Netspar Network for Studies on Pensions, Aging and Retirement

Alessandro Bucciol

Measuring Self-Control Problems: A Structural Estimation

Discussion Paper 01/2009 - 002
January 3, 2009

MEASURING SELF-CONTROL PROBLEMS: A STRUCTURAL ESTIMATION*

ALESSANDRO BUCCIOL†

University of Amsterdam and Netspar

This draft: January 3, 2009

Abstract

We perform a structural estimation of the preference parameters in a buffer-stock consumption model augmented with temptation disutility. We adopt a two-stage Method of Simulated Moments methodology to match our simulated moments with those observed in the US Survey of Consumer Finances. To identify the parameters we use liquid and quasi-liquid (retirement) wealth holdings at different ages as target moments. We find evidence of a significantly positive degree of temptation, and estimates under the unrestricted model are more precise than estimates under the restricted model with no temptation. Overall, a model with temptation preferences explains better the observed target moments.

JEL classification codes: D91; E21; G11.

Keywords: temptation disutility; method of simulated moments; non-liquid assets; consumption-saving behaviour; life-cycle models.

* I am grateful to Guglielmo Weber for his advice on this and earlier versions. Comments from Christopher Carroll, Alexander Michaelides, and Stephan Meier and two anonymous referees were also of significant benefit. I further thank Viola Angelini and the seminar participants at MIT, University of Padua, University of Venice, and IFS. Financial support from MIUR (PRIN 2005, grant no. 2005132539_001) is gratefully acknowledged. I am responsible for any error.

† Contact information: Alessandro Bucciol, University of Amsterdam, Dept. of General Economics, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands. Telephone: +31 20.525.4236. Fax: +31 20.525.4254. Email: a.bucciol@uva.nl.

1. Introduction

A recent branch of the economic literature has been investigating the impact of temptation on economic decisions. Following the seminal work of Gul and Pesendorfer (2001), a number of authors have been using the temptation disutility model in various contexts, most noticeably asset pricing (Krusell et al., 2002; DeJong and Ripoll, 2007) and addiction (Gul and Pesendorfer, 2007). The use of the model in household consumption-saving decisions is instead hindered by unclear evidence on temptation and self-control problems in life-cycle choices. Although experimental results seem to support the hypothesis of a tempting behaviour (see Ameriks et al., 2007; Brown et al., 2007), it is challenging to assess whether agents' preferences exhibit temptation from empirical consumption and saving data. The reason is that the predictions of the temptation disutility model are undistinguishable from those of the standard life-cycle model for a given choice of the preference parameters, unless further elements are taken into account.

The aim of this paper is to assess whether temptation exists in household life-cycle decisions, and evaluate its quantitative importance by focusing on real data on liquid and quasi-liquid asset holdings. An implication of the Gul and Pesendorfer's model is that tempted agents have a desire to commit their actions. Since resisting the temptation to make choices with immediate rewards is costly, they are better-off when they save in assets that cannot be withdrawn immediately or used as a collateral for borrowing; in so doing they eschew situations of exposure to more tempting alternatives. So far, the literature has explored the impact of temptation on one specific commitment asset: Social Security. It is a well-known result that this programme reduces the welfare of standard agents in dynamically efficient economies (see, e.g., Imrohoroglu et al., 1999). Simulations from overlapping-generation models augmented with temptation disutility, however, show that Social Security generates less severe welfare losses in tempted agents than in non-tempted agents (Kumru and Thanopoulos, 2008), and in some cases it may even create welfare gains (Bucciol, 2008). Social Security is not useful for our purpose as it forces to a mandatory saving, identical for tempted as well as non-tempted agents. We instead focus on other assets providing commitment. Relevant assets in the portfolios of US households are housing and retirement plans (see Laibson, 1997, and the reference therein). Since the choice to hold housing assets may also be driven by consumption motives, to isolate the pure effect of temptation we consider only retirement plans and similar quasi-liquid financial assets.

We use a "buffer-stock" model of life-cycle consumption and investment choices (Deaton, 1991; Carroll, 1992) that accounts for uncertainty on labour income and death age. The model is augmented with temptation disutility as in Gul and Pesendorfer (2004), and nests the standard CRRA utility as a special case. Investment can be made at any time between liquid and quasi-liquid assets; the quasi-liquid asset limits the withdrawal of funds before retirement. This significantly reduces the cost of self-control for tempted agents. We use a two-stage Method of Simulated Moments (henceforth MSM) procedure (Pakes and Pollard, 1989) to structurally estimate the three preference parameters – the degree of temptation, the discount factor, and the coefficient of relative risk aversion – that make the model match the holdings of liquid and quasi-liquid financial wealth observed in US household survey data. In the first stage of our MSM procedure we calibrate the model's exogenous parameters to describe the real US economy. In the second stage we use this information to estimate the preference parameters from the comparison between simulated and observed moments. Our procedure accounts for first-stage uncertainty; Gourinchas and Parker (2002) and Laibson et al. (2007) highlight that ignoring uncertainty on the calibration of the exogenous parameters significantly reduces the standard errors of the point estimates. We then compare the estimates of the (unrestricted) temptation disutility model with those of a standard (restricted) life-cycle model of buffer-stock saving, in which temptation is absent by definition. This way we are able to formally compare the two models, using *t*-tests and over-identification tests.

This paper follows the literature on the MSM estimation of the preference parameters in life-cycle models, including the prominent works of Gourinchas and Parker (2002), Cagetti (2003), and Laibson et al. (2007). Gourinchas and Parker (2002) estimate the discount factor and relative risk aversion coefficient of a standard life-cycle model from US Consumer and Expenditure Survey (CEX) consumption data. Cagetti (2003) runs a similar analysis using US Panel Study of Income Dynamics (PSID) and Survey of Consumer Finances (SCF) data on total wealth holdings. The closest work to ours is Laibson et al. (2007). The authors consider a quasi-hyperbolic discount model and estimate the short-run and long-run discount factors from US PSID and SCF data on consumption, wealth, and credit card borrowing. In most of their analysis the coefficient of risk aversion is set exogenously. They obtain statistically different short-run and long-run discount factors, and find that the model with quasi-hyperbolic discount functions explains the data better than the standard model.

In this paper we run a similar exercise, using as target variables liquid and quasi-liquid wealth holdings from the SCF, separately by age groups. In contrast to Laibson et al. (2007), we provide a joint estimate of the three preference parameters in our model (the degree of temptation, the discount factor, and the coefficient of risk aversion). Estimates in Laibson et al. (2007) show to be sensitive to the choice of the coefficient of risk aversion, and taking it as endogenous allows us to account for the potential interplay among these variables.

Using our benchmark calibration we obtain reasonable estimates for the preference parameters in the restricted model (discount factor: 0.9312, relative risk aversion: 1.6845), but the coefficient of risk aversion is estimated with extremely poor precision. This suggests that there is an identification problem when matching liquid and quasi-liquid wealth holdings with the restricted model. With the unrestricted model we estimate a degree of temptation of 0.1596, a discount factor of 0.9960, and a risk aversion of 0.8957. The degree of temptation is significantly different from 0, and the standard errors of all the parameters are small. In particular, risk aversion is significantly below 1. This large change of the parameter between restricted and unrestricted models supports our choice to treat it as endogenous in the analysis. Overall, the model augmented with temptation disutility provides a better fit to the observed US data and generates a MSM loss function that is more than three times smaller than the one generated under the restricted model. Our results do not change qualitatively if we vary the calibration of the exogenous parameters. Quantitatively, the estimates are affected by the parameters describing the features of the quasi-liquid asset.

The paper is organised as follows. Section 2 introduces the theoretical model. Section 3 reviews the two-step MSM methodology for estimating the parameters in the model. Section 4 presents the target variables derived from SCF data, the calibration of the exogenous parameters (first-stage estimation), and the MSM estimates of the parameters (second-stage estimation). A comparison of the behaviour predicted by the restricted and unrestricted models with the estimated preference parameters is then made. The Section ends with a robustness check around some critical parameters. Section 5 concludes, and the Appendix A provides mathematical details on the life-cycle model and the MSM.

2. The model

2.1. Setting

The economic framework consists of a life-cycle model where time is discrete and t denotes adult age (effective age minus 24). Every period corresponds to one year, and an agent lives for a maximum of $D = 66$ periods (age 90), K as a worker and $D - K$ as a retiree. For sake of simplicity the retirement age K is taken as exogenous. We allow for uncertainty in death age as in Hubbard et al. (1995). At any time t, individuals face a probability π_{t+1} of surviving until the next period. Death is certain after period D, that is, $\pi_{D+1} = 0$.

The timing of the model is the following. The agent starts any period t with a given amount of wealth, and receives an income Y_t. All the variables are expressed in real terms. Wealth is made of liquid and quasi-liquid assets. The rules governing the quasi-liquid asset are the following:

- Contributions can be made in any period before retirement age K;
- Withdrawals from the vested account balance H_t are unlimited after age K, and limited to a fraction $\delta < 1$ before K.

This asset can be seen as a stylised retirement plan. For instance, the US 401(k) plans impose severe restrictions on the withdrawal of funds before age 59½. Withdrawal is limited to few specific reasons (e.g., medical expenses, home purchase, college tuition), and usually incurs a severe penalty tax. Many 401(k) plans, however, easily allow one to take loans at a low interest rate, up to a given fraction of the plan vested account balance.

Cash-on-hand $X_t \geq 0$ is the sum of liquid wealth and income, and grows at a constant gross rate of return R^X; quasi-liquid wealth $H_t \geq 0$ instead grows at a constant gross rate of return R^H. Given X_t and H_t, the agent determines the optimal level of consumption C_t, and the investment I_t in the quasi-liquid asset. Given the quasi-liquid asset features, the following constraints to the investment choice must hold:

$$\begin{cases} I_t \geq -\delta H_t & t \leq K \\ I_t \in [-H_t, 0] & t > K \end{cases} \tag{1}$$

What is left (if any) after consumption and investment in the quasi-liquid asset goes into cash-on-hand. If the agent cannot die with debt, uncertainty in death age binds at any time the sum of consumption and investment in the quasi-liquid asset to be no larger than cash-on-hand, i.e.,

$$C_t + I_t \leq X_t \tag{2}$$

2.2. Preferences

The instantaneous payoff function follows the temptation disutility model of Gul and Pesendorfer (2001):

$$U\left(C_t, X_t, H_t\right) = u\left(C_t\right) - \left(v\left(\tilde{C}_t\right) - v\left(C_t\right)\right) \tag{3}$$

where $u(\cdot)$ and $v(\cdot)$ are two von Neumann-Morgenstern utility functions, C_t denotes consumption at time t of non-durable goods and $\tilde{C}_t = \arg\max_{C_t}\left\{v\left(C_t\right)\right\}$ is the most tempting consumption alternative, i.e., the largest consumption that is achievable in period t. Using equation (1), the most tempting alternative in period t is

$$\tilde{C}_t = \max_{I_t}\left\{X_t - I_t\right\} = \begin{cases} X_t + \delta H_t & t \leq K \\ X_t + H_t & t > K \end{cases} \tag{4}$$

Agent's intertemporal preferences at time t are described by the following expected utility function,

$$\begin{aligned} &U\left(C_t, X_t, H_t\right) \\ &+E_t\left[\sum_{s=t+1}^{D} \beta^{s-t} \frac{1}{\pi_s}\left(\prod_{j=t+1}^{s} \pi_j\right)\left(\pi_s U\left(C_s, X_s, H_s\right) + \left(1-\pi_s\right) B\left(X_s, H_s\right)\right)\right] \end{aligned} \tag{5}$$

with β discount factor, and $B(\cdot)$ utility from leaving a bequest.

In what follows we characterise $u(\cdot)$ with a Constant Relative Risk Aversion (CRRA) utility function,

$$u\left(C_t\right) = \frac{\left(C_t\right)^{1-\rho} - 1}{1-\rho} \tag{6}$$

where $\rho > 0$ is the coefficient of relative risk aversion. We also restrict our attention to a rescaled CRRA utility function for $v(\cdot)$,

$$v(C_t) = \lambda u(C_t) \tag{7}$$

where $\lambda \geq 0$, the *degree of absolute temptation*, measures the agent's sensitivity to the tempting alternative. In particular, preferences are of the standard CRRA type when $\lambda = 0$. The larger λ is, the more temptation to consume takes control of the individual's behaviour. Under the specification in (6) and (7), the payoff function (3) is consistent with balanced growth (DeJong and Ripoll, 2007) and makes tempted households risk averse in consumption but risk loving in cash-on-hand (since $-v(\tilde{X}_t)$ in (3) is convex). An implication is that variations in consumption make the household worse-off, whereas variations in cash-on-hand make the household better-off as they reduce the utility cost of resisting temptation.

Finally, we follow the standard literature and assume for sake of simplicity that the function $B(\cdot)$ is a rescaled CRRA utility function of bequeathable wealth:

$$B(X_t, H_t) = \alpha u(X_t + \kappa H_t) \tag{8}$$

where the parameter α dictates the degree of altruism, and the parameter κ indicates the fraction of quasi-liquid wealth that is bequeathed after a decease.

2.3. Labour income

Following Zeldes (1989), labour income Y_t at any time t before retirement is exogenously described by a combination of deterministic and random components:

$$Y_t = P_t Z_t, \qquad t = 1, \ldots, K \tag{9}$$

$$P_t = G_t P_{t-1} N_t, \qquad P_0 = 1 \tag{10}$$

with G_t deterministic function of age and other individual characteristics, P_t permanent component with innovation N_t, and Z_t transitory component. We assume that the permanent shock $n_t = \log(N_t)$ and the transitory shock $z_t = \log(Z_t)$ are independent and identically normally distributed with parameters $\{-0.5\sigma_n^2, \sigma_n^2\}$ and $\{-0.5\sigma_z^2, \sigma_z^2\}$ respectively; their mean is different from zero to allow N_t and Z_t to be equal to one on average. Earnings in retirement are a constant fraction a of permanent labour income in the last working year:

$$Y_t = aP_t = aP_K\,, \qquad t = K+1,\ldots,D \tag{11}$$

2.4. The optimization problem

We represent the agent's problem using the recursive value function

$$V_t\left(X_t, H_t, P_t\right) = \max_{C_t, I_t} \begin{Bmatrix} U\left(C_t, X_t, H_t\right) \\ +\pi_{t+1}\beta E_t\left[V_{t+1}\left(X_{t+1}, H_{t+1}, P_{t+1}\right)\right] \\ +\left(1-\pi_{t+1}\right)\beta E_t\left[B\left(X_{t+1}, H_{t+1}, P_{t+1}\right)\right] \end{Bmatrix} \tag{12}$$

subject to the dynamic budget constraint for cash-on-hand,

$$X_{t+1} \le R^X\left(X_t - C_t - I_t\right) + Y_{t+1}\,, \qquad X_{t+1} \ge 0 \tag{13}$$

and for quasi-liquid wealth,

$$H_{t+1} \le R^H\left(H_t + I_t\right), \qquad H_{t+1} \ge 0 \tag{14}$$

with $X_0 = H_0 = 0$, and constraints (1)-(4) and (6)-(11). The control variables of the problem are $\left\{C_t, I_t\right\}_{t=1}^{D}$, and the state variables are $\left\{t, X_t, H_t, P_t\right\}_{t=1}^{D}$. Following Carroll (1992) we standardise the variables dividing by permanent income P_t, denoting the new variables by lower-case letters (e.g., $c_t = C_t / P_t$). This allows us to reduce the dimensionality of the state space (see Section A.1 for details). Since the problem does not admit a closed-form solution, the policy functions $c_t = c\left(x_t, h_t, t\right)$ and $i_t = i\left(x_t, h_t, t\right)$ are derived numerically using backward induction.

At any age $t < D$, disregarding corner solutions the optimizing behaviour is implicit in the system of two equations (see Section A.2)

$$\begin{cases} \dfrac{\partial U}{\partial c_t} = \pi_{t+1}E_t\left[\beta_{t+1}R^X\left(\dfrac{\partial U}{\partial c_{t+1}} + \dfrac{\partial U}{\partial x_{t+1}}\right)\right] + \left(1-\pi_{t+1}\right)E_t\left[\beta_{t+1}R^X\dfrac{\partial B}{\partial x_{t+1}}\right] \\ \dfrac{\partial U}{\partial c_t} = \pi_{t+1}E_t\left[\beta_{t+1}R^H\left(\dfrac{\partial U}{\partial c_{t+1}} + \dfrac{\partial U}{\partial h_{t+1}}\right)\right] + \left(1-\pi_{t+1}\right)E_t\left[\beta_{t+1}R^H\dfrac{\partial B}{\partial h_{t+1}}\right] \end{cases} \tag{15}$$

where $\beta_{t+1} = \beta\left(P_{t+1}/P_t\right)^{-\rho}$. Using the definition of $U(.)$ and $B(.)$ in equations (3), (6)-(8), the system in (15) can be rewritten as

$$\begin{cases} (c_t)^{-\rho} = \pi_{t+1} E_t \left[\beta_{t+1} R^X \left((c_{t+1})^{-\rho} - \tau (\tilde{c}_{t+1})^{-\rho} \right) \right] \\ \qquad + (1-\pi_{t+1}) E_t \left[\beta_{t+1} R^X \frac{\alpha}{1+\lambda} (x_{t+1} + \kappa h_{t+1})^{-\rho} \right] \\ (c_t)^{-\rho} = \pi_{t+1} E_t \left[\beta_{t+1} R^H \left((c_{t+1})^{-\rho} - \left(1-(1-\delta)\mathbf{1}_{\{\mathbf{t<k}\}}\right) \tau (\tilde{c}_{t+1})^{-\rho} \right) \right] \\ \qquad + (1-\pi_{t+1}) E_t \left[\beta_{t+1} R^H \frac{\alpha}{1+\lambda} \kappa (x_{t+1} + \kappa h_{t+1})^{-\rho} \right] \end{cases} \tag{16}$$

The parameter $\tau = \lambda/(1+\lambda) \in [0,1)$, the *degree of relative temptation*, measures the importance of temptation relative to consumption. The two equations in (16) differ from the standard case in the absence of temptation (obtained when $\lambda = \tau = 0$). The first equation (the Euler equation) informs that the marginal cost of giving up one unit of current consumption must be equal to the marginal benefit of consuming the proceeds of the extra saving in the next period, minus the marginal cost of resisting the additional temptation in the next period. Hence the cost of saving is higher for tempted individuals than for non-tempted ones, everything else being equal. The second equation arises from the decision to invest in the quasi-liquid asset. Here the negative effect of temptation is mitigated before retirement, when $\delta < 1$ and the asset is less liquid. As an extreme case, if the asset is completely illiquid ($\delta = 0$), the marginal cost of resisting the additional temptation in the next period does not enter the equation.

3. Estimation methodology

Our model generates a closed-form expression for the agent's first-order conditions. In principle, one could use the log-linear approximation of the first (Euler) equation in (16) to estimate the coefficients on risk aversion ρ and temptation τ directly with instrumental variables or generalised method of moments. Assuming that the conditional survival probability $\pi_{t+1} = 1$ – a reasonable approximation for most individual ages – and that the market return R^X varies over time, $R_t^X \neq R_s^X$ for $t \neq s$, the two parameters can be inferred from the equation (see Huang et al., 2005):

$$\ln\left(\frac{c_{t+1}}{c_t}\right) = \theta_0 + \theta_1 \ln R_{t+1}^X + \theta_2 \ln \frac{\tilde{c}_{t+1}}{c_{t+1}} + \varepsilon_{t+1}$$

where ε_{t+1} is the error,

$$\theta_1 = \frac{1}{\rho} \qquad \theta_2 = \frac{\lambda}{(1+\lambda)\chi^{\rho} - \lambda}$$

and $\chi = \tilde{c}/c$ is the steady-state ratio of temptation (wealth) to consumption. Consistent estimates are obtained provided that *i)* the Euler equation in (16) holds true, *ii)* long time series of individual consumption and wealth data are available, and *iii)* the error ε_{t+1} is not correlated with the instruments[1]. The three assumptions, however, may not hold true. First, the Euler equation can be incorrect, possibly because the choice of the utility functions (6) and (7) is wrong. Second, long time series of individual data may not be available for all the variables of interest. Third, in some cases the error may correlate with the instruments. Simulation exercises in Attanasio and Low (2004) show that the error is uncorrelated with common instruments unless many impatient agents are close to consuming all their cash-on-hand most of the time. In this case consumption in any given period is sensitive to income shocks and very volatile; as a consequence it is difficult to disentangle the contribution of each preference parameter in the variance of consumption growth. This issue is particularly relevant for the temptation disutility model, which induces agents to rationally procrastinate savings and consume at young ages most of their income.

Given these potential concerns we prefer to resort on a two-stage MSM procedure; this method allows us to estimate jointly the three household preference parameters (τ, β, ρ) by matching empirical and simulated moments. Simulations require a preliminary calibration of the model's exogenous parameters. The approach we follow incorporates uncertainty in the setting of these parameters as well as uncertainty in the observation data. While MSM can generate efficiency losses[2], it is more robust than alternative methods as it does not depend on the specific functional form of the Euler equation, nor does it require the use of time series data and instrumental variables.

In the first stage of our MSM procedure we set the exogenous parameters of the model, such as the rates of return to liquid and quasi-liquid assets, the income profile, and the survival probabilities. We also measure the uncertainty associated with this calibration. In the second stage we simulate life-cycle choices from the model of Section 2, conditional on the exogenous and endogenous (preference) parameters. From the simulated wealth profiles we

[1] Instruments are needed to deal with potential endogeneity and measurement error problems. Common instruments are lagged values of interest rate, income and consumption growth.
[2] It uses less information on the heterogeneity of the observed data and their underlying process.

calculate the moments used in the MSM estimation procedure. The procedure, described in detail below, determines the preference parameters that minimise the distance between simulated and empirical patterns of liquid and quasi-liquid wealth accumulation. Uncertainty in first-stage parameters propagates to the standard errors of the preference parameters estimated in the second stage.

3.1. Method of simulated moments

Our structural estimation is based on the Method of Simulated Moments developed in Pakes and Pollard (1989). In the model there are nuisance parameters χ that we treat as exogenous (e.g., the interest rate on liquid and quasi-liquid assets, the variance of the income risk) and other J parameters θ that we consider endogenous (the three preference parameters τ, β and ρ). In this paper we make assumptions on the nuisance parameters, and use these assumptions to estimate θ following a two-stage procedure similar to Gourinchas and Parker (2002) and Laibson et al. (2007).

In the first stage, we calibrate the nuisance parameters χ with $\hat{\chi}$, to reflect characteristics of the US economy (see Section 4.2). The estimates $\hat{\chi}$ are associated with the variance $\hat{\Omega}_{\chi} = \text{var}(\hat{\chi})$[3]. The second stage applies the MSM to estimate the unknown vector θ, conditional on $\hat{\chi}$ and $\hat{\Omega}_{\chi}$. Let m_i be a $L \times 1$ vector containing the observations of the L target variables over the $i-$eth sample unit, with $i = 1,...,N$[4]. Define the simulation analogue to $m_i, i = 1,...,N$ by $m(\theta, \chi)$, constructed as an average over M simulations. The moment condition implies that in expectation $E\left[g_i(\theta_0, \chi_0)\right] = E\left[m(\theta_0, \chi_0) - m_i\right] = 0$ where (θ_0, χ_0) is the true parameter vector. Let W be a positive definite $L \times L$ weighting matrix. Define

$$q(\theta, \chi) = \left(\frac{1}{N}\sum_{i=1}^{N} g_i(\theta, \chi)\right)' W \left(\frac{1}{N}\sum_{i=1}^{N} g_i(\theta, \chi)\right) \tag{17}$$

[3] As regards the nuisance parameters for which we do not have information on the variance of the estimates, we perform additional robustness check in Section 4.5.

[4] We assume here the same number of observations for each target variable. In practice we have a different number of observations for each moment. In the application we conservatively use the maximum number of observations; this inflates the standard errors of our estimates.

as the (scalar) weighted sum of squared deviations of simulated moments from their corresponding empirical values. The MSM chooses θ to minimise the loss function $q(\theta, \hat{\chi})$[5],

$$\hat{\theta} = \arg\min_{\theta} q(\theta, \hat{\chi}) \tag{18}$$

Pakes and Pollard (1989) show that, under some regularity conditions satisfied here, $\hat{\theta}$ is a consistent estimator of θ_0 and is asymptotically normally distributed with asymptotic variance

$$\Omega_\theta = \left(G_\theta' W G_\theta\right)^{-1} G_\theta' W \left(\Omega_g \left(1 + \frac{N}{M}\right) + G_\chi \Omega_\chi G_\chi'\right) W G_\theta \left(G_\theta' W G_\theta\right)^{-1} \tag{19}$$

where

$$G_\theta = \frac{\partial g_i(\theta_0, \chi_0)}{\partial \theta} \qquad G_\chi = \frac{\partial g_i(\theta_0, \chi_0)}{\partial \chi} \tag{20}$$

$$\Omega_g = E\left[g_i(\theta_0, \chi_0) g_i(\theta_0, \chi_0)'\right] \qquad \Omega_\chi = \mathrm{var}(\chi) \tag{21}$$

Equation (19) incorporates in $1 + N/M$ the correction for simulation error, and in $G_\chi \Omega_\chi G_\chi'$ the correction for "first-stage uncertainty". Failing to account for these two sources of uncertainty significantly biases the standard errors (see Gourinchas and Parker, 2002; Laibson et al., 2007).

We use equation (19) to calculate standard errors for the estimates of θ. All the matrixes are estimated consistently from sample data (the variances Ω_g and Ω_χ) or with numerical analogues (the derivatives G_θ and G_χ)[6]. We use two alternative weighting matrixes in our estimation. First, we choose a matrix that does not depend on the fitted model:

$$W_0 = \left(\Omega_g\right)^{-1} \tag{22}$$

This matrix (henceforth "robust") is inherited from the GMM literature, and is optimal when the first-stage correction does not matter.

Second, we employ the "optimal" weighting matrix

$$W_1 = \left(\Omega_g \left(1 + \frac{N}{M}\right) + G_\chi \Omega_\chi G_\chi'\right)^{-1} \tag{23}$$

[5] We perform the minimization with Matlab's Nelder-Mead simplex algorithm. This algorithm is slower but more robust than derivative-based methods.

[6] We take numerical derivatives on both sides of the optimum.

This matrix is efficient since it reduces the variance of the estimates to $\left(G_\theta' W_1 G_\theta\right)^{-1}$. We estimate it using the consistent parameters obtained with the robust weighting matrix (22). A growing literature questions the small-sample validity of optimal weighting. The optimal weighting matrix can indeed be more efficient than the robust weighting matrix, but also more biased (see for instance Altonji and Segal, 1996). For this reason we report the estimates obtained using both matrixes (22) and (23).

If the model is correct, we are able to perform tests on over-identifying restrictions using the statistic

$$\xi\left(\hat{\theta},\hat{\chi}\right) = N\left(\frac{1}{N}\sum_{i=1}^{N} g_i\left(\hat{\theta},\hat{\chi}\right)\right)' \hat{W}\left(\frac{1}{N}\sum_{i=1}^{N} g_i\left(\hat{\theta},\hat{\chi}\right)\right) \xrightarrow{d} \chi^2_{L-J} \qquad (24)$$

4. Structural estimation

4.1. Target moments

We compare the predictions of the temptation disutility model with data on wealth drawn from the US Survey of Consumer Finances (SCF). The SCF collects detailed information on assets and liabilities for a cross-section of US households. The data we use are exceptionally good because the survey *i)* uses a sample designed specifically to support wealth estimation, and *ii)* handles missing data computing five complete data sets of imputed values.

We focus on the waves from 1995 to 2004 (four waves) to disentangle age, year and cohort effects in the construction of each target variable. From the waves we collect data on socio-demographic characteristics, normal income[7] and financial assets. We neglect non-financial assets such as housing since their accumulation is also driven by reasons that we do not model explicitly.

In principle, we could use a variety of different target moments, provided that some include information on quasi-liquid wealth (otherwise the temptation model cannot be identified). To keep the analysis simple, we restrict our attention to the following four indicators:

- the ratio between liquid wealth and permanent income, average between ages 41 and 45;

[7] The survey defines normal income as an income that is not «unusually high or low compared to what [a household] would expect in a "normal" year». This information is not available in waves before 1995.

- the ratio between quasi-liquid wealth and permanent income, average between ages 41 and 45;
- the ratio between liquid wealth and permanent income, average between ages 61 and 65;
- the ratio between quasi-liquid wealth and permanent income, average between ages 61 and 65.

In so doing we capture the trend in the age profiles of liquid and quasi-liquid wealth. We ignore younger and older ages since we do not model explicitly educational choices and health shocks. If our model is correctly specified, however, analysing different moments should not change our estimated parameter values.

The financial assets we consider are transaction accounts (checking and saving accounts, money market accounts and mutual funds, and accounts to brokerages), saving bonds, certificates of deposit, mutual funds, bonds, stock, retirement plans (IRA, 401(k), thrift and pension plans), life insurances (in cash value) and other assets (including, e.g., trusts, annuities, loans from the household). In our definition only retirement plans form quasi-liquid wealth; all the remaining assets form liquid wealth.[8] We compute the amount of liquid and quasi-liquid wealth for each household, and divide these measures by normal income, interpreting normal income as being permanent income. We regress with OLS each ratio over a third-degree polynomial in age and other socio-demographic characteristics of the household head (marital status, number of children, race, education), plus cohort (five-year wide dummies) and year (log of real per capita GDP) variables. Regressions are corrected for imputations and sampling weights, and include only households whose head is working and in adult age (25-65). We use these coefficients to subtract from our ratios the cohort and year effects, using as reference the cohort of those born in 1961-1965 and surveyed in 1995. From these corrected ratios we construct sample averages for households with head in age 41-45 and 61-65. In so doing we take into account all the five imputations and use the SCF population weights. Disregarding population weights would bias the moments as the SCF over-samples wealthy households.

Our final sample consists of 1,763 households in age 41-45, and 1,455 in age 61-65. Table 1 reports the target moments in our sample, together with their standard errors. The sum

[8] In a separate analysis we considered a second definition, where quasi-liquid wealth also includes life insurances and other non-liquid assets. Our results remain unchanged with this definition.

of liquid and quasi-liquid wealth in ages 61-65 (1.5399 and 1.1884 respectively) roughly corresponds to the value (2.60) estimated in Laibson et al. (2007) from the waves 1983-1998 of the SCF. Using our definition, households hold on average more liquid than quasi-liquid wealth, and wealth relative to income in ages 61-65 is around three times higher than in ages 41-45. We observe larger standard errors in the estimates of the liquid wealth holdings, especially over the ages 61-65. This larger variability arises because some of the assets we treat as liquid (in particular trusts and annuities) are held by few households in our sample, and often account for a large fraction of the wealth of these households.

We use the standard errors to estimate the moments' covariance matrix Ω_g. The matrix is block-diagonal as we do not observe households over years, and moments computed from households in ages 41-45 and 61-65 are uncorrelated. To estimate the matrix we also use the observed correlations between same-age liquid and quasi-liquid wealth moments (0.136 in ages 41-45 and 0.231 in ages 61-65).

4.2. First-stage parameters

We explored the literature to find reasonable values and capture the main economic, demographic and political aspects of the US economy. Table 2 reports the exogenous parameters of the model, and the calibration we adopt (estimates and standard errors).

We take as survival probabilities $\{\pi_t, t = 2,\ldots,D+1\}$ those resulting from 1960 cohort life tables as computed by the Social Security Administration. We associate a different survival probability to each age t, and set $\pi_{D+1} = 0$. The number of working years is set to $K = 39$ (retirement age 63), with standard deviation 0.042, as results from our calculations based on the SCF 2004 question on the year of retirement.

The deterministic income growth consists of three parameters, (ϕ_1, ϕ_2, ϕ_3) taken from Scholz et al. (2006). The coefficients multiply respectively a constant, age, and age squared over 100. Scholz et al. (2006) derive the parameters from HRS data, separately for different groups. We use the estimates from the largest group in their dataset, the one including households with one income earner, and whose head is married and without college education. These estimates are similar to those in Engen et al. (1999) based on PSID data.

The pension replacement rate is set to $a = 41.3$ percent as in the intermediate figure from the 2008 OASDI Trustees Report (Social Security Administration, 2008). The report pro-

jects rates between 27.3 and 55.7 percent, depending on the scale of earnings patterns. Similar estimates, with a median value around 37 percent, are found in Engel et al. (1999) and Scholz et al. (2006) from HRS data. The standard deviations of permanent and transitory shocks are set to $\sigma_n = 0.147$ and $\sigma_z = 0.210$ respectively, from Carroll and Samwick (1997) using PSID data. Similar values are used in Gourinchas and Parker (2002).

The real return to liquid assets is set to $R^X = 1.034$, following Gourinchas and Parker (2002). The authors compute it as the average real return on Moody's AAA municipal bonds over the period from January 1980 to March 1993. It is more difficult to obtain an estimate of the historical return to quasi-liquid assets. A large empirical literature documents that quasi-liquid assets provide, on average, significantly higher returns than liquid assets (see Brennan and Subrahmanyam, 1996; Pàstor and Stambaugh, 2003).[9] It is however complicated to pin down a single number as indicative of quasi-liquid assets' real return; the return is intimately related to the asset's size and composition. Larger plans generate higher returns than smaller plans, as they generally have access to a wider variety of investment options and economies of scale. Munnell et al. (2006) report for 401(k) plans of intermediate size a median nominal return of 7.8 percent per year over the period 1988-2004. The real rate of return, obtained after discounting for inflation, is around 4.8 percent. Similar estimates with a nominal return around 8 percent are reported in Watson Wyatt (2008) over the period 1995-2006. In our benchmark analysis we set $R^H = 1.05$, thus 1.6 percentage points above the return to liquid assets.

In the real world many 401(k) plans allow employees to take loans from their retirement plans. In general these loans charge a small interest rate and are allowed for a maximum of 50 percent of the plan vested account balance. We therefore set to $\delta = 0.5$ the maximum fraction of quasi-liquid wealth that is available for withdrawal before retirement. We also set to $\kappa = 1$ the fraction of quasi-liquid wealth that is passed to the beneficiaries after death, according to the standard rules of retirement assets. Finally, the degree of altruism is set to $\alpha = 0$ following Hurd (1989). This implies that all bequests are accidental, generated by an uncertain lifespan.

Overall, we have seven exogenous parameters with an associated standard error (see Table 2). The information on the standard errors is used to compute the MSM standard errors.

[9] Possible reasons for higher returns in retirement plans are a favourable tax treatment, and access to asset classes otherwise not available to a small investor.

Since the estimates in Table 2 arise from different sources, we treat them as independent and therefore construct the matrix Ω_χ as a diagonal matrix whose elements are the squared values of the standard errors in Table 2, times the number of observations in the sample. We also have six further (groups of) parameters with no associated standard error. Such parameters (the rate of return to quasi-liquid assets, the fraction of quasi-liquid wealth that can be withdrawn, the fraction that can be bequeathed, the degree of altruism, the replacement rate and the survival probabilities) are difficult to elicit empirically. In Section 4.5 we check the robustness of our results to changes in these parameters.

4.3. Second-stage parameters

Table 3 shows the output of our benchmark MSM estimates for the CRRA (restricted) and temptation (unrestricted) models; the number of simulations in our analysis is set to $M = 2{,}000$. The table reports results using the robust and optimal weighting matrixes of equations (22) and (23). Since estimates vary only marginally with either choice, in our discussion we will focus only on estimates obtained with the optimal weighting matrix.

With the restricted model we find a discount factor of 0.9312 and a risk aversion of 1.6845. These estimates are in line with previous structural estimations[10], but exhibit large standard errors (especially the coefficient of risk aversion). Large standard errors are indicative of the model's failure to identify the preference parameters from the given target variables. Indeed, we see that the model is unable to match the target moments in ages 61-65, with the simulated liquid wealth being 31 percent larger than the empirical one ($2.0204/1.5399 - 1$) and the simulated quasi-liquid wealth being 93 percent larger ($2.2893/1.1884 - 1$). The unrestricted model generates values closer to the empirical data, although it does not match the targets perfectly: the simulated liquid (quasi-liquid) wealth in ages 61-65 is indeed 21 (24) percent larger that its empirical counterpart. With this model we however obtain much smaller standard errors for all the coefficients; this suggests that identification of the preference parameters can be achieved from target liquid and quasi-liquid wealth holdings once the unrestricted model is taken into account. We find a degree of relative temptation of 0.1596 – that is, nearly 16 percent. The estimate is signifi-

[10] Estimates of the discount factor in comparable cases range from 0.917 (Cagetti, 2003, from SCF mean data on high school graduates ignoring real wealth) to 0.957 (Gourinchas and Parker, 2002, benchmark analysis), and estimates of risk aversion from 1.397 (Gourinchas and Parker, 2002) to 4.45 (Cagetti, 2003).

cantly different from zero ($0.1596/0.0077 = 20.798 > 1.96$), and between the values found in DeJong and Ripoll (2007) and Huang et al. (2005).[11] We also find a discount factor of 0.9960 and a risk aversion of 0.8957. Such parameters change markedly after the inclusion of the degree of temptation in the model; in fact, with the new estimates the discount factor is not significantly different from 1, and the coefficient of risk aversion is significantly below 1. This evidence supports our choice to estimate the preference parameters jointly.

The minimised MSM loss function is trivially lower with the unrestricted model, which includes a further free parameter. The difference in magnitude is however remarkable: the loss function is 0.1866 with the unrestricted model, which is only 31 percent of the loss function with the restricted model (0.5952). The test of over-identifying restrictions however rejects the null hypothesis at usual significance levels, for both the restricted and unrestricted model. This suggests that the temptation model may not be enough to explain the observed wealth holdings, although it provides a closer match than the CRRA model.

4.4. Illustrative life-cycle simulation

To get an intuition of the real gain from using the model with temptation disutility rather than the model with CRRA utility, we simulate the life-cycle behaviour predicted in the two cases with the preference parameters estimated using the optimal weighting matrix (last two columns of Table 3). Results are based on 2,000 simulations over random realizations of labour income shocks and death age.

Figure 1 shows the average life-cycle consumption profile (standardised by permanent income) predicted with the two models over the working life. Both models generate a smooth, hump-shaped profile, where consumption reaches its top at around age 50. Consumption then falls to accommodate the reduction in income after retirement. Figure 2 draws the average liquid wealth holdings over the working life. Liquid wealth increases with age, reaching a peak at around age 60; it then reduces to support consumption at retirement. The predicted age-wealth profile is almost identical in the two models, and the

[11] DeJong and Ripoll (2007) argue that a degree of relative temptation of 0.073 $\left(0.0787/\left(1+0.0787\right)\right)$ is necessary to reconcile the theory with data on stock-price volatility, risk-free rates and equity premiums. Huang et al. (2005) obtain statistically significant GMM estimates of risk aversion and temptation from a log-linearised Euler equation. In their benchmark analysis their estimates are respectively $\rho = 0.6734$ $\left(1/1.485\right)$ and $\tau = 0.2058$ $\left(9.31^{0.6734}/\left(1+1/0.048\right)\right)$.

simulated values match adequately the empirical wealth holdings at ages 41-45 and 61-65. The main difference is that tempted agents accumulate less wealth at young ages. They indeed prefer to postpone the accumulation of savings to reduce their exposure to tempting consumption alternatives.

We observe different predictions of the models with CRRA utility and temptation disutility only in Figure 3. The figure shows the average holdings of quasi-liquid wealth over the working life. In both models quasi-liquid wealth starts accumulating after age 30, increases up to around age 60, and then reduces. The size of the holding is however different in the two models. Were the discount factor and risk aversion coefficients identical in the two models, we would observe fewer portfolio holdings under the CRRA model. Yet the two coefficients are not the same in the two models, and the CRRA model predicts larger holdings than the temptation model. Neither model matches exactly the target variables: both under-estimate the target at ages 41-45, and both over-estimate the target at ages 61-65. In this case the CRRA model generates quasi-liquid wealth holdings that are nearly twice as large as the empirical data; predicted holdings from the temptation disutility model are instead 24 percent larger.

In summary, a model with temptation disutility generates consumption and liquid wealth profiles similar to standard models with CRRA utility; predictions from both models adequately match empirical liquid wealth holdings, but fail to match quasi-liquid wealth holdings. The model with temptation disutility, nevertheless, generates a size of quasi-liquid wealth that is closer to the values found empirically in ages 61-65.

4.5. Sensitivity analysis

We conclude the section with a robustness check of our results to different calibrations of the exogenous parameters for which standard errors are not available. In what follows we report only the estimates in the temptation disutility model resulting with the optimal weighting matrix. Estimates in the restricted CRRA model always provide a poorer fit to the data, and are associated to large standard errors.

We first concentrate our robustness check on the parameters describing the features of the quasi-liquid asset. The concern is that our conclusions may change if the main characteristics of the quasi-liquid asset (its return and the fraction of wealth that can be withdrawn before retirement) in the model are closer to those of the liquid asset. In a similar exercise

Laibson et al. (2007) observe their estimates change markedly when they modify the asset returns. In particular they find no support for their (quasi-hyperbolic discounting) model when returns on liquid and quasi-liquid assets are closer. We run a similar analysis and consider two cases in which the return to the quasi-liquid asset is reduced to $R^H = 1.04$ and raised to $R^H = 1.06$ rather than the benchmark $R^H = 1.05$ (see Table 4). The quasi-liquid asset thus provides a return of $R^H - R^X = 0.6$ or 2.6 percent in excess to the liquid asset. When the return R^H increases, the quasi-liquid asset is more attractive for everybody. In our model this implies that a lower degree of temptation is enough to justify the investment and match the target variables. Our point estimate for τ is indeed 0.1075 when $R^H = 1.06$ and 0.2199 when $R^H = 1.04$. The degree of temptation is however significantly different from 0 in both cases. The withdrawal limit before retirement δ is also a crucial parameter in our framework, as a retirement asset with a higher limit is closer to being liquid. Everything else being equal, an agent in the model should exhibit a higher degree of temptation to generate the target wealth holdings, since the quasi-liquid asset offers less protection against self-control problems. In the last two columns of Table 4 we report the estimation results for the cases $\delta = 0.25$ and $\delta = 0.75$ (whereas the benchmark analysis assumes $\delta = 0.50$). We indeed see that, as the withdrawal limit increases, the estimate of the degree of temptation moves from 0.0825 ($\delta = 0.25$) to 0.3063 ($\delta = 0.75$); in both cases the coefficient is significantly positive. We however note that the estimates of all the preference parameters become less precise as δ increases; when the quasi-liquid asset offers only weak commitment, an identification problem of the model with temptation disutility may arise.

We then focus on a second set of exogenous parameters affecting the propensity to save. The first three columns of Table 5 show the output of the MSM estimation when we vary the pension replacement rate ($a = 0.27$ and 0.56 rather than 0.413) and the survival probabilities (1980 rather than 1960 cohort life tables). With a higher replacement rate, agents have a lower incentive to save for precautionary reasons. In contrast, the higher probability of being alive until elder ages makes agents save more to support their retirement consumption. Changes in these variables are likely to affect estimates of the discount factor β and the risk aversion ρ, but should affect the degree of temptation only marginally, through its interplay with β and ρ. We indeed see in Table 5 that the parameter varies between 0.1580 and 0.1716, and is significantly positive in all the cases.

We conclude our analysis with an estimate of the coefficients when the parameters governing the bequest utility vary. There is no consensus in the literature on the size of the degree of altruism α. We choose a parameter $\alpha = 1.5$ to reflect a marginal propensity to consume in the last period (when death is certain) roughly equal to 50 percent[12]. The last two columns of Table 5 report the MSM estimates when $\alpha = 1.5$, assuming that all the quasi-liquid wealth is bequeathed ($\kappa = 1$, as in the benchmark case) or only half of it is bequeathed ($\kappa = 0.50$). Agents with a bequest motive receive a utility gain from accumulating wealth, as they may leave it to their heirs once they pass away; this partially offsets the utility cost of exerting self-control. From the table we however see that the effect on the estimates is negligible, as the degree of temptation remains significantly positive and around the benchmark estimate $\tau = 0.1596$; estimates of the coefficient of risk aversion, nevertheless, are associated to a large standard error.

5. Concluding remarks

In this paper we use a MSM procedure to structurally estimate the three parameters describing preferences in a life-cycle model featuring temptation disutility. It is known that the consumption-saving predictions of models with and without temptation disutility are empirically undistinguishable for a proper choice of the preference parameters. Tempted agents however have a desire to commit their actions. In particular they eschew situations of exposure to more tempting alternatives by saving in quasi-liquid assets. To identify our parameters we thus consider as target variables field data on liquid and quasi-liquid (retirement) wealth holdings, and construct a model where the agent may choose how to allocate savings in two assets, one of which allows only limited withdrawal before retirement.

We estimate a significantly positive degree of temptation, and the unrestricted model generates a loss function that is more than three times smaller than the one under the restricted model where temptation is absent. Our results do not change qualitatively if we choose a different calibration of the exogenous parameters, but are affected quantitatively by the parameters describing the features of the quasi-liquid assets. While providing a better fit than the restricted model, a model featuring temptation preferences does not match exactly the observed data on quasi-liquid wealth holdings. In particular, it under-estimates

[12] The case $\alpha = 0$ trivially implies a marginal propensity to consume of 1, because the person would consume everything and leave no bequests.

quasi-liquid wealth at young ages, and over-estimates it near retirement. Hence temptation preferences contribute only partially to explain the holding of retirement wealth that is observed in the US, provided that the model is correctly specified and the target moments are adequately calculated.

Our economic environment is stylised and future work should enrich the realism of the framework, especially including other forms of non-liquid savings (such as housing) that may influence portfolio decisions. In so doing, the fit of the model to the observed data is likely to improve. A further robustness check should analyse whether different target moments change the estimated parameter values; if the model is correctly specified, the choice of the target moments does not alter the estimates. A comparison between temptation and quasi-hyperbolic discounting models is also worthwhile, because a quasi-hyperbolic agent also finds commitment valuable (see Laibson, 1997, and Angeletos et al., 2001).

A. Appendix

A.1. Reducing the dimensionality of the state space

To reduce the dimensionality of the state space, we standardise the variables dividing by permanent income P_t. A clear link emerges between value functions based on non-standardised (upper-case letters) and standardised (lower-case letters) variables. Let us first define with C_t^* and I_t^* the optimal consumption and investment choices made in period t, and with lower-case letters any variable that in period t is divided by permanent income P_t. At terminal age $t = D$,

$$\begin{aligned} V_D\left(X_D, H_D, P_D\right) &= U\left(C_D^*, X_D, H_D\right) + \beta E_D\left[B\left(X_{D+1}, H_{D+1}, P_{D+1}\right)\right] \\ &= \left(P_D\right)^{1-\rho} U\left(c_D^*, x_D, h_D\right) + \beta\left(P_{D+1}\right)^{1-\rho} E_D\left[B\left(x_{D+1}, h_{D+1}\right)\right] \\ &= \left(P_D\right)^{1-\rho}\left(U\left(c_D^*, x_D, h_D\right) + \beta\left(\frac{P_{D+1}}{P_D}\right)^{1-\rho} E_D\left[B\left(x_{D+1}, h_{D+1}\right)\right]\right) \\ &= \left(P_D\right)^{1-\rho} V_D\left(x_D, h_D\right) \end{aligned} \tag{25}$$

given specifications (3), (6)-(8) of the functions $U(\cdot)$ and $B(\cdot)$.

At age $t = D-1$,

$$V_{D-1}\left(X_{D-1},H_{D-1},P_{D-1}\right)=\left(P_{D-1}\right)^{1-\rho}\left(\begin{array}{l}U\left(c_{D-1}^{*},x_{D-1},h_{D-1}\right)\\+\pi_D E_{D-1}\left[\beta\left(\frac{P_{D-1}}{P_D}\right)^{1-\rho}V_D\left(x_D,h_D\right)\right]\\+\left(1-\pi_D\right)E_{D-1}\left[\beta\left(\frac{P_{D-1}}{P_D}\right)^{1-\rho}B_D\left(x_D,h_D\right)\right]\end{array}\right) \tag{26}$$
$$=\left(P_{D-1}\right)^{1-\rho}V_{D-1}\left(x_{D-1},h_{D-1}\right)$$

The dynamic budget constraints can be rephrased as

$$x_D\le R^X\left(x_{D-1}-c_{D-1}^{*}-i_{D-1}^{*}\right)\frac{P_{D-1}}{P_D}+a \tag{27}$$

and

$$h_D\le R^H\left(h_{D-1}+i_{D-1}^{*}\right)\frac{P_{D-1}}{P_D} \tag{28}$$

It can be shown that the following correspondence holds at a generic time t:

$$V_t\left(X_t,H_t,P_t\right)=\left(P_t\right)^{1-\rho}V_t\left(x_t,h_t\right) \tag{29}$$

where

$$\begin{aligned}V_t\left(x_t,h_t\right)&=U\left(c_t^{*},x_t,h_t\right)\\&\quad+\pi_{t+1}E_t\left[\beta_{t+1}^{*}V_{t+1}\left(x_{t+1},h_{t+1}\right)\right]+\left(1-\pi_{t+1}\right)E_t\left[\beta_{t+1}^{*}B\left(x_{t+1},h_{t+1}\right)\right]\end{aligned} \tag{30}$$

subject to the budget constraints

$$x_{t+1}\le R^X\left(x_t-c_t^{*}-i_t^{*}\right)\frac{P_t}{P_{t+1}}+\frac{Y_{t+1}}{P_{t+1}} \tag{31}$$

$$h_{t+1}\le R^H\left(x_t+i_t^{*}\right)\frac{P_t}{P_{t+1}} \tag{32}$$

with

$$\beta_{t+1}^{*}=\beta\left(\frac{P_{t+1}}{P_t}\right)^{1-\rho} \tag{33}$$

$$\frac{P_t}{P_{t+1}}=\begin{cases}\dfrac{1}{G_{t+1}N_{t+1}} & t<K\\ 1 & t\ge K\end{cases}\qquad\qquad\frac{Y_{it+1}}{P_{t+1}}=\begin{cases}Z_{it+1} & t<K\\ a & t\ge K\end{cases} \tag{34}$$

Since $V_t\left(X_t, H, P_t\right)$ and $V_t\left(x_t, h_t\right)$ differ by just a scale factor, it is equivalent to maximise either function. To reduce the dimensionality of the state space, we base our analysis on equation (30).

A.2. Solution to the life-cycle problem

To derive the system of equations (15), we compute the first-order condition of $V_t\left(x_t, h_t\right)$ with respect to c_t,

$$\begin{aligned}
&\frac{\partial V_t\left(x_t, h_t\right)}{\partial c_t} = 0 \Rightarrow \\
&\frac{\partial U\left(c_t, x_t, h_t\right)}{\partial c_t} = \pi_{t+1} E_t\left[\beta_{t+1} R^X \frac{\partial V_{t+1}\left(x_{t+1}, h_{t+1}\right)}{\partial x_{t+1}}\right] \\
&\qquad + \left(1-\pi_{t+1}\right) E_t\left[\beta_{t+1} R^X \frac{\partial B\left(x_{t+1}, h_{t+1}\right)}{\partial x_{t+1}}\right]
\end{aligned} \tag{35}$$

and with respect to i_t,

$$\begin{aligned}
&\frac{\partial V_t\left(x_t, h_t\right)}{\partial i_t} = 0 \Rightarrow \\
&\pi_{t+1} E_t\left[R^X \frac{\partial V_{t+1}\left(x_{t+1}, h_{t+1}\right)}{\partial x_{t+1}}\right] + \left(1-\pi_{t+1}\right) E_t\left[R^X \frac{\partial B\left(x_{t+1}, h_{t+1}\right)}{\partial x_{t+1}}\right] = \\
&\pi_{t+1} E_t\left[R^H \frac{\partial V_{t+1}\left(x_{t+1}, h_{t+1}\right)}{\partial h_{t+1}}\right] + \left(1-\pi_{t+1}\right) E_t\left[R^H \frac{\partial B\left(x_{t+1}, h_{t+1}\right)}{\partial h_{t+1}}\right]
\end{aligned} \tag{36}$$

We then use the envelope theorem, deriving $V_t\left(x_t, h_t\right)$ with respect to x_t,

$$\begin{aligned}
\frac{\partial V_t\left(x_t, h_t\right)}{\partial x_t} &= \frac{\partial U\left(c_t, x_t, h_t\right)}{\partial x_t} \\
&\quad + \pi_{t+1} E_t\left[\beta_{t+1} R^X \frac{\partial V_{t+1}\left(x_{t+1}, h_{t+1}\right)}{\partial x_{t+1}}\right] \\
&\quad + \left(1-\pi_{t+1}\right) E_t\left[\beta_{t+1} R^X \frac{\partial B\left(x_{t+1}, h_{t+1}\right)}{\partial x_{t+1}}\right]
\end{aligned} \tag{37}$$

and with respect to h_t,

$$\frac{\partial V_t(x_t,h_t)}{\partial h_t} = \frac{\partial U(c_t,x_t,h_t)}{\partial h_t} + \pi_{t+1}E_t\left[\beta_{t+1}R^H \frac{\partial V_{t+1}(x_{t+1},h_{t+1})}{\partial h_{t+1}}\right] + (1-\pi_{t+1})E_t\left[\beta_{t+1}R^H \frac{\partial B(x_{t+1},h_{t+1})}{\partial h_{t+1}}\right] \tag{38}$$

Using (35) and (37) it turns out that

$$\frac{\partial V_t(x_t,h_t)}{\partial x_t} = \frac{\partial U(c_t,x_t,h_t)}{\partial c_t} + \frac{\partial U(c_t,x_t,h_t)}{\partial x_t} \tag{39}$$

Using (35), (36) and (38), furthermore,

$$\frac{\partial V_t(x_t,h_t)}{\partial h_t} = \frac{\partial U(c_t,x_t,h_t)}{\partial c_t} + \frac{\partial U(c_t,x_t,h_t)}{\partial h_t} \tag{40}$$

Hence we can rewrite the FOCs using (39) and (40):

$$\frac{\partial U}{\partial c_t} = \pi_{t+1}E_t\left[\beta_{t+1}R^X\left(\frac{\partial U}{\partial c_{t+1}} + \frac{\partial U}{\partial x_{t+1}}\right)\right] + (1-\pi_{t+1})E_t\left[\beta_{t+1}R^X \frac{\partial B}{\partial x_{t+1}}\right] \tag{41}$$

$$\frac{\partial U}{\partial c_t} = \pi_{t+1}E_t\left[\beta_{t+1}R^H\left(\frac{\partial U}{\partial c_{t+1}} + \frac{\partial U}{\partial h_{t+1}}\right)\right] + (1-\pi_{t+1})E_t\left[\beta_{t+1}R^H \frac{\partial B}{\partial h_{t+1}}\right] \tag{42}$$

References

Altonji, J.G., and **L.M. Segal** (1996), "Small-Sample Bias in GMM Estimation of Covariance Structures", *Journal of Business Economics and Statistics*, Vol. 14(3), pp. 353-366.

Ameriks, J., A. Caplin, J. Leahy, and **T. Tyler** (2007), "Measuring Self-Control Problems", *American Economic Review*, Vol. 97(3), pp. 966-972.

Angeletos, G.M., D.I. Laibson, A. Repetto, J. Tobacman, and **S. Weinberg** (2001), "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation", *Journal of Economic Perspectives*, Vol. 15(3), pp. 47-68.

Attanasio, O., and **H. Low** (2004), "Estimating Euler Equations", *Review of Economic Dynamics*, Vol. 7(2), pp. 406-435.

Brennan, M.J., and **A. Subrahmayam** (1996), "Market Microstructure and Asset Pricing: on the Compensation from Illiquidity in Stock Returns", *Journal of Financial Economics*, Vol. 41(3), pp. 441-464.

Brown, A.L., C.F. Camerer, and **Z.E. Chua** (2007), "Learning and Visceral Temptation in Dynamic Savings Experiments", mimeo.

Bucciol, A. (2008), "Social Security, Dynamic Efficiency and Self-Control Problems", QSPS Discussion Paper n. 2008-5.

Cagetti, M. (2003), "Wealth Accumulation over the Life Cycle and Precautionary Saving", *Journal of Business and Economic Statistics*, Vol. 21(3), pp. 339-353.

Carroll, C.D. (1992), "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence", *Brookings Papers on Economic Activity*, Vol. 1992(2), pp. 61-156.

Carroll, C.D., and **A.A. Samwick** (1997), "The Nature of Precautionary Wealth", *Journal of Monetary Economics*, Vol. 40(1), pp. 41-71.

Deaton, A. (1991), "Saving and Liquidity Constraints", *Econometrica*, Vol. 59(5), pp. 1221-1248.

DeJong, D.N., and **M. Ripoll** (2007), "Do Self-Control Preferences Help Explain the Puzzling Behavior of Asset Prices?", *Journal of Monetary Economics*, Vol. 54(4), pp. 1035-1050.

Engen, E.M., W.G. Gale, and **C. Uccello** (1999), "The Adequacy of Household Saving", *Brookings Papers on Economic Activity*, Vol. 1999(2), pp. 65-165.

Gourinchas, P., and **J.A. Parker** (2002), "Consumption over the Life-Cycle", *Econometrica*, Vol. 70(1), pp. 47-89.

Gul, F., and **W. Pesendorfer** (2001), "Temptation and Self-Control", *Econometrica*, Vol. 69(6), pp. 1403-1435.

Gul, F., and **W. Pesendorfer** (2004), "Self-Control and the Theory of Consumption", *Econometrica*, Vol. 72(1), pp. 119-158.

Gul, F., and **W. Pesendorfer** (2007), "Harmful Addiction", *Review of Economic Studies*, Vol. 74(1), pp. 147-172.

Huang, K.H.D., Z. Liu, and **J.Q. Zhu** (2005), "Temptation and Self-Control: Some Evidence and Applications", SED meeting paper n. 770.

Hubbard, G., J.S. Skinner, and **S. Zeldes** (1995), "Precautionary Saving and Social Insurance", *Journal of Political Economy*, Vol. 103(2), pp. 360-399.

Hurd, M.D. (1989), "Mortality Risk and Bequests", *Econometrica*, Vol. 57(4), pp. 779-813.

Imrohoroglu, A., S. Imrohoroglu, and **D.H. Joines** (1999), "Social Security in an Overlapping Generations Economy with Land", *Review of Economic Dynamics*, Vol. 2(3), pp. 638-655.

Krusell, B., P. Kuruscu, and **A. Smith** (2002), "Time Orientation and Asset Prices", *Journal of Monetary Economics*, Vol. 49(1), pp. 107-135.

Kumru, C., and **A. Thanopoulos** (2008), "Social Security and Self-Control Preferences", *Journal of Economic Dynamics & Control*, Vol. 32(3), pp. 757-778.

Laibson, D.I. (1997), "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics*, Vol. 112(2), pp. 443-477.

Laibson, D.I., A. Repetto, and **J. Tobacman** (2007), "Estimating Discount Function with Consumption Choices over the Lifecycle", mimeo.

Munnell, A.H., M. Soto, J. Libby, and **J. Prinzivalli** (2006), "Investment Returns: Defined Benefit Vs. 401(k) Plans", CRR Issue in Brief n. 52.

Pakes, A., and **D. Pollard** (1989), "Simulation and the Asymptotics of Optimization Estimators", *Econometrica*, Vol. 57(5), pp. 1027-1057.

Pàstor, L., and **R.F. Stambaugh** (2003), "Liquidity Risk and Expected Stock Returns", *Journal of Political Economy*, Vol. 111(3), pp. 642-685.

Scholz, J.K., A. Seshadri, and **S. Khitatrakun** (2006), "Are Americans Saving "Optimally" for Retirement?", *Journal of Political Economy*, Vol. 51(3), pp. 607-643.

Social Security Administration (2008), "The 2008 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds", April 10.

Watson Wyatt (2008), "Defined Benefit Vs. 401(k) Plans: Investment Returns from 2003-2006", Insider January.

Zeldes, S.P. (1989), "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence", *Quarterly Journal of Economics*, Vol. 104(2), pp. 275-298.

Table 1. Target moments

Description	Name	Moment	Standard error
Liquid wealth over income; ages 41-45	LIQ 41-45	0.5711	0.0176
Quasi-liquid wealth over income; ages 41-45	QLIQ 41-45	0.3570	0.0100
Liquid wealth over income; ages 61-65	LIQ 61-65	1.5399	0.0690
Quasi-liquid wealth over income; ages 61-65	QLIQ 61-65	1.1884	0.0253

Note: Number of observations: 1,763 (age 41-45) and 1,455 (age 61-65).

Table 2. Calibration of the exogenous parameters (1st stage MSM)

Description	Symbol	Value	Std . error	Source
Survival probability	$\{\pi_t, t = 2,\ldots, D+1\}$	1960 cohort	-	Social Security Administration website
Number of working years	K	39	0.042	SCF (authors' calculation)
Deterministic income profile, constant	ϕ_1	6.742	0.018	Scholz et al. (2006)
Deterministic income profile, linear trend	ϕ_2	0.174	0.002	Scholz et al. (2006)
Deterministic income profile, quadratic trend	ϕ_3	-0.195	0.003	Scholz et al. (2006)
Pension replacement rate	a	0.413	-	Social Security Administration (2008)
Std. dev. permanent income risk	σ_n	0.147	0.010	Carroll and Samwick (1997)
Std. dev. transitory income risk	σ_z	0.210	0.013	Carroll and Samwick (1997)
Liquid asset return	R^X	1.034	0.003	Gourinchas and Parker (2002)
Quasi-liquid asset return	R^H	1.05	-	
Withdrawal limit before retirement	δ	0.50	-	
Degree of altruism	α	0	-	
Fraction of quasi-liquid wealth that can be bequeathed	κ	1	-	

Table 3. Benchmark structural estimation (2nd stage MSM)

	Data	Robust *W*		Optimal *W*	
		CRRA	Temptation	CRRA	Temptation
Parameters estimates $\hat{\theta}$					
τ		0	0.1594	0	0.1596
		(-)	(0.0036)	(-)	(0.0077)
β		0.9308	0.9965	0.9312	0.9960
		(0.0542)	(0.0089)	(0.0523)	(0.0041)
ρ		1.6834	0.8857	1.6845	0.8957
		(1.5398)	(0.1054)	(1.4977)	(0.0235)
Target moments					
LIQ 41-45	0.5711	0.5582	0.5183	0.5647	0.5354
QLIQ 41-45	0.3570	0.2522	0.1230	0.2523	0.1355
LIQ 61-65	1.5399	1.9946	1.8013	2.0204	1.8652
QLIQ 61-65	1.1884	2.2885	1.4015	2.2893	1.4789
Loss function		0.6068	0.1875	0.5952	0.1866
Test over-ident. restr. (p-value)		0.0000	0.0000	0.0000	0.0000

Note: MSM standard errors corrected for first-stage uncertainty in parentheses.

Table 4. Robustness: quasi-liquid asset features

Optimal *W*	**Benchmark**	R^H		δ	
		= 1.04	**= 1.06**	**= 0.25**	**= 0.75**
Parameters estimates $\hat{\theta}$					
τ	0.1596	0.2199	0.1075	0.0825	0.3063
	(0.0077)	(0.0009)	(0.0122)	(0.0016)	(0.0598)
β	0.9960	1.0122	0.9774	0.9902	0.9945
	(0.0041)	(0.0131)	(0.0082)	(0.0012)	(0.0312)
ρ	0.8957	0.8968	0.9494	0.7152	0.1503
	(0.0235)	(0.1124)	(0.1617)	(0.0124)	(0.1956)
Target moments					
LIQ 41-45	0.5354	0.5245	0.4905	0.5081	0.5259
QLIQ 41-45	0.1355	0.1681	0.0908	0.1215	0.1191
LIQ 61-65	1.8652	1.7881	1.7352	1.7643	1.7406
QLIQ 61-65	1.4789	1.4295	1.3883	1.3857	1.3685
Loss function	0.1866	0.1355	0.2336	0.1862	0.1858
Test over-ident. restr. (p-value)	0.0000	0.0000	0.0000	0.0000	0.0000

Note: Estimates using the optimal weighting matrix. MSM standard errors corrected for first-stage uncertainty in parentheses.

Table 5. Robustness: replacement rate and bequest parameters

Optimal *W*	a		**Surv. prob.**	$\alpha = 1.5$	
	= 0.27	**= 0.56**	**1980 cohort**	κ **= 1**	κ **= 0.50**
Parameters estimates $\hat{\theta}$					
τ	0.1580	0.1716	0.1621	0.1581	0.1604
	(0.0018)	(0.0219)	(0.0018)	(0.0060)	(0.0088)
β	0.9919	1.0236	1.0068	0.9909	0.9913
	(0.0045)	(0.0031)	(0.0051)	(0.0349)	(0.0225)
ρ	0.8952	0.6968	0.7763	0.9017	0.8927
	(0.0417)	(0.0992)	(0.0425)	(0.3866)	(0.2272)
Target moments					
LIQ 41-45	0.4661	0.5540	0.5270	0.4813	0.4604
QLIQ 41-45	0.1360	0.1033	0.1332	0.1308	0.1324
LIQ 61-65	2.0037	1.3315	1.8755	2.0026	1.9650
QLIQ 61-65	1.4829	1.1892	1.4689	1.4506	1.4631
Loss function	0.1968	0.1796	0.1881	0.1949	0.1964
Test over-ident. restr. (p-value)	0.0000	0.0000	0.0000	0.0000	0.0000

Note: Estimates using the optimal weighting matrix. MSM standard errors corrected for first-stage uncertainty in parentheses.

Figure 1. Average simulated consumption over age

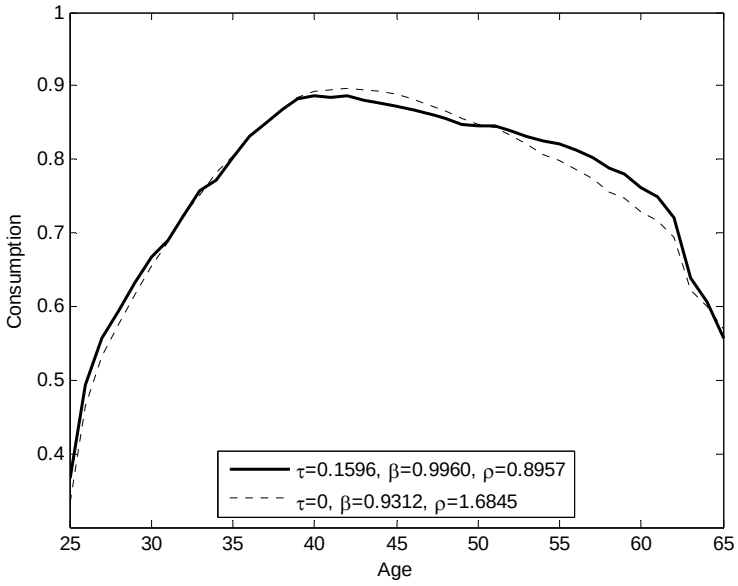

Figure 2. Average simulated liquid wealth over age

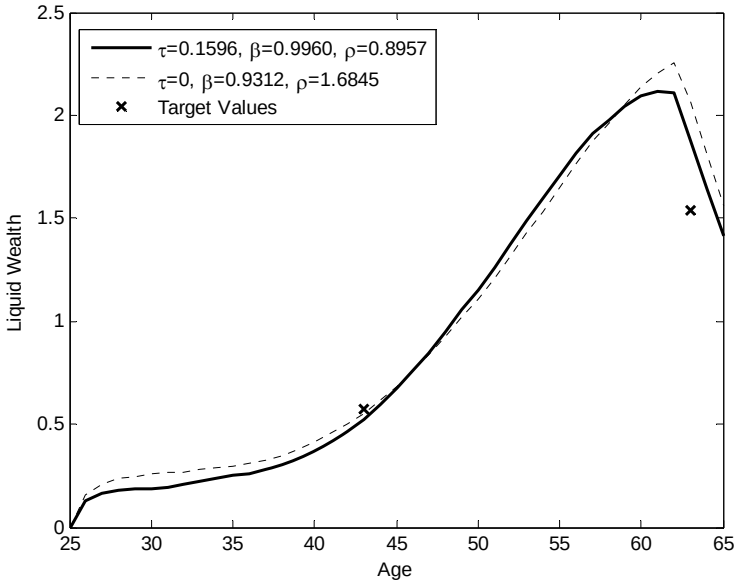

Figure 3. Average simulated quasi-liquid wealth over age