(\hat{\mathbf{R}}_{n,t} - \mathbf{R}_{n,t} \right)$ and $\mathbf{I}_5$ is a 5×5 identity matrix and thus we put equal weight on the returns irrespective of the bond from which it is calculated. The utility curvature parameter, γ , has large impact on the model Sharpe ratio and therefore it is appropriate to calibrate γ such that the simulated Sharpe ratio of the aggregate market in the model equals the Sharpe ratio in the data. The parameter b has only little impact on the Sharpe ratio and thus the Sharpe ratio is mainly determined from γ . However, since γ and b both impact simulated bond holding returns, we have to calibrate these parameters simultaneously. Thus, for each value of b we search for the value of γ that matches the empirical Sharpe ratio.¹⁶

The remaining parameter δ is calibrated such that at $s_t = \bar{s}$, the real risk-free rate in the model equals the real risk-free rate in the data, i.e. δ is calibrated from Eq. (8) with $s_t = \bar{s}$ as:

$$\delta = \exp \left\{ -\bar{r}_f + \gamma g - \frac{\gamma(1-\phi) - b}{2} \right\},$$

where $\bar{r}_f$ is the mean real risk-free rate in the data. As is seen from the formula, this parameter can be calculated when γ and b are known and can therefore be calculated inside the minimization routine for each guess of γ and b . From this calibration we ensure that the mean real risk-free rate in the model (approximately) matches the mean real risk-free rate in the data.

Wachter (2006) calibrates b to match a cross section of yields and this ensures a positive

¹⁵See Section 4.2 for the calculation of bond prices and Eq. (9) for the calculation of returns as a function of bond prices.

¹⁶Alternatively, one could add a penalizing term to the criterion function (with a large weight) to ensure matching the Sharpe ratio in the data. The two methods, naturally, yield "identical" results.

value of b . Our calibration method does not ensure this by construction. However, as is seen from Table 2, we still obtain positive values of b implying an upward-sloping yield curve. From Table 2 it is also seen that we obtain a higher value of the utility curvature parameter, γ , than in previous calibrations of this model. We get $\gamma = 3.70$ for assetholders, $\gamma = 4.68$ for non-assetholders, and we obtain an intermediate value for aggregate consumption. This indicates that assetholders are less risk averse than non-assetholders, which is likely to be the case since assetholders are defined as individuals investing in risky assets and since this group is likely to have higher income than non-assetholders.¹⁷ Both Campbell and Cochrane (1999) and Wachter (2006) find $\gamma = 2.00$, but this low value of gamma is not confirmed for any of our consumption series. One obvious explanation is the different data periods used.

As mentioned, we calibrate b to match a cross section of real bond holding returns and obtain a positive value for all consumption series. From Eq. (8) it is seen that for $b > 0$ the real risk-free rate is negatively correlated with surplus consumption, implying that real bonds have positive risk premia.

The subjective discount factor, δ , is higher for assetholders than for non-assetholders, which indicates that non-assetholders are less patient. This might be one of the reasons why they do not hold assets in the first place, since they are less likely to want to save up compared to assetholders.

The mean long-run surplus consumption ratio, $\bar{s}$, is mainly determined from the volatility of the consumption series which is seen directly from Eq. (5). Intuitively, this follows from the assumption of a non-negative surplus consumption ratio. When the volatility of consumption is high, the consumption must, in general, venture further above the habit index such that when there are large drops in consumption, it will remain above the habit index. For the aggregate market we obtain $\bar{s} = -3.05$, which is similar to the values of -2.85 found in Campbell and Cochrane (1999) and -3.25 found in Wachter (2006). For the more volatile consumption series of non-assetholders and in particular assetholders, we obtain much larger levels of $\bar{s} = -0.93$ and $\bar{s} = -0.79$, respectively. Naturally, similar results are obtained for

$s_{\max}$.

¹⁷The relative risk aversion in the model is given by γ/S_t , implying that the value of γ is a main determinant of the relative risk aversion.

4.2 Simulation

In this section we explain how the model is simulated both in order to calibrate the model and to assess the properties and implications of our results. Let $P_{n,t}$ denote the price of a real bond (a bond whose pay-off is one unit of the consumption good) with n periods to maturity at time t . The price of such a real bond is determined recursively by the Euler equation and can be written as:

$$P_{n,t} = E_t [M_{t+1} P_{n-1,t+1}]. \quad (10)$$

At maturity a real bond pays one unit of consumption implying that the boundary condition is given as $P_{0,t} = 1$. Knowing the price of the real bond, we can calculate the yield of the n period bond at time t as:

$$y_{n,t} = -\frac{1}{n} \ln(P_{n,t}). \quad (11)$$

There is no closed form expression for yields with more than two periods to maturity, and we therefore use numerical methods in order to calculate the bond prices. Notice that the surplus consumption ratio, s_t , is the only state variable in determining real bond prices in this model and the bond prices at time t can then be written as a function of s_t by substituting Eq. (7) into Eq. (10) to obtain:

$$P_n(s_t) = E_t [\exp \{ \ln(\delta) - \gamma g - \gamma(1 - \phi)(\bar{s} - s_t) - \gamma(\lambda(s_t) + 1) \sigma \varepsilon_{t+1} \} P_{n-1}(s_{t+1})], \quad (12)$$

where $\varepsilon_{t+1} \sim NID(0, 1)$ and with boundary condition $P_{0,t} = 1$. Based on Eq. (12) we simulate prices on real bonds and we construct real bond returns from Eq. (9).

In the calibration we also match the Sharpe ratio in the model with the Sharpe ratio in the data. Since the aggregate wealth in this economy is equal to the market portfolio with dividends equal to the aggregate consumption, the price-consumption ratio can be used to calculate the return and Sharpe ratio on the market. To do this, we simulate the

price-dividend ratio given by:

$$\begin{aligned}\frac{P_{n,t}^e}{C_t}(s_t) &= E_t \left[M_t \frac{C_{t+1}}{C_t} \frac{P_{n-1,t+1}^e}{C_{t+1}}(s_{t+1}) \right] \\ &= E_t \left[\exp \left\{ \frac{\ln(\delta) + (1-\gamma)g - \gamma(1-\phi)(\bar{s} - s_t)}{(1-\gamma(\lambda(s_t) + 1))\sigma\varepsilon_{t+1}} \right\} \frac{P_{n-1,t+1}^e}{C_{t+1}}(s_{t+1}) \right],\end{aligned}\quad (13)$$

where $P_{n,t}^e$ is the price at time t of an asset paying C_{t+n} in n periods. Notice that $P_{n,t}^e$ pays no dividends (coupons) and can thus be thought of as a zero coupon bond on future endowment. As with real bonds, the surplus consumption ratio s_t is the only state variable and the relevant boundary condition is $P_{0,t}^e/C_t = 1$. The price-consumption (price-dividend) ratio of the market is now given as the infinite sum of zero coupon bonds on future price-consumption ratios, i.e.

$$\frac{P_t}{C_t} = \sum_{n=1}^{\infty} \frac{P_{n,t}^e}{C_t}.$$

Practically, one has to choose a truncation point and we include 300 terms in our simulation.¹⁸

We can now use the simulated price-dividend ratios to calculate the log returns, needed to calculate the Sharpe ratio, as:

$$\begin{aligned}r_{t+1}^m &= \ln(R_{t+1}^m) \equiv \ln\left(\frac{P_{t+1} + C_{t+1}}{P_t}\right) \\ &= \ln\left(\frac{C_{t+1}}{C_t} \left(\frac{P_{t+1}}{C_{t+1}} + 1\right) \left(\frac{P_t}{C_t}\right)^{-1}\right) \\ &= g + \ln\left(\frac{P_{t+1}}{C_{t+1}} + 1\right) - \ln\left(\frac{P_t}{C_t}\right).\end{aligned}\quad (14)$$

With the above equations for yields (11), real zero coupon bond prices (12), and price-dividend ratios (13), we can calculate the unconditional Sharpe ratio in the model as:

$$SR = \frac{E[r^m - y_1]}{\sigma(r^m - y_1)}.\quad (15)$$

The simulation procedure can now be summarized as follows. First, 100,000 quarters of consumption data are simulated through Eq. (3) based on the parameter values of g and σ

¹⁸A simulation study shows that adding more terms has little effect on our results.

from Table 1. Secondly, the time series of the surplus consumption ratio s_t is generated by feeding the consumption draws through Eq. (2). Thirdly, real bond prices and the price-dividend ratio are calculated from Eqs. (12) and (13), respectively. In the final step, returns and the Sharpe ratio are calculated through Eqs. (9) and (15), respectively.

To calculate the price of real zero coupon bonds (12) and the price-dividend ratios (13) of the aggregate market, we apply the series method also applied in Wachter (2006).¹⁹ Wachter (2005) shows that this method has significant advantages over the more traditional fixed point method used in Campbell and Cochrane (1999) and for practical purposes it is computationally more efficient. Calculating the expectations in Eqs. (12) and (13) involves numerically calculating an integral on the form:

$$\int_{-\infty}^{\infty} p(v) F dv,$$

where the integrand, F , is either given from Eq. (12) for calculating real zero coupon prices or Eq. (13) for calculating the price-dividend ratio. We follow Wachter (2005) and calculate the integral using Gauss-Legendre 40-point quadrature and bound the integral by -8 and $+8$ standard deviations. The value of $P_{n-1}(s_{t+1})$ or $\frac{P_{n-1,t+1}^e}{C_{t+1}}(s_{t+1})$ is calculated at each step by interpolation on a grid, depending on the updated state variable, s_{t+1} . We use the finest grid on S_t from Wachter (2005) which consists of 101 evenly distributed points $]0, S_{\max}]$ and 900 logarithmically distributed points between e^{-300} and the lowest of the 101 points on $]0, S_{\max}]$. This makes the grid relatively finer when S is close to zero, since asset prices are relatively more sensitive to variations of S in this region.

In Figure 1 we have depicted the simulated 3 month and 5 year real yield curves as a function of the surplus consumption ratio, S_t . It is seen that the yield curves for the non-assetholders are more sensitive to changes in the surplus consumption ratio than the yield curves for assetholders. This feature follows since the calibrated value of the negative loading of the interest rate on S_t , b , is calibrated to a higher value for non-assetholders compared to assetholders. This also implies that for similar variation in the surplus consumption ratio,

¹⁹One period real yields can be calculated directly from Eq. (8), while longer maturity bond prices (or yields) and the price-dividend ratio must be calculated using numerical methods. Appendix B in Wachter (2005) shows that the absence of an explicit solution does not arise from the discrete time assumption.

the standard deviation on the non-assetholders' yields will be larger than for assetholders' yields.

Since we have calibrated b to be positive, one would, in general, expect the 5 year yield to be above the 3 month yield and this is also confirmed from the figure. However, this relationship is reversed for very small values of the surplus consumption ratio. That is, in deep recessions the model predicts a flat or slightly decreasing yield curve. This result has also been reported in previous calibrations and might be a feature of the model. Also notice that the model is able to predict negative real yields, which is often observed in the data. However, the non-assetholder data seems to be able to generate unreasonably high and low real yields with real yields in excess of 100% in extreme recessions and -10% in extreme booms.

From Figure 1 it is not possible to determine if mean yields are higher for assetholders or non-assetholders and the standard deviations of yields cannot be assessed either, since the two data sets generate different series of the surplus consumption ratio in terms of both volatility and magnitude. Therefore, we calculate the mean and standard deviation of 3 month and 1-5 year yields for assetholders, non-assetholders, as well as for the aggregate consumption series. The results are shown in Table 3. The 3 month yield is matched relatively well to its empirical value of 2.245% for all consumption measures. However, this is expected given that the calibration ensures that data is matched at the long-run steady state value of the surplus consumption ratio, $\bar{s}$. From the table it is seen that when using non-assetholder consumption data, the model is able to match yields better than assetholder data in the short end where the simulated 1 year yield is 2.578% for assetholders and 2.703% for non-assetholders, which should be compared to 2.867% in the data. However, non-assetholder data overshoots dramatically in the long end of the yield curve where a value of 4.622% is simulated for the 5 year yield compared to 3.869% in the data. With assetholder data the model implies a 5 year yield of 3.991%, which is very close to its empirical value. The assetholder series also achieves much better results than the aggregate consumption series in matching mean yields. However, it is somewhat surprising that using non-assetholder data generates better results for mean yields compared to using aggregate consumption.

From Table 3 it is also seen that the simulated standard deviations are significantly closer to those observed in the data for asseholders compared to both non-asseholder and aggregate consumption data. Simulating the model using the asseholder data, we obtain standard deviations of 3.164% for the 1 year yield and 2.964% for the 5 year yield, compared to 2.911% and 2.823%, respectively, in the data, and for all maturities the simulation using asseholder data gives results for the standard deviation of yields close to the empirical data. In comparison, the standard deviations of yields calculated using non-asseholder data are around 40% larger than empirical standard deviations. In general, it is seen that the standard deviations of empirical yields are decreasing in maturity. This feature is matched by both asseholder and non-asseholder data, but not so for aggregate consumption. Matching the decreasing standard deviation with aggregate consumption does not seem to be possible. This result was also found in Wachter (2006).

The top row in the table reports the sum of squared differences between model simulated moments and data moments. It includes the 1-5 year yields, since the 3 month yield is matched in approximation by construction. The sum of squared errors is of considerably less magnitude for both means and standard deviations for asseholders than for non-asseholders and aggregate consumption. Finally, it is seen from the table that for all consumption series the yield curve is steeper than the yield curve indicated from the data. This is particularly prominent for the non-asseholder and the aggregate consumption series and if bonds with longer time to maturity are simulated, the model will significantly overshoot the observed yields.²⁰

Our hypothesis of asseholder consumption data containing much more information about real bond yields is confirmed by these results. It is somewhat puzzling though that the aggregate consumption data does not consistently yield better simulation results than non-asseholder data. The ability of the model to match data moments of real yields with asseholder consumption is impressive in the sense that the model has been calibrated to match bond holding returns and not means or standard deviations of yields.

As indicated from Figure 2, the volatility of consumption is an important factor in de-

²⁰This feature of a too steep yield curve is also seen clearly from the results for nominal yields in Table 4 in Wachter (2006).

termining the slope of the yield curve. High volatility implies a lower slope and this is part of the reason why the aggregate consumption measure has a too high slope. It is also the case that the non-assetholder yield curve is much steeper than the assetholder yield curve. Once again, we see from the figure that part of the explanation is due to the lower volatility of non-assetholder consumption. From the figure it is also clear that the model produces a (close to) linear yield curve structure on average.

We now turn our focus to equity data. Figure 3 shows the price-dividend ratio as a function of the surplus consumption ratio S_t and it is seen that it increases in S_t on the entire interval for both the assetholder and non-assetholder data. This is in accordance with empirical data since the price-dividend ratio is correlated with the business cycles of the economy (see e.g. Fama and French (1989)). The two series are quite similar; however, the price-dividend ratio seems slightly steeper for the non-assetholder data although one has to take into account that $S_{\max}^{NAH} < S_{\max}^{AH}$. This could indicate that the price-dividend ratio is more volatile for the non-assetholder data than for the assetholder data. This is confirmed in Table 4 where it is seen that the volatility of the price-dividend ratio is 0.260 for non-assetholder data while it is 0.232 for assetholder data. This value should be compared to a value of 0.389 in the data and it thus seems as if the limited participation series are not able to capture the volatility of the empirically observed price-dividend ratio. The Sharpe ratio is matched to the empirical value of 0.533 by construction. However, the mean and standard deviation of excess returns both seem to be matched well by all series too. In particular, the aggregate consumption data matches these moments well.

In previous calibrations of the model a low value of the utility curvature parameter, γ , has not translated into a low value of the Arrow-Pratt measure of relative risk aversion. However, from the last row in Table 4 it is seen that the relative risk aversion at steady state, $\gamma/\bar{S}$, needed to explain both the risk-free rate puzzle of Weil (1989) and the equity premium puzzle of Mehra and Prescott (1985) is about 8 for assetholder consumption data, i.e. the relative risk aversion at steady state does not exceed 10, as considered plausible by Mehra and Prescott (1985). Using aggregate consumption data, the model requires unreasonably large values of relative risk aversion, which is consistent with the findings in previous literature.

4.3 Relative Risk Aversion

The previous section showed that the model performs much better when using assetholder consumption compared to the conventional aggregate consumption measure. In particular, the data implies that a reasonable low relative risk aversion is sufficient to explain the risk premium puzzle. Similar results have been hard to obtain in previous studies. One exception, though, is Malloy et al. (2009) who also utilize limited participation data, but rely on recursive utility with long-run consumption.

It is fruitful to investigate the driving factors of our results further. We therefore study the partial effect of varying the parameters by simulating the model and calculating the relative risk aversion for a grid of values. Figure 4 shows the partial analysis for the parameters σ , ϕ , γ and b . The top left plot shows that the relative risk aversion is decreasing in the volatility of consumption for all measures of consumption. From this it seems that it is the high volatility observed in limited participation data that makes the model able to explain the risk premium puzzle with a low value of relative risk aversion. We also see that the difference in the other parameters over the different consumption measures has limited effect if σ is high, while it has more effect for a low σ .

The top right plot shows that for all consumption measures the relative risk aversion is decreasing in the habit persistence parameter and more so for aggregate consumption than for limited participation consumption. For very high habit persistence the relative risk aversion is unidentified and we therefore exclude the unidentified region. The bottom left plot shows that for all consumption measures the relative risk aversion is increasing in the utility curvature parameter as expected, while the bottom right plot shows that the relative risk aversion is decreasing in b . The plots for ϕ , γ , and b make it clear that the relative risk aversion is much more sensitive to changes in these parameters for the aggregate consumption compared to the limited participation measures. This indicates that the high volatility of the limited participation measures makes the relative risk aversion less sensitive to the remaining parameters. The two parameters not considered in Figure 4, g and δ , have no influence on the relative risk aversion.

Besides the steady state level of relative risk aversion, which the sensitive analysis in

Figure 4 focuses on, it is also interesting to examine how the relative risk aversion varies over time in the model. In Figure 5 we show the time series of the relative risk aversion using aggregate and assetholder consumption, respectively. From the figure it is clear that the overall level of risk aversion is much higher for aggregate consumption than for assetholder consumption, as already discussed. With aggregate consumption the relative risk aversion is in the range from 54 to 174, whereas with assetholder consumption the relative risk aversion is in the range from 6.8 to 9.3. The figure also illustrates that with aggregate consumption we observe very rapid variations in the relative risk aversion. As an example of this, the relative risk aversion is about 175 in 1997 and then falls to about 75 in 2000, which is a dramatic change in relative risk aversion within a relatively short time period. With assetholder consumption, on the other hand, we observe much more economically plausible changes in relative risk aversion over time. Thus, the figure illustrates some important differences in the time series of relative risk aversion between aggregate and assetholder consumption. However, the figure also shows that the two consumption measures are similar in the sense that they share the same business cycle pattern. In fact, the correlation between the two series of relative risk aversion is as high as 0.49.

5 GMM Estimation

In the previous section we have shown by calibration and simulation that the model works better in a variety of different dimensions using assetholder consumption in comparison to non-assetholder and aggregate consumption. In order to obtain further information about the performance of the model under the alternative consumption measures, we now turn to formal estimation and testing of the model using GMM. We use the following moment

conditions:

$$0_{N \times 1} = E [\mathbf{R}_{t+1} \delta e^{-\gamma\{g+(\phi-1)(s_t-\bar{s})+[1+\lambda(s_t)]v_{t+1}\}} - 1] \quad (16)$$

$$0 = E \left[r_{f,t+1} + \log(\delta) - \gamma g + \frac{\gamma(1-\phi) - b}{2} + b(s_t - \bar{s}) \right] \quad (17)$$

$$0 = E \left[y_{2,t} - y_{1,t} + \frac{1}{2} \begin{bmatrix} -0.5(\gamma(1-\phi) - b) \\ +(\gamma(1-\phi) + b(\phi-2))(s_t - \bar{s}) \\ +0.5\sigma^2 [b\lambda(s_t) - \gamma - \gamma\lambda(s_t)]^2 \end{bmatrix} \right] \quad (18)$$

From the unconditional Euler equation, and using the stochastic discount factor in Eq. (7), we form the moment conditions in Eq. (16) where $\mathbf{R}_{t+1}$ contains real gross returns on a vector of N assets. The purpose is to examine whether the model has the ability to explain the variation in returns on a cross section of bonds. We want to make sure that the parameters are estimated such that the model tries to fit the risk-free rate. Thus, using the specification of the risk-free rate in Eq. (8), we include the moment condition in Eq. (17). We are also interested in estimating the parameters such that the model matches the yield spread, and therefore we incorporate a moment condition for the two period yield spread, $y_{2,t} - y_{1,t}$, in Eq. (18).²¹ Taken together, the choice of moment conditions enables us to examine whether the model has the ability to jointly explain the returns on a cross section of bonds, the risk-free rate as well as the yield spread.

We collect the sample moment conditions in a vector $\mathbf{g}_T$, and we then estimate the parameters by minimizing the objective function $\mathbf{g}_T' \mathbf{W} \mathbf{g}_T$. As weighting matrix, $\mathbf{W}$, we use the identity matrix to make sure that GMM pays equal attention to all moment conditions. By using a pre-specified weighting matrix, such as the identity matrix, it becomes possible to compare the magnitude of the estimated pricing errors across the different consumption measures that we apply. Such a comparison would not be possible if, instead of the identity matrix, the statistically optimal weighting matrix of Hansen (1982) were used. This is due to the fact that it places the highest weight on the linear combination of moments with the lowest variance, and this linear combination may change from one consumption measure to

²¹See Møller (2009) for a derivation of the two period yield spread.

the next.²²

There is no observable data on the s_t process due to the fact that the habit level of the households is not a directly observable process. In order to generate the s_t process, we use the steady state value as the initial value, and based on data on consumption and a set of parameter starting values, we obtain the s_t process recursively using Eqs. (2), (3) and (4). The GMM procedure then iterates over the parameters until convergence. We treat δ , γ , ϕ , and b as free parameters to be estimated in the GMM system.

Table 5 shows the GMM estimation results. The cross section of asset returns is given by the aggregate stock market return and returns on government bonds with maturities of 1, 2, 5, 7, and 10 years. With the moment conditions for the risk-free rate and the yield spread, we have eight moment conditions in total, which gives us four overidentifying restrictions because there are four parameters to be estimated.

Panel A of Table 5 gives the GMM estimates of the parameters. The subjective discount factor δ is estimated to be less than 1 for all consumption measures, which is in line with the calibration analysis. The utility curvature parameter γ is estimated to be positive, but it is not significantly different from zero for any consumption measure. As in the calibration, the overall level of the risk aversion depends strongly on the consumption measure. The steady state risk aversion is 19.4, 28.4, and 155.5 when using assetholder, non-assetholder, and aggregate consumption, respectively. These levels of risk aversion are somewhat higher than in the calibration analysis, but it is still the case that assetholder consumption produces the lowest steady state risk aversion.²³ The GMM estimates of the persistence parameter ϕ are high, especially for aggregate consumption, and statistically significant in all cases. The high values of ϕ imply that habit only adjusts very slowly to changes in consumption, which is consistent with our calibration analysis as well as previous literature. Finally, we see from Panel A of Table 5 that the parameter b is estimated with a lot of uncertainty, suggesting

²²Since we do not use the statistically most efficient weighting matrix, but a suboptimal weighting matrix in a statistical sense, the standard way of calculating standard errors and the J -test of overidentifying restrictions does not apply. Hence, we use the general formulas that apply to suboptimal estimates, see Chapter 11 in Cochrane (2005).

²³As argued by Parker (2001), the sample period with CEX data is "one of unusually high returns in many years" and adjusting for this would naturally lead to lower risk aversion estimates.

that GMM has difficulties in identifying b .

Now we turn to the fit of the model as presented in Panel B of Table 5. We see that aggregate and non-assetholder consumption both generate a close fit of the aggregate stock market return, while assetholder consumption slightly underestimates the aggregate stock market return. When it comes to pricing bonds, however, it is clear that assetholder consumption performs better than the two other consumption measures. Average bond returns increase in maturity from 1.01% at the 1 year maturity to 1.72% at the 10 year maturity, but with non-assetholder as well as aggregate consumption the model basically generates no action along the bond maturity dimension, as the model predicted average bond returns range from 1.41% to 1.50%. With assetholder consumption the model predicted average bond returns increase in maturity from 1.37% at the 1 year maturity to 1.74% at the 10 year maturity, implying that assetholder consumption improves the fit of the model although the pricing errors are still sizeable for short-term bonds. Besides the cross section of returns, the GMM estimation also includes moments for the risk-free rate and the yield spread. As expected, all three consumption measures generate low and reasonable average risk-free rates. With aggregate consumption, however, the negative GMM estimate of b implies that the model predicts a counterfactual negative slope of the yield curve.

The overall evidence of the GMM estimation suggests that, regardless of consumption measure, the model has some difficulties in fully explaining the imposed moment conditions. The GMM estimation also reveals that the model's ability to price bonds strongly depends on the consumption measure. With assetholder consumption the model does a reasonable job pricing bonds, while with non-assetholder and aggregate consumption the model implies rather large pricing errors. Consistent with this evidence, the J -test of overidentifying restrictions statistically rejects the model at conventional significance levels when using aggregate and non-assetholder consumption, while with assetholder consumption the model is just rejected at a 5% level, but not at a 1% level.

6 Conclusion

In this paper we investigate habit-based asset pricing taking into account limited asset market participation. We use calibration and simulation techniques as well as GMM estimation to evaluate the performance of the model. We show that the ability of the model to price bonds improves substantially when using the consumption of households who do actually invest in assets in comparison to using the consumption of non-assetholders and aggregate consumption. In particular, with assetholder consumption the model provides a much better fit of means and standard deviations of bond yields as well as bond returns compared to non-assetholder or aggregate consumption.

Furthermore, we show that with assetholder consumption the habit model has the ability to explain the large equity premia using an economically plausible level of relative risk aversion. We find that the necessary level of relative risk aversion is only about 8 at steady state if we use the consumption of those households who invest in asset markets. This more plausible level of relative risk aversion is an intuitive result as those households who invest in asset markets are exposed to more risk than those who do not.

The main driving force behind our results is a higher level of volatility of consumption for households investing in asset markets. We show that this higher level of volatility helps in explaining the slope of the empirical yield curve and is necessary to achieve a reasonable low level of relative risk aversion.

References

- Battistin, E., 2004. Errors in survey reports of consumption expenditures. Working paper, Institute for Fiscal Studies, London.
- Brav, A., Constantinides, G., Geczy, C., 2002. Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy* 110, 793-824.
- Campbell, J.Y., Cochrane, J.H., 1995. By force of habit: A consumption-based explanation of aggregate stock market behavior. NBER Working Paper No. 4995.
- Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205-251.
- Cochrane, J.H., 2005. *Asset Pricing*. Revised edition, Princeton University Press.
- Cochrane, J.H., Piazzesi, M., 2005. Bond risk premia. *American Economic Review* 95, 138-160.
- Fama, E.F., Bliss, R.R., 1987. The information in long-maturity forward rates. *American Economic Review* 77, 680-692.
- Fama, E.F., French, K.R., 1989. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25, 23-49.
- Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50, 1029-1054.
- Malloy, C.J., Moskowitz, T.J., Vissing-Jørgensen, A., 2009. Long-run stockholder consumption risk and asset returns. *Journal of Finance* 64, 2427-2479.
- Mankiw, N.G., Zeldes, S.P., 1991. The consumption of stockholders and nonstockholders. *Journal of Financial Economics* 29, 97-112.
- Mehra, R., Prescott, E.C., 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15, 145-161.
- Møller, S.V., 2009. Habit persistence: Explaining cross-sectional variation in returns and time-varying expected returns. *Journal of Empirical Finance* 16, 525-536.

- Parker, J. A., 2001. The consumption risk of the stock market. *Brookings Papers on Economic Activity* No. 2, 279-333.
- Wachter, J.A., 2005. Solving models with external habit. *Finance Research Letters* 2, 210-226.
- Wachter, J.A., 2006. A consumption-based model of the term structure of interest rates. *Journal of Financial Economics* 79, 365-399.
- Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics* 24, 401-421.

Table 1: Summary statistics of consumption data.

	AH	NAH	Agg.
<i>Correlations</i>			
AH	1.0000	0.3416	0.1527
NAH		1.0000	0.1178
AGG			1.0000
<i>Parameters</i>			
g (%)	0.0657	-0.0975	0.5170
σ (%)	3.0099	2.1906	0.3567

The first panel shows correlations between assetholder (AH), non-assetholder (NAH) and aggregate (Agg.) consumption growth. The second panel shows means and standard deviations of the consumption series. The quarterly consumption growth rate, g , and standard deviation of quarterly consumption growth, σ , are in percentages.

Table 2: Calibrated parameter values.

	AH	NAH	Agg.
<i>Parameters</i>			
γ	3.6954	4.6830	4.1591
b	0.0544	0.0778	0.0319
ϕ	0.9689	0.9689	0.9689
<i>Derived parameters</i>			
δ	0.9671	0.9569	0.9677
$\bar{s}$	-0.7936	-0.9319	-3.0466
$s_{\max}$	-0.3959	-0.5094	-2.5477

The first panel shows the calibrated values of the independent parameters of the model. The second panel shows parameters derived from the independent parameters. δ is jointly determined by the values of γ and b such that at $s_t = \bar{s}$ the real risk-free rate given in Eq. (8) equals the mean real risk-free rate in the data. $\bar{s}$ and $s_{\max}$ are calculated from Eqs. (5) and (6), respectively.

Table 3: Means and standard deviations of real zero coupon bonds.

Maturity	Mean				Std. dev.			
	AH	NAH	Agg.	Data	AH	NAH	Agg.	Data
sqr. err.	0.2309	0.8346	2.3023	—	0.1570	6.1605	6.2878	—
3m	2.2867	2.3030	2.3132	2.2450	3.1644	4.1831	3.7062	2.7793
1y	2.5779	2.7078	2.7431	2.8673	3.1370	4.1331	3.7865	2.9105
2y	2.9531	3.2243	3.3244	3.2442	3.0978	4.0629	3.8910	2.9154
3y	3.3137	3.7152	3.9131	3.5200	3.0558	3.9891	3.9915	2.8877
4y	3.6597	4.1808	4.5067	3.7294	3.0110	3.9119	4.0870	2.8548
5y	3.9912	4.6217	5.1024	3.8688	2.9636	3.8319	4.1762	2.8227

The table shows means and standard deviations of real zero coupon bonds in the model and in the data. The 3 month yield could just as well have been calculated from Eq. (8), but the reported 3 month yields are from simulated bond prices. Longer maturity bonds are also calculated based on Eq. (12). The calibration error is calculated as the sum of squared differences between simulated moments and data moments. Mean yields are annualized and standard deviations are quarterly standard deviations of annualized yields.

Table 4: Statistics for the equity market.

Statistic	AH	NAH	Agg.	Data
$Sharpe$	0.5329	0.5329	0.5328	0.5328
$E(r_m - r_f)$	0.0693	0.0733	0.0859	0.0832
$\sigma(r_m - r_f)$	0.1300	0.1375	0.1623	0.1561
$\sigma(p - d)$	0.2316	0.2601	0.3206	0.3886
$\gamma/\bar{S}$	8.1718	11.8917	87.5228	—

The table shows statistics for the aggregate stock market from the model and in the data. Returns are continuously compounded and the mean and standard deviation of market excess returns are in annualized percentages. In the data the price-dividend ratio is calculated by dividing the price of the market index by the sum of the dividends over the previous year. In the model the price-dividend ratio is calculated from Eq. (13).

Table 5: GMM estimation results.

<i>Panel A: Parameter estimates</i>				
Parameter		AH	NAH	Agg.
δ		0.8412	0.8153	0.9579
		(0.3557)	(0.2019)	(0.2488)
γ		3.0064	2.5732	20.6170
		(3.8165)	(1.5885)	(13.7540)
ϕ		0.8812	0.8429	0.9919
		(0.2300)	(0.1908)	(0.1475)
b		0.0162	0.0164	-0.1407
		(0.6021)	(0.6644)	(3.0023)
<i>Panel B: Model fit</i>				
Moment	Data	AH	NAH	Agg.
Market	2.67%	2.24%	2.59%	2.67%
10y	1.72%	1.74%	1.48%	1.50%
7y	1.74%	1.65%	1.43%	1.45%
5y	1.57%	1.64%	1.50%	1.45%
2y	1.20%	1.47%	1.45%	1.43%
1y	1.01%	1.37%	1.41%	1.42%
3m	0.59%	0.36%	0.63%	0.62%
Spread	0.06%	0.07%	0.06%	-0.08%
J -stat		9.883	18.178	14.187
		(0.042)	(0.001)	(0.007)

Panel A shows parameter estimates with standard errors in parentheses. Panel B shows the model fit. J -stat is Hansen's test of overidentifying restrictions with p-values in parentheses.

Figure 1: Bond yields as a function of surplus consumption.

The figure depicts yields as a function of the surplus consumption ratio S_t . 3 month (1 period) yields are shown in red and 5 year yields are shown in blue.

Figure 2: Yields as a function of volatility and time to maturity.

Figure 3: Simulated price-dividend ratios as a function of surplus consumption.

The figure depicts price-dividend ratios as a function of the surplus consumption ratio S_t .

The simulated price-dividend ratios are annualized by dividing by four.

Figure 4: Sensitivity analysis of the relative risk aversion.

The figure contains plots of partial analysis of the effect of changes in the parameters σ , ϕ , γ and b on the relative risk aversion at steady state. The relative risk aversion is calculated by simulating the model at 21 values for σ on the grid 0.003-0.05, ϕ on 0.95-0.983, γ on 2-13, and b on 0.01-0.09.

Figure 5: Time-varying relative risk aversion.

The figure shows the relative risk aversion using assetholder consumption (left scale) and aggregate consumption (right scale).

- 2010-32: Mogens Bladt and Michael Sørensen: Simple simulation of diffusion bridges with application to likelihood inference for diffusions
- 2010-33: Fernando Baltazar-Larios and Michael Sørensen: Maximum likelihood estimation for integrated diffusion processes
- 2010-34: Leonidas Tsiaras: The Forecast Performance of Competing Implied Volatility Measures: The Case of Individual Stocks
- 2010-35: Leonidas Tsiaras: Dynamic Models of Exchange Rate Dependence Using Option Prices and Historical Returns
- 2010-36: Robinson Kruse and Rickard Sandberg: Linearity Testing in Time-Varying Smooth Transition Autoregressive Models under Unknown Degree of Persistency
- 2010-37: Tom Engsted and Thomas Q. Pedersen: The log-linear return approximation, bubbles, and predictability
- 2010-38: Thomas Q. Pedersen: Predictable return distributions
- 2010-39: Rasmus Tangsgaard Varneskov: The Role of Dynamic Specification in Forecasting Volatility in the Presence of Jumps and Noisy High-Frequency Data
- 2010-40: Antonis Papapantoleon and David Skovmand: Picard Approximation of Stochastic Differential Equations and Application to Libor Models
- 2010-41: Ole E. Barndorff-Nielsen, Fred Espen Benth and Almut E. D. Veraart: Modelling electricity forward markets by ambit fields
- 2010-42: Robinson Kruse and Philipp Sibbertsen: Long memory and changing persistence
- 2010-43: Dennis Kristensen: Semi-Nonparametric Estimation and Misspecification Testing of Diffusion Models
- 2010-44: Jeroen V.K. Rombouts and Lars Stentoft: Option Pricing with Asymmetric Heteroskedastic Normal Mixture Models
- 2010-45: Rasmus Tangsgaard Varneskov and Valeri Voev: The Role of Realized Ex-post Covariance Measures and Dynamic Model Choice on the Quality of Covariance Forecasts
- 2010-46: Christian Bach and Stig Vinther: Habit-based Asset Pricing with Limited Participation Consumption