

Tilburg University

Intertemporal Consumption with Directly Measured Welfare Functions and Subjective Expectations

Kapteyn, A.; Kleinjans, K.; van Soest, A.H.O.

Publication date:
2008

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Kapteyn, A., Kleinjans, K., & van Soest, A. H. O. (2008). *Intertemporal Consumption with Directly Measured Welfare Functions and Subjective Expectations*. (CentER Discussion Paper; Vol. 2008-85). *Econometrics*.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



No. 2008–85

**INTERTEMPORAL CONSUMPTION WITH DIRECTLY
MEASURED WELFARE FUNCTIONS AND SUBJECTIVE
EXPECTATIONS**

By Arie Kapteyn, Kristin Kleinjans, Arthur van Soest

October 2008

ISSN 0924-7815

Intertemporal Consumption with Directly Measured Welfare Functions and Subjective Expectations¹

Arie Kapteyn, RAND

Kristin Kleinjans, University of Aarhus, USC, and RAND

Arthur van Soest, Tilburg University and RAND

Abstract

Euler equation estimation of intertemporal consumption models requires many, often unverifiable assumptions. These include assumptions on expectations and preferences. We aim at reducing some of these requirements by using direct subjective information on respondents' preferences and expectations. The results suggest that individually measured welfare functions and expectations have predictive power for the variation in consumption across households. Furthermore, estimates of the intertemporal elasticity of substitution based on the estimated welfare functions are plausible and of a similar order of magnitude as other estimates found in the literature. The model favored by the data only requires cross-section data for estimation.

JEL Codes: D91, D84, D12

Key words: Expectations, Consumption, Euler equations

¹ The authors are grateful to Rob Alessie for making available the cleaned income data used in the paper and to two anonymous referees for helpful and constructive comments.

1. Introduction

Modern empirical studies of intertemporal allocation of consumption usually rely on Euler equations. For the estimation of such models, one typically needs panel data on consumption, assumptions on how respondents form their expectations, and a parameterization of preferences (see, for example, Hall 1978, Browning and Lusardi 1996, Carroll 2001, and Attanasio and Low 2004). In this paper, we aim at reducing some of these estimation requirements by using subjective data on income expectations and data on the income levels respondents say they need to attain a given satisfaction level. The latter are taken as points on a contemporaneous utility function of consumption. We investigate if directly measured utility functions and expectations have explanatory power for consumption and savings behavior. An affirmative answer is useful because it implies that preference distributions can be estimated more easily than with traditional approaches, avoiding the need to make arbitrary assumptions about expectations. Moreover, we find that the data slightly favor a model that can be estimated on cross section data and does not require panel data.

In contrast to the conventional approach, both utility functions and expectations are measured directly by asking subjective questions to survey respondents. By combining the information thus obtained with data on consumption (computed from income and saving), we are able to test if these directly measured utility functions and expectations have explanatory power for consumption behavior. Economists have long been skeptical of the use of subjective responses in questionnaires that do not refer to objective phenomena, and to which extent such responses help to explain behavior is an open question. We investigate whether these models can be used to explain consumption and to analyze the sensitivity of saving and consumption to the interest rate (i.e., the intertemporal rate of substitution).

As a specification of preferences, we adopt the individual welfare function, a concept introduced by Van Praag (1968) and operationalized in numerous papers since, including Van Praag (1971), Van Praag and Kapteyn (1973), Kapteyn and Wansbeek (1985), Groot et al. (2004), Van Praag and Ferrer-i-Carbonell (2004), and Rablen (2008). In particular, we use the individual welfare function of income (henceforth

WFI), which in our dynamic framework is interpreted as a welfare function of consumption. In a static context, the WFI represents the satisfaction an individual attaches to a certain income (or consumption) level, measured on a continuous scale from 0 to 1. In Section 3, we will describe in some detail how the WFI is constructed from answers to a set of relatively straightforward questions.

An Euler equation relates the marginal utility of current consumption to the expected marginal consumption of the next period (which, in this paper, is next year). Thus, writing down an Euler equation for intertemporal allocation of consumption does not only require knowledge of the utility function but also of expectations. Therefore, we do not only elicit the individual utility function of consumption directly, but we also use direct information on expectations, following the approach pioneered by Dominitz and Manski (1997). We use data from the Dutch DNB household panel survey, which has the unique feature that it includes questions on both expectations of future income and incomes needed to attain given satisfaction levels, thus enabling us to directly measure both individual expectations and preferences.

Different assumptions can be made regarding the evolution of preferences over time. More precisely, in solving the intertemporal consumption problem, a consumer has to make assumptions about his or her future preferences. A “myopic” consumer may assume that tomorrow’s preferences are the same as today’s. A (super) rational consumer, on the other hand, may be able to predict tomorrow’s preferences perfectly.

Our results suggest that the individually measured welfare functions and expectations have predictive power for the cross-section variation in consumption. Estimates of the intertemporal elasticity of substitution based on the estimated welfare functions are of a similar order of magnitude as those found in the literature.

The remainder of the paper is structured as follows. In the next section, we describe the panel data used in this paper. In Section 3, the welfare function of income (or consumption) and its measurement are described and in Section 4, we explain how we measure expectations. In Section 5, the Euler equations are derived under the assumptions that intratemporal utility can be described by a lognormal welfare function and that the subjective distribution of future consumption follows a lognormal distribution. The empirical strategy is explained in Section 6. We discuss the empirical

results in Section 7, where we also compare implied intertemporal elasticities of substitution to existing findings in the literature. Section 8 concludes.

2. Data

The DNB Household Survey (DHS), formerly known as the CentER Savings Survey, is a Dutch panel survey that started in 1993. The survey is conducted by CentERdata and administered over the Internet. If a potential participant has no Internet access, he or she is provided with access through a so-called set top box (also called Web TV or Internet Player) that connects to the Internet via the telephone and a television set. The survey consists of six modules and asks a variety of questions about demographics, health, income, assets, and economic and psychological concepts. For our research, data up to and including the year 2007 were available.

Over the years several changes took place in technology used and sample selection. The current set-up of the panel is in place since 2000 and hence we use the waves from 2001 onward.² Appendix A describes the sample selection (Table A.1) and the number of observations by survey year (Table A.2). We use the unbalanced sample so the number of observations changes across waves, due to both attrition and refreshment. Our final sample has 9,293 observations of 3,075 individuals.

The survey does not directly measure consumption and the consumption measure used in this paper is unfortunately not ideal: Consumption was constructed as the difference between (self-reported) income and savings. For savings, we used the answers to two questions. If the first question, “*Did your household put any money aside in the past 12 months?*”³ was answered affirmatively, the respondent was asked a second question: “*About how much money has your household put aside in the past 12 months?*” The respondent was asked to choose one of seven different brackets, which

² In 1997 the panel was moved from its original institute associated with the University of Amsterdam to CentERdata. In 2000 the technology used for the interviewing of respondents was thoroughly modernized. Each of these changes led to substantial sample loss. Starting in 2001, the panel was gradually rebuilt.

³ Before 2004, the wording of the question was slightly different: “*Did you put any money aside in the past 12 months?*” The follow-up question is the same in all waves.

differed across some of the waves. From these answers and the self-reported (net) income, we constructed lower and upper limits on consumption. Note that, depending on the answers, one of the two limits might be missing. For instance, if a respondent reports that no money was put aside during the last twelve months, this means that savings can have been zero or negative and hence consumption must have been at least equal to income.

Table 1: Summary Statistics

Variable	N	Mean
Age	9,293	50.2 (14.7)
Female	9,293	0.45
Education	8,763	
Primary or less		0.05
Lower level		0.25
Intermediate vocational		0.22
Intermediate general		0.14
Higher vocational		0.04
University		0.29
Other		0.01
Income ¹	9,293	31,231 (39,127)
Savings		
Lower limit	6,870	2,660 (4,423)
Upper limit	8,916	5,165 (6,493)
Consumption		
Lower limit	8,916	25,404 (29,443)
Upper limit	6,870	30,027 (44,108)

¹ *Income is measured after tax. The sample we work with includes observations with at least one non-missing consumption bracket. All income measures are annual and expressed in Euros of the year 2006. Standard deviations are given in parentheses (except for dummy variables).*

Table 1 provides some summary statistics for our sample. As is explained more fully in Appendix A, the number of observations varies across variables, partly because questions were asked in different modules that were administered during different time periods. As a result of this, some respondents answer some modules, but not others.

This explains, for example, the number of observations with missing educational level. The lower numbers of observations for the limits on savings and consumption are the result of the nature of the questions on savings discussed above. There is at least one limit available for each observation.

3. Welfare Functions of Income and Their Measurement

A WFI is measured by asking respondents in a survey a so-called Income Evaluation Question (IEQ). The formulation of the IEQ varies somewhat across surveys. The IEQ used in the survey on which the current paper is based reads as follows:

The next question again concerns the net income of the household, that is, the net income of all household members taken together. Consider the current situation of your household when answering this question.

Which NET income of the household would you, IN YOUR SITUATION, find very bad, bad, insufficient, sufficient, good, very good? Please provide annual incomes.

<i>VERY BAD</i>	<i>if the annual income would be about:</i>	<i>€.....</i>
<i>BAD</i>	<i>if the annual income would be about:</i>	<i>€.....</i>
<i>INSUFFICIENT</i>	<i>if the annual income would be about:</i>	<i>€.....</i>
<i>SUFFICIENT</i>	<i>if the annual income would be about:</i>	<i>€.....</i>
<i>GOOD</i>	<i>if the annual income would be about:</i>	<i>€.....</i>
<i>VERY GOOD</i>	<i>if the annual income would be about:</i>	<i>€.....</i>

The reference to the respondent's own situation invites the respondent to take into account any factor that influences income satisfaction. Such factors may include family composition, labor market status, health, etc. Numerous papers have been written dealing with these factors. See, for instance, Kapteyn and Wansbeek (1985) or Van Praag and Ferrer-i-Carbonell (2004) for overviews. In the current paper, the determinants of the answers are not our concern, but rather how these answers can be used to explain consumption.

Van Praag's theory assumes an isomorphism between utility theory and probability theory. Utility is assumed to be measurable on a $[0,1]$ -scale.⁴ In order to use the answers to the above IEQ to estimate a utility function (as Van Praag calls it, a welfare function), one needs to assign numerical values between zero and one to the verbal labels "very bad", "bad", "insufficient", "sufficient", "good", and "very good". Based on an information maximization argument, Van Praag (1971) proposes to assign numerical values such that each label represents an equal part of the $[0,1]$ -interval. The basic idea is that a respondent maximizes the information conveyed by the answers if he or she partitions the $[0,1]$ scale into equal intervals. One can formalize this argument, but basically if a respondent would not partition the $[0,1]$ interval this way then there would be parts of the interval that have a higher information content than others (see, for instance, Kapteyn and Wansbeek, 1985, for a detailed explanation). In the formulation used here, the equal interval assumption means that "very bad" is assigned the value $1/12$, "bad" $3/12$, "insufficient" $5/12$, "sufficient" $7/12$, "good" $9/12$, and "very good" $11/12$.

The IEQ asks for income levels providing a certain welfare level, and the underlying theory in, e.g., Van Praag (1971) is static. In the standard life-cycle model, utility in a given period depends on consumption in that period. In the intertemporal context of saving and consumption decisions, it therefore seems natural to interpret the (static) welfare function of income as a welfare function of within period consumption. The interpretation of the IEQ is then that respondents' answers reflect the consumption levels that would yield a certain amount of within period utility, assuming that the question refers to a situation without (positive or negative) savings, so that the reported income amounts are actually amounts of total consumption expenditure. An equivalent formulation would be to assume that the income levels elicited in the IEQ are levels of permanent income, which in a standard life cycle model equals consumption.

⁴ Of course the $[0,1]$ interval is arbitrary, but the crucial assumption is that utility is bounded from above and from below. So, for instance, CARA or CRRA utility functions would not fit in this framework.

From now on, we will interpret the WFI in that way, i.e. as representing the utility of consumption in period t . In order to use preferences as an explanation for consumption choices, we will follow Van Praag (1971) and summarize the WFI in two preference parameters (μ_t and σ_t). Denote consumption in period t by x_t . Using the isomorphism with probability theory and invoking a Central Limit Theorem, Van Praag shows that under certain conditions the WFI can be approximated by a lognormal distribution function:⁵

$$U(x_t, \mu_t, \sigma_t) = \int_0^{x_t} \frac{1}{z} \frac{1}{\sigma_t \sqrt{2\pi}} \exp\left[-\frac{1}{2} \left(\frac{\ln z - \mu_t}{\sigma_t}\right)^2\right] dz.$$

The parameters μ_t and σ_t are individual preference parameters, which can be estimated from the IEQ above in a straightforward way: Let Φ be the standard normal distribution function, and denote the answers to the IEQ at time t by z_{jt} , $j=1, \dots, 6$. Furthermore, denote the numerical labels corresponding to the verbal labels “very bad”, “bad”, etc. by w_j , $j=1, \dots, 6$; that is, $w_1 = \frac{11}{12}$, $w_2 = \frac{9}{12}$, etc. Then we have:

$$w_j = \Phi\left(\frac{\ln z_{jt} - \mu_t}{\sigma_t}\right), \quad j=1, \dots, 6.$$

Rewriting and allowing for measurement error ε_{it} in the responses then gives:

$$\ln z_{jt} = \mu_t + \sigma_t \Phi^{-1}(w_j) + \varepsilon_{it}, \quad j=1, \dots, 6.$$

⁵ Throughout, the index representing the individual is suppressed. Without attempting a full expose of the Central Limit Theorem argument, the basic idea can be sketched as follows. Total consumption can be seen as the consumption of a bundle of characteristics. Each of the characteristics is evaluated by a partial welfare function, taking values between 0 and 1. An efficient representation of characteristics would be where they are independent. In a framework that is isomorphic to probability theory, this means that the evaluation of a bundle of characteristics is the product of the partial welfare functions of each of the characteristics. Next, a production technology of the following form is assumed:

$f_1(x_1) \cdot f_2(x_2) \dots f_n(x_n) = x$, where x_1, x_2, \dots, x_n are characteristics and x is total expenditures. Equivalently, $\ln f_1(x_1) + \ln f_2(x_2) + \dots + \ln f_n(x_n) = \ln x$. Once again invoking the isomorphism with probability theory, we can interpret $\ln x$ as the sum of n independently distributed random variables. Thus, by using the Central Limit Theorem, one finds that $\ln x$ is “distributed” as a normal variable. In the context of a utility framework, this means that the functional form of the evaluation of $\ln x$ is a normal distribution; by definition the evaluation of x then has the functional form of a lognormal distribution function.

Under the assumption that the measurement errors satisfy the classical linear model assumptions, OLS for each individual and each time period (on six observations) gives unbiased estimates for all (individual and year specific) parameters μ_t and σ_t . Note that we estimate parameters μ_t and σ_t for each respondent and for each time period separately.⁶

Table 2: Income Evaluation and Estimated Utility Function Parameters


Variable	Mean
Income	31,231 (39,127)
Log income	10.22 (0.48)
Income Evaluation	
Very bad	14,152 (7,317)
Bad	17,656 (7,949)
Insufficient	21,169 (9,012)
Sufficient	26,876 (11,303)
Good	33,654 (14,695)
Very good	47,887 (36,438)
Utility Function Parameters	
μ_t	10.02 (0.38)
σ_t	0.44 (0.22)

Working sample. N=9,293. All income measures are annual and expressed in Euros of 2006. Standard deviations in parentheses.

⁶ The linear estimating equation allows for various specification tests and comparisons with alternative functional forms. Over the years a substantial number of tests has been performed. These tests confirm that the lognormal distribution gives a very good approximation to the functional form of the WFI and supports the statistical assumptions underlying the linear model (and its estimation by OLS). See for instance: Van Herwaarden et al. (1977) and Kapteyn (1977).

Table 2 shows the means and standard deviations of the required incomes at each verbal label of the income evaluation question, as well as of the estimated utility function parameters for our sample. To make a comparison with actual income levels possible, the table also includes statistics for actual net household income. The average actual income level is close to what on average is considered a “good” income, in line with the notion that, on average, Dutch people are rather satisfied with their actual income. Table A.2 in the appendix shows how median actual and required incomes (in 2006 prices) evolve over time. Variation is limited and probably reflects variation in sample composition.⁷

Figure 1: Some Estimated Utility Functions



⁷ Van de Stadt et al. (1985) have analyzed the dynamics and the autocorrelation patterns in μ_t , emphasizing the roles of changes in family size and μ_{t-1} , where the latter can be interpreted as habit formation. Clark (1999) finds strong evidence of habit formation in subjective job and wage satisfaction.

Figure 1 shows some welfare functions based on the estimates in Table 2 for individuals with mean $\mu = 10.02$. In the current context, this means that a consumption level of $\exp(10.02) = \text{€}22,472$ is evaluated at satisfaction level 0.5, between “insufficient” and “sufficient.” Clearly, a consumer with a larger value of μ needs a higher consumption level to reach a given satisfaction level. The parameter σ determines how steep or flat a welfare function is. The smaller σ , the steeper is the welfare function. We have drawn three welfare functions, all with $\mu = 10.02$, but with different values of σ , namely 0.20 (the 10th percentile), 0.44 (the sample mean), and 0.72 (the 90th percentile).

4. Measurement of Expectations

Most surveys soliciting subjective expectations about future outcomes ask for point estimates. Since future outcomes are intrinsically uncertain, a single point estimate provides incomplete information – it says nothing about the dispersion of the respondent’s subjective distribution of the future outcome. Moreover, it is not clear which point estimate respondents give in answer to any such question; this could be, for example, the mode, the median, or the mean (cf., e.g., Manski, 2004).

This problem can be overcome if information is solicited about the subjective probability distribution of the future outcome considered. Dominitz and Manski developed an approach in which respondents first give upper- and lower bounds on their future outcomes, and are then asked for the probabilities that outcomes lie in specific intervals which are subsets of the interval between the upper and lower bound. They applied this methodology to income expectations in the Survey of Economic Expectations (SEE); see, for example, Dominitz and Manski (1997). Das and Donkers (1999) applied this method to income expectations of Dutch households, using the same data source as we do, but earlier waves.

Specifically, respondents in the DHS were asked the following questions:

“We would like to know a bit more about what you expect will happen to the net income of your household in the next 12 months.

What do you expect to be the LOWEST total net income your household may realize in the next 12 months? Please use digits only, no dots or commas.”

“What do you expect to be the HIGHEST total net income your household may realize in the next 12 months?”

Follow up questions were asked about subjective probabilities that income lies within five equally sized intervals between the reported lowest and highest possible income. Many respondents did not answer some of these probability questions, even when they did provide the lowest and highest income asked for in the question. We have therefore only used the maximum and minimum incomes provided. The subjective income distributions are imputed by linearly interpolating probabilities for the five intervals between the maximum and minimum, where the probability that income is below the lowest income is set to zero and the probability that the income is below the maximum is set equal to one.⁸

Following the procedure of Dominitz and Manski (1997) but using the imputed probabilities rather than the reported probabilities, we then take the imputed probability distribution as approximations to a lognormal distribution function. Thus, for each respondent, the lognormal distribution that gives the best fit to the imputed probabilities is determined using non-linear least squares on the six points. For each respondent, this gives least squares estimates of the log median and the log standard deviation of her or his subjective distribution (the two parameters of the lognormal distribution).

Table 3 shows the summary statistics for actual income, the range of possible future incomes (the lowest and highest possible amounts), and the estimated respondent-specific parameters of the lognormal distribution.

⁸ We repeated our analysis using instead the responses to the probability questions for the subsample of individuals who answered at least two of these questions. The results were very similar (not shown but available from the authors upon request).

Table 3: Summary Statistics for subjective income distributions

Variable	Mean
Income	31,231 (39,127)
Lowest expected net income	26,625 (12,817)
Highest expected net income	44,844 (1,068,575)
Log-median of subjective distribution in next year	10.21 (0.45)
Log-standard deviation of subjective distribution in next year	0.07 (0.10)
Median of subjective distribution in next year	35,001 (477,023)
Standard deviation of subjective distribution in next year	26,594 (572,191)
Number of observations	9,293

All measures are annual and expressed in Euros of 2006. Standard deviations in parentheses.

The mean of the lowest possible future net incomes is significantly lower than the mean actual income. Similarly, the mean of the highest possible future net incomes is significantly higher than the mean actual income. The average range between highest and lowest possible amount is about €22,000. There is substantial variation here, however. For example, for 12.6% of the sample minimum and maximum amount are identical, indicating no subjective uncertainty in future income. In general, the log standard deviations of the subjective income distributions are quite low, with a mean of 0.07. Table A.2 in the appendix shows how median and dispersion vary over time. Variation over time is limited and probably mainly reflects variation in sample composition.⁹

⁹ For an extensive analysis of the subjective income distributions and their determinants, see Das and Donkers (1999). They find that past income is the main determinant of the median of the subjective income distribution, together with labor force status. The same factors have an effect on subjective income uncertainty, as well as gender (women are less uncertain than men), and education (with the largest uncertainty for the highest education level).

5. Models of intertemporal consumption

Consider the following standard intertemporal utility maximization problem of a consumer at time t :

$$\text{Max } E_t \sum_{\tau=t}^T \left(\frac{1}{1+\delta} \right)^{\tau-t} U(x_\tau, \mu_\tau, \sigma_\tau) \text{ subject to } \sum_{\tau=t}^T \left(\frac{1}{1+r} \right)^{\tau-t} x_\tau = M, \quad (1)$$

where:

E_t is the expectation operator conditional on all information available at time t ,

r is the real interest rate,

δ is the time preference rate,

M are total lifetime resources,

x_τ is consumption in period τ , and

T is the time horizon.

$U(x_\tau, \mu_\tau, \sigma_\tau)$ is the utility of consumption x_τ ; in line with the discussion in Section 3, this function is assumed to be the distribution function of a lognormal distribution with parameters μ_τ and σ_τ :

$$U(x_\tau, \mu_\tau, \sigma_\tau) = \int_0^{x_\tau} \frac{1}{z} \frac{1}{\sigma_\tau \sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{\ln z - \mu_\tau}{\sigma_\tau} \right)^2 \right] dz. \quad (2)$$

Note that this specification implies that preferences can change over time (through μ_τ and σ_τ). Our empirical strategy described below will deal with this automatically, since it does not restrict the variation of μ_τ and σ_τ over time or across respondents.

The first order conditions of consumption smoothing are:

$$U'(x_t, \mu_t, \sigma_t) = \frac{1+r}{1+\delta} E_t U'(x_{t+1}, \mu_{t+1}, \sigma_{t+1}). \quad (3)$$

To be able to characterize the solution to the first order conditions, we make assumptions about the distribution of the random variables that determine the expectation on the right hand side of (3). We will assume that this leads to a

distribution of x_{t+1} that is lognormal with parameters m_{t+1} and s_{t+1}^2 .¹⁰ Under these assumptions, the following expression for consumption x_t can be derived (see Appendix B for the derivation):

$$\ln x_t = \mu_t - \sigma_t^2 + \sigma_t \sqrt{(\sigma_t^2 - 2\mu_t) - (\sigma_{t+1}^2 - 2\mu_{t+1}) + \ln \frac{s_{t+1}^2 + \sigma_{t+1}^2}{\kappa^2 \sigma_t^2} + \frac{[m_{t+1} - (\mu_{t+1} - \sigma_{t+1}^2)]^2}{s_{t+1}^2 + \sigma_{t+1}^2}}. \quad (4)$$

Here we have defined $\kappa \equiv \frac{1+r}{1+\delta}$.

In order to use (4) as the basis for the empirical work, we need to replace the individual specific parameters by their estimates.¹¹ The parameter estimates for μ_t and σ_t were described in Section 3. For the future utility function in period $t+1$, we will use the estimates μ_{t+1} and σ_{t+1} in two out of the three models we consider (Models 1 and 3, defined in the next section), assuming that the consumer realizes that preferences change over time and accounts for this when making consumption decisions. In one model (Model 2), we will take $\mu_t = \mu_{t+1} = \mu$ and $\sigma_t = \sigma_{t+1} = \sigma$. That is, we assume that the consumer makes the consumption decision in period t under the (incorrect) assumption that preferences in periods t and $t+1$ are the same. Thus, the consumer plans as if preferences do not change; this is what we call “myopic” behavior in Section 1. In this model, (4) simplifies to

$$\ln x_t = \mu - \sigma^2 + \sigma \sqrt{\ln \left(\frac{s_{t+1}^2 + \sigma^2}{\kappa^2 \sigma^2} \right) + \frac{[m_{t+1} - (\mu - \sigma^2)]^2}{s_{t+1}^2 + \sigma^2}}. \quad (5)$$

The other choice we have to make is what to do with the parameters of the respondent’s subjective future consumption distribution used in making the consumption decision, m_{t+1} and s_{t+1} . In Models 1 and 2, we will replace these by their

¹⁰ Empirical evidence for the lognormality of the distribution of consumption is provided by Battistin et al. (2007).

¹¹ We will ignore the potential measurement error due to the fact that these are estimates.

estimates discussed in Section 4. Thus we interpret the subjective future income distribution as a distribution of future consumption. Clearly, this can only be approximately correct, since after all the whole idea of intertemporal consumption smoothing is to break the contemporaneous link between income and consumption. Still, we expect the future income distribution and future consumption distribution to be related, so that the future income distribution can be seen as a proxy for the future consumption distribution. The reason for this correlation is uncertainty: shocks in (permanent) income will lead to shocks in (permanent) consumption.

An alternative to Models 1 and 2 is based upon the theory of preference formation introduced by Kapteyn (1977). This theory states that the respondent's utility of consumption is purely driven by the rank in a perceived cross-sectional consumption distribution. One variant of that may be that parameters of the utility function in the next period coincide with the subjective distribution of consumption, i.e.: $m_{t+1} = \mu_{t+1}$ and $s_{t+1} = \sigma_{t+1}$. This is the assumption we will use in Model 3. In this case, we obtain:

$$\ln x_t = \mu_t - \sigma_t^2 + \sigma_t \sqrt{(\sigma_t^2 - 2\mu_t) - (\sigma_{t+1}^2 - 2\mu_{t+1}) + \ln 2 \frac{\sigma_{t+1}^2}{\kappa^2 \sigma_t^2} + \frac{1}{2} \sigma_{t+1}^2} . \quad (6)$$

In the empirical analysis, we will estimate all three non-nested models and test them against each other.

6. Empirical Strategy

As described above, we will denote the three different specifications as Models 1 through 3. To be precise:

- Model 1: (4) with direct estimates of μ_t , μ_{t+1} , σ_t , σ_{t+1} , m_{t+1} , and s_{t+1} ;
- Model 2: (5) with direct estimates of μ , σ , m_{t+1} , and s_{t+1} ;
- Model 3: (6) with direct estimates of μ_t , σ_t , μ_{t+1} , and σ_{t+1} .

An unusual feature of the models (4)-(6) is that they contain only one unknown parameter: κ . The parameters μ_t , μ_{t+1} , σ_t , σ_{t+1} , m_{t+1} , and s_{t+1} are estimated in a preliminary step for each individual separately, as described in Section 3. Thus in (4)-

(6) these are data. The task at hand is to investigate the plausibility of the models and to compare their empirical fit.

There are several ways to investigate the plausibility of the specifications. The first approach is to introduce parameters β_1 , β_2 and β_3 and consider empirical models of the following form:

$$\ln x_t = \beta_1 \mu_t + \beta_2 \sigma_t^2 + \beta_3 \sigma_t \sqrt{f_t(\kappa)} + u_t, \quad (7)$$

where u_t is an i.i.d. error term and $f_t(\kappa)$ is the term under the square root in (4), (5) or (6). If one of the models specified here is correct, we would expect that $\beta_1 = 1$, $\beta_2 = -1$ and $\beta_3 = 1$. Thus, we estimate each of the three specifications and consider the estimates of β_1 , β_2 and β_3 .

Since saving is measured in brackets (see Section 2), we used interval regression to estimate the models, i.e., we estimated ordered response models with normally distributed errors u_t and known cut-off points (the boundaries of the brackets in the questions about how much money one has put aside). Furthermore, as we are using a panel, we account for the fact that individuals may be in the data more than once by using a random effects specification - u_t consists of a time invariant random effect and an idiosyncratic error term, both assumed to be normally distributed and independent of regressors and of other errors.

Conditional on κ , model (7) is linear. Thus, we perform a grid search over values of κ and estimate (7) conditional on each value of κ . We then select the value of κ that gives the best likelihood value.

In a second approach, we compare the performance of the three different specifications by conducting several non-nested tests. The basic idea of these tests can be summarized by the following “artificial regression” (cf. Davidson and MacKinnon, 1981, 1993):

$$\ln x_t = \alpha m_1(z_t) + (1 - \alpha) m_2(z_t) + \varepsilon_t, \quad (8)$$

where $m_1(z_t)$ and $m_2(z_t)$ are alternative models explaining $\ln x_t$, and z_t are the explanatory variables. If $\alpha = 1$, $m_1(z_t)$ is the correct model, whereas if $\alpha = 0$, $m_2(z_t)$ is correct. We can rewrite (8) as

$$\ln x_t - m_2(z_t) = \alpha[m_1(z_t) - m_2(z_t)] + \varepsilon_t . \quad (9)$$

Applying this to (7), and using superscripts to denote Models 1 and 2, we obtain:

$$\begin{aligned} \ln x_t - \beta_1^2 \mu_t - \beta_2^2 \sigma_t^2 - \beta_3^2 \sigma_t \sqrt{f_t^2(\kappa^2)} = \\ \alpha \left([\beta_1^1 - \beta_1^2] \mu_t + [\beta_2^1 - \beta_2^2] \sigma_t^2 + \sigma_t [\beta_3^1 \sqrt{f_t^1(\kappa^1)} - \beta_3^2 \sqrt{f_t^2(\kappa^2)}] \right) + \varepsilon_t , \end{aligned} \quad (10)$$

where $f_t^i(\kappa^i)$, $i=1,2$, denotes the arguments of the square root in the various models. If we assume that the β s are equal to their theoretical values, then we obtain the following simplification:

$$\ln x_t - \mu_t + \sigma_t^2 - \sigma_t \sqrt{f_t^2(\kappa^2)} = \alpha \left(\sigma_t [\sqrt{f_t^1(\kappa^1)} - \sqrt{f_t^2(\kappa^2)}] \right) + \varepsilon_t . \quad (11)$$

We will perform the tests both with the β s set equal to their theoretical values and with estimated β s.

7. Empirical Results

Table 4 provides the results for the first test, showing the estimates of the β s for the best fitting $\kappa \equiv \frac{1+r}{1+\delta}$. These estimates of κ are smaller than one for all three models ($\kappa=0.76$ for Model 1; $\kappa=0.87$ for Model 2; and $\kappa=0.78$ for Model 3), but with fairly wide confidence intervals.¹² In view of the definition of κ , its estimates for Models 1,

¹² As noted, we have estimated κ by means of a grid search. The standard errors presented in the table are therefore conditional on the value of κ listed at the top of each column. A confidence interval for κ can be obtained in a straightforward way, by recognizing that a likelihood ratio test can be used to test any value of κ . If we apply this approach to model 2 (which as will be seen below is our preferred model), we find that $\kappa=1$ cannot be rejected at a 5% significance level.

2 and 3 imply high subjective discount rates. This is a fairly typical finding in the literature.¹³

For all three models the estimates for the coefficient β_1 of μ_t have the signs predicted by theory and are very close to the theoretical value of 1. The coefficient of σ_t^2 , which should be equal to -1 according to the theory, is positive for Model 1, and negative for Models 2 and 3. Similarly, the coefficients for the square root terms, which should be equal to 1, tend to deviate from their theoretical values towards zero, but are closest to 1 for Models 2 and 3. In summary, Models 2 and 3 produce estimates that are closer to the theoretical predictions than Model 1.

Table 4: Random Effect Interval Regressions

	Model 1 (4)	Model 2 (5)	Model 3 (6)
κ	0.76	0.87	0.78
μ	1.008** (0.002)	1.000** (0.001)	0.997** (0.002)
σ^2	0.089* (0.041)	-0.470** (0.038)	-0.306** (0.051)
$\sigma\sqrt{\cdot}$	0.108** (0.024)	0.508** (0.024)	0.573** (0.044)
Observations	5,871	9,293	5,723
Number of id	2,148	3,075	2,128
Log Likelihood	-12,438.22	-19,423.39	-11,985.06

Standard errors in parentheses. * significant at the 5% level; ** significant at the 1% level.

Table 5 provides the results for two variants of the second test using the best fitting values of κ . The first variant uses (10). That is, the three separate models are estimated first and produce estimates of the β s, after which α is estimated. The

¹³ See Frederick et al. (2002) for an overview.

second variant uses (11). That is, it estimates α while restricting the β s to their theoretical values.

In the columns headed “Model 1”, Model 1 represents $m_2(z_t)$ in (8), whereas $m_1(z_t)$ is represented by the other two models. For the comparison of Models 1 and 2, we find that the estimated values of α are closer to 1 than to 0 in both variants of the test. As a result, we reject Model 1 against Model 2. The comparison of Models 1 and 3 yields similar results, though the estimates of α are not as close to 1 as for the test of Model 1 versus Model 2. In any case, all results suggest that the empirical performance of Model 1 is inferior to that of the other models.

Table 5: Non-nested Tests of Models 1, 2 and 3 against each other

Test of	Against	Model 1 (4)		Model 2 (5)
		Model 2 (5)	Model 3 (6)	Model 3 (6)
(10)	α	1.127* (0.067)	0.866* (0.085)	0.298* (0.056)
(11)	α	1.031* (0.025)	0.742* (0.031)	0.480* (0.029)

$\alpha = 1$ means the alternative model (2nd row) is accepted; $\alpha = 0$ means the benchmark model (1st row) is accepted. Standard errors in parentheses.

* significant at the 1% level.

In the columns headed “Model 2”, Model 2 plays the role of $m_2(z_t)$ in (8). Here, both tests suggest a preference for Model 2. For the variant not imposing the theoretical values of the β s, the comparison of Model 2 against Model 3 yields an estimate of α well below 0.5, suggesting a preference for Model 2.¹⁴ For the test restricting the parameter values to be equal to the ones predicted by the theoretical model, the estimate of α suggests only a slight preference for Model 2.

¹⁴ Strictly speaking, only values of α equal to one or zero are conclusive evidence in favor of one model or the other; a value of α in between zero and one would suggest a mixture of two models outperforms both models.

Thus, as so often in non-nested tests the results are not entirely clear-cut. The empirical performance of Model 1 is clearly inferior to that of either Model 2 or Model 3, but a choice between Models 2 and 3 has to be more tentative. Yet Table 5 seems to suggest that Model 2 is somewhat better. The evidence in Table 4 is somewhat less conclusive, but also there on balance the parameter estimates for Model 2 are a little closer to their expected theoretical values than for Model 3.

The results raise the question why the parameter estimates deviate from their theoretical values. The obvious explanation is measurement error. It is known that the measurement of σ is not very precise (Kapteyn, 1977), and hence, both the terms σ^2 on the right hand side of equations (5) and (6) and the square roots, which contain σ , suffer from measurement error, which will tend to bias their associated coefficients towards zero.

Finally, we will compare the implications of our estimates with results found in the literature. An important parameter of interest in intertemporal models of consumption smoothing is the intertemporal elasticity of substitution (IES); see, e.g., Hall (1988) or Kapteyn and Teppa (2003). The IES for the lognormal utility function is equal to:

$$IES = \frac{\sigma_i^2}{\ln x_i - (\mu_i - \sigma_i^2)} .$$

We calculate the IES based on Models 2 and 3, using predicted consumption and adding two random draws from a normal distribution to account for the random effect and the estimation error term. Table 6 presents the distribution of the individual estimates of the IES (i.e., the median and the first and third quartiles of the distribution) by education level and income quartile. We find that the median IES is 0.201 for Model 2 and 0.203 for Model 3. The median is increasing in education and income. The results for Models 2 and 3 are very similar.

The values of our IES estimates are similar to what other authors using microeconomic data have found. For example, Hall (1988), using several different data sets from the US and different estimators, estimated values of the IES for aggregate consumption close to zero. Barsky et al. (1997) report average lower and upper bounds on the absolute value of the IES of 0.007 and 0.36. Yogo (2004) estimates the IES for

eleven developed countries using different specifications and finds that in all cases the IES is closer to zero than to one. His estimates of the IES for the Netherlands lie between -0.25 and 0.24.¹⁵ Kapteyn and Teppa (2007), using an approach similar to that of Barsky et al. (1997), find an IES of about 0.5. The fact that our estimates are in line with the earlier estimates is reassuring, since, after all, our approach is consistent with (but simpler to implement than) conventional Euler equation estimation.

Table 6: Distribution IES by Education and Income: Model 2

	P25		P50		P75	
	Model 2	Model 3	Model 2	Model 3	Model 2	Model 3
All	- 0.076	-0.045	0.201	0.203	0.522	0.516
Education						
Primary or less	-0.054	-0.042	0.213	0.211	0.523	0.490
Lower level	-0.110	-0.043	0.164	0.176	0.475	0.477
Interm. vocational	-0.079	-0.035	0.191	0.196	0.541	0.510
Interm. general	-0.056	-0.034	0.215	0.204	0.507	0.502
Higher vocational	0.014	-0.087	0.230	0.215	0.530	0.519
University	-0.064	-0.057	0.223	0.207	0.536	0.533
Other	-0.046	-0.439	0.222	0.190	0.508	0.418
Missing	-0.067	0.045	0.247	0.280	0.564	0.665
Income						
1. Quartile	-0.085	-0.029	0.184	0.202	0.496	0.506
2. Quartile	-0.089	-0.074	0.184	0.189	0.509	0.483
3. Quartile	-0.082	-0.037	0.192	0.198	0.500	0.491
4. Quartile	-0.006	-0.048	0.244	0.226	0.587	0.579

Income quartiles derived by wave.

Model 2: N=9,293. There are 530 missing observations in education, after imputing missing values from education in other waves. $\kappa = 0.83$.

Model 3: N=5,723. There are 312 missing observations in education, and only 32 with education “other”, after imputing missing values from education in other waves. $\kappa = 0.78$.

8. Concluding Remarks

Empirical Models of intertemporal allocation of consumption usually rely heavily on Euler equations. Estimating such models requires several strong assumptions. In this

¹⁵ Research in macroeconomics typically finds numbers close to zero; see Guvenen (2006) for an attempt to reconcile the different results.

paper, we aim at reducing some of these estimation and identification requirements by exploiting subjective data. Specifically, we investigate if directly measured utility functions and expectations help to explain behavior. It appears that a rather straightforward application of the estimation of Euler equations to individually measured welfare functions and expectations generates results that are in line with the theory. Parameter estimates have the correct sign, although some of the parameters seem to deviate substantially from their theoretical values. A plausible explanation is measurement error in some of the right hand side variables. This warrants further research. We find that estimates of the Intertemporal Elasticity of Substitution are of similar order of magnitude as found in the literature using very different approaches.

It seems clear that various improvements in the empirical implementation are possible. In particular, the measurement of expectations has been rather crude. Our empirical work was, furthermore, hampered by the fact that saving was measured in brackets, rather than continuously. Continuous measurement would have made the analysis simpler and more powerful.

Having said that, the current results appear sufficiently promising to further pursue this line of research. We also observe that the data seem to favor Model 2. It is important to note that for the estimation of Model 2, only cross-section data are needed. Both the measurement of expectations and of welfare functions is in principle quite straightforward and should therefore simplify empirical work that tries to improve our understanding of intertemporal decision-making in consumption.

References

- Attanasio, O. and H. Low, 2004. Estimating Euler Equations. *Review of Economic Dynamics* 7(2), 406-435.
- Barsky, R. B., M. S. Kimball, F. T. Juster, and M. D. Shapiro, 1997. Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study. *The Quarterly Journal of Economics* 112(2), 537-579.
- Battistin, E., R. Blundell, and A.D. Lewbel, 2007. Why is Consumption more Lognormal than Income? Gibrat's law revisited. Working Paper, Boston College.

- Browning, M. and A. Lusardi, 1996. Household Saving: Micro Theories and Micro Facts. *Journal of Economic Literature* XXXIV (December), 1797-1855.
- Carroll, C., 2001. Death to the Log-Linearized Consumption Euler equation! (And Very Poor Health to the Second-Order Approximation). *Advances in Macroeconomics* 1(1), Article 6.
- Clark, A. 1999. Are wages habit-forming? Evidence from micro data. *Journal of Economic Behavior and Organization*. 39(2), 179-200.
- Das, M., and B. Donkers, 1999. How Certain Are Dutch Households About Future Income? An Empirical Analysis. *The Review of Income and Wealth* 45(3), 325-338.
- Davidson, R and J. G. MacKinnon, 1981. Several Tests for Model Specification in the Presence of Alternative Hypotheses. *Econometrica* 49(3), 781-793.
- Davidson, R. and J. G. MacKinnon, 1993. *Estimation and Inference in Econometrics*. Oxford: Oxford University Press.
- Dominitz, J. and C. F. Manski, 1997. Using Expectations Data to Study Subjective Income Expectations. *Journal of the American Statistical Association* 92(439), 855-867.
- Frederick, S., G. Loewenstein, and T. O'Donoghue, 2002. Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature* XL (June), 351-401.
- Groot, W., H. Maassen van den Brink, and E. Plug, 2004. Money for Health: The Equivalent Variation of Cardiovascular Diseases. *Health Economics* 13(9), 859-872.
- Guvenen, F., 2006. Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective. *Journal of Monetary Economics* 53(7), 1451-1472.
- Hall, R. E., 1978. Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy* 86(6), 971-987.
- Hall, R. E., 1988. Intertemporal Substitution in Consumption. *Journal of Political Economy* 96(2), 339-357.

- Kapteyn, A., 1977 A Theory of Preference Formation, Ph.D. thesis, Leiden University.
<http://arno.uvt.nl/show.cgi?fid=76998>
- Kapteyn, A. and F. Teppa, 2003. Hypothetical Intertemporal Consumption Choices. *The Economic Journal* 113 (March), C140-C152.
- Kapteyn, A. and T. Wansbeek, 1985. The Individual Welfare Function. *Journal of Economic Psychology* 6(4), 333-363.
- Manski, C. F., 2004. Measuring Expectations. *Econometrica* 72(5), 1329-1376.
- Rablen, M., 2008. Relativity, Rank and the Utility of Income. *The Economic Journal* 118 (April), 801-821.
- Van de Stadt, H., A. Kapteyn, and S. van de Geer, 1985. The relativity of utility: evidence from panel data. *The Review of Economics and Statistics* 67(2), 179-187.
- Van Herwaarden, F. G., A. Kapteyn, and B. M. S. van Praag, 1977. Twelve Thousand Individual Welfare Functions: a Comparison of Six Samples in Belgium and The Netherlands. *European Economic Review* 9(3), 283-300.
- Van Praag, B. M. S., 1968. *Individual Welfare Functions and Consumer Behavior*. Amsterdam North-Holland Publishing Company, Contributions to Economic Analysis, Nr. 57.
- Van Praag, B. M. S., 1971. The Welfare Function of Income in Belgium: An Empirical Investigation. *European Economic Review* 2(3), 337-369.
- Van Praag, B. M. S. and A. Kapteyn, 1973. Further Evidence on the Individual Welfare Function of Income: an Empirical Investigation in The Netherlands. *European Economic Review* 4(1), 32-62.
- Van Praag, B. M. S. and A. Ferrer-i-Carbonell, 2004. *Happiness Quantified*. Oxford: Oxford University Press.
- Yogo, M., 2004. Estimating the Elasticity of Intertemporal Substitution When Instruments are Weak. *Review of Economics and Statistics* 86(3), 797-810.

Appendix A

This appendix explains the sample selection in more detail and shows some summary statistics for the working sample. Table A.1 shows the sample selection and the number of observations of the working sample. In what follows, we explain why some of the observations were dropped. In particular, one of the sample selection steps might cause some concern. A high number of observations (1,708) were dropped because of unusable answers to the income satisfaction questions. Of these, 701 had an implausibly low answer (< 11) to the question of which income the respondent would consider good or very good. The remaining observations (1,007) were dropped because the answers were missing or had coding errors without an obvious correction, or because the answers did not make sense given the wording of the questions (for example when respondents stated very low amounts which could be monthly amounts although the question asked about annual amounts, but a correction would have yielded amounts far outside actual incomes). To reduce the number of observations lost, we recoded obvious errors (such as the wrong number of zero digits). Of the final 9,293 observations used, 160 had some or all of the answers recoded. We also corrected the lowest and/ or highest expected income in 2 cases of seemingly missing zeros.

Table A.1: Sample Selection

Initial Observations (2001-2007) with income satisfaction questions answered	11,736
Answers to income satisfaction questions appearing to be topcoded	31
Unusable income satisfaction answers	1,708
Top-coded or missing income expectations	125
First question about savings (yes/ no) not answered	44
Erroneous income expectations	430
Highest expected income $< 5,000$ Euros	30
Missing or negative consumption bounds	75
Final Sample	9,293

Table A.2 shows the summary statistics for the estimated parameters as well as the lowest and highest expected income by year and for the entire sample.

Table A.2: Summary Statistics: Medians, by Year and for Total

Variable	2001	2002	2003	2004	2005	2006	2007	Total
Income	27,423	27,756	28,519	27,741	24,569	29,449	28,107	27,566
Utility Function Parameters								
μ_t	9.938	9.918	10.083	10.052	10.048	10.052	10.052	10.023
σ_t	0.410	0.406	0.436	0.409	0.374	0.396	0.383	0.403
Lowest expected net income	24,996	24,216	26,105	24,684	24,354	24,000	23,575	24,684
Highest expected net income	29,995	29,059	31,326	30,855	29,509	30,000	29,469	30,000
Subjective distribution in next year								
Log-median	10.2221	10.189	10.284	10.230	10.199	10.203	10.208	10.221
Log-standard deviation	0.036	0.039	0.042	0.040	0.039	0.039	0.042	0.039
Median	27,470	26,613	29,268	27,731	26,882	26,996	27,111	27,470
Standard deviation	12,588	11,946	15,016	12,984	11,338	12,054	12,012	12,509
N	1,206	1,227	1,299	1,435	1,420	1,350	1,356	9,293

All income measures are annual and expressed in Euros of 2006.

Appendix B

This appendix derives equation (4) in Section 5. For convenience of notation, subscripts are suppressed. We have

$$E_t U'(x_{t+1}, \mu_{t+1}, \sigma_{t+1}) = \frac{1}{2\pi} \int_0^\infty \frac{1}{x} \frac{1}{\sigma} \exp\left[-\frac{1}{2} \left(\frac{\ln x - \mu}{\sigma}\right)^2\right] \frac{1}{x} \frac{1}{s} \exp\left[-\frac{1}{2} \left(\frac{\ln x - m}{s}\right)^2\right] dx \quad (12)$$

Applying a transformation of variables $z = \ln x$, this can be written as

$$\frac{1}{2\pi} \int_{-\infty}^\infty \frac{1}{s} \frac{1}{\sigma} \exp\left[-z - \frac{1}{2} \left(\left(\frac{z - \mu}{\sigma}\right)^2 + \left(\frac{z - m}{s}\right)^2 \right)\right] dz. \quad (13)$$

It can be verified directly that this can be further rewritten as:

$$\begin{aligned} & \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{s^2 + \sigma^2}} \exp\left\{-\frac{1}{2} \left[\frac{(m - \mu)^2 - \sigma^2 s^2 + 2\mu s^2 + 2m\sigma^2}{s^2 + \sigma^2} \right]\right\} \\ & \frac{1}{\sqrt{2\pi}} \int_{-\infty}^\infty \frac{\sqrt{s^2 + \sigma^2}}{s\sigma} \exp\left\{-\frac{1}{2} \left[\frac{z - \frac{\mu s^2 + m\sigma^2 - \sigma^2 s^2}{s^2 + \sigma^2}}{\frac{\sigma s}{\sqrt{s^2 + \sigma^2}}} \right]^2\right\} dz. \end{aligned} \quad (14)$$

The second term is the integral of a normal density, and hence equal to one. So we obtain (reintroducing the subscripts):

$$E_t U'(x_{t+1}, \mu_{t+1}, \sigma_{t+1}) = \frac{1}{\sqrt{2\pi}} \frac{1}{\sqrt{s_{t+1}^2 + \sigma_{t+1}^2}} \exp\left\{-\frac{1}{2} \left[\frac{(m_{t+1} - \mu_{t+1})^2 - \sigma_{t+1}^2 s_{t+1}^2 + 2\mu_{t+1} s_{t+1}^2 + 2m_{t+1} \sigma_{t+1}^2}{s_{t+1}^2 + \sigma_{t+1}^2} \right]\right\} \quad (15)$$

It is of some interest to consider the case $s=0$, i.e. where there is no uncertainty about the future. In that case, the expression should reduce to the marginal utility of consumption in period $t+1$. Indeed we get

$$\begin{aligned} E_t U'(x_{t+1}, \mu_{t+1}, \sigma_{t+1}) &= \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma_{t+1}} \exp\left\{-\frac{1}{2} \left[\frac{(\ln x_{t+1} - \mu_{t+1})^2 + 2 \ln x_{t+1} \sigma_{t+1}^2}{\sigma_{t+1}^2} \right]\right\} = \\ & \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma_{t+1}} \frac{1}{x_{t+1}} \exp\left\{-\frac{1}{2} \left[\frac{(\ln x_{t+1} - \mu_{t+1})^2}{\sigma_{t+1}^2} \right]\right\} \end{aligned} \quad (16)$$

which is the lognormal density function.

Define $\kappa \equiv \frac{1+r}{1+\delta}$. The general solution for $\ln x$ that can now be derived from

(2) and (3) is given by

$$\ln x_t = \mu_t - \sigma_t^2 + \sigma_t \sqrt{(\sigma_t^2 - 2\mu_t) - (\sigma_{t+1}^2 - 2\mu_{t+1}) + \ln \frac{s_{t+1}^2 + \sigma_{t+1}^2}{\kappa^2 \sigma_t^2} + \frac{[m_{t+1} - (\mu_{t+1} - \sigma_{t+1}^2)]^2}{s_{t+1}^2 + \sigma_{t+1}^2}} . \quad (17)$$