

Solving Constrained Consumption–Investment Problems by Simulation of Artificial Market Strategies

Björn Bick, Holger Kraft

Department of Finance, Goethe University Frankfurt am Main, 60323 Frankfurt am Main, Germany
{bick@finance.uni-frankfurt.de, holgerkraft@finance.uni-frankfurt.de}

Claus Munk

Department of Finance, Copenhagen Business School, DK-2000 Frederiksberg, Denmark, cm.fi@cbs.dk

Utility-maximizing consumption and investment strategies in closed form are unknown for realistic settings involving portfolio constraints, incomplete markets, and potentially a high number of state variables. Standard numerical methods are hard to implement in such cases. We propose a numerical procedure that combines the abstract idea of artificial, unconstrained complete markets, well-known closed-form solutions in affine or quadratic return models, straightforward Monte Carlo simulation, and a standard iterative optimization routine. Our method provides an upper bound on the wealth-equivalent loss compared to the unknown optimal strategy, and it facilitates our understanding of the economic forces at play by building on closed-form expressions for the strategies considered. We illustrate and test our method on the life-cycle problem of an individual who receives unspanned labor income and cannot borrow or short sell. The upper loss bound is small, and our method performs well in comparison with two existing methods.

Key words: optimal consumption and investment; numerical solution; labor income; incomplete markets; artificially unconstrained markets; welfare loss

History: Received June 30, 2011; accepted May 24, 2012, by Wei Xiong, finance. Published online in *Articles in Advance* September 11, 2012.

1. Introduction

Utility-maximizing consumption and investment strategies are notoriously difficult to compute when markets are incomplete and strategies are constrained. Closed-form solutions are only known in unrealistic special cases. Numerical dynamic programming is frequently used but suffers from the curse of dimensionality. The existing alternative numerical methods are complex. Little is known about the precision of any of these numerical methods. This paper introduces a simple numerical approach combining (i) the idea of artificially unconstrained and complete markets, (ii) well-known closed-form solutions for unconstrained consumption/portfolio problems in affine or quadratic settings with time-additive power utility, (iii) straightforward Monte Carlo simulation to evaluate various simple consumption and investment strategies, and (iv) a standard optimization routine. We refer to our approach as SAMS, short for *simulation of artificial markets strategies*.

In addition to its relative simplicity, SAMS has a number of attractive features. First, SAMS applies to high-dimensional models as long as the relevant state variables have affine or quadratic dynamics, which is assumed in most existing models. Second, the consumption and investment strategy produced

by SAMS is given in closed form (involving some parameters that we optimize over as a part of the approach) and is thus easy to interpret. Third, in contrast to the mainstream numerical methods, SAMS also delivers an upper bound on the welfare loss the individual incurs by using the strategy suggested by our procedure instead of the unknown optimal strategy.

For concreteness we focus in most of the paper on a classical life-cycle problem: a power-utility individual receives an unspanned labor income stream and has access to trade in a risk-free asset and a stock, but the investment in each asset must be between 0 and 100% of current financial wealth. This is a prime example of a problem with no closed-form solution, but with incomplete markets and a portfolio constraint that for realistic parameter values is binding in a substantial part of the (time, state) space.¹ SAMS produces a relatively simple closed-form, near-optimal consumption and investment strategy. The upper bound on the

¹Optimal unconstrained strategies have been derived in closed form for some settings with negative exponential utility and normally distributed income (Svensson and Werner 1993, Henderson 2005, Christensen et al. 2012) and for settings with deterministic or spanned labor income (Hakansson 1970, Bodie et al. 1992).

welfare loss from following this strategy depends on the assumed income–stock correlation and the ratio of initial financial wealth to initial annual income. In our benchmark parametrization of the model with a 50-year time horizon, the upper bound is below 0.5% of total wealth for most of the combinations of the correlation and the wealth–income ratio, and the highest upper bound is 1.1%. We show that the precision of our method is at least as good as two existing alternative methods, and we explain why our approach can handle problems in higher dimensions more efficiently than these alternative methods.

Concerning the economic properties of the solution to this life-cycle problem, we demonstrate that the optimal fraction of financial wealth invested in the stock *during* retirement, where income is assumed to be risk free, depends on the correlation between labor income and stock returns *before* retirement. If this correlation is high, the individual invests less in the stock market before retirement, so at retirement financial wealth will often be small compared to the present value of the risk-free retirement income. To obtain the desired risk–return balance, the individual will therefore invest a relatively large fraction of financial wealth in the stock after retirement. Furthermore, before retirement the human wealth depends on financial wealth because of the unspanned income risk and the portfolio constraints. We show that the human wealth is increasing in the ratio of financial wealth to initial income and is decreasing in the stock–income correlation.

Finally, we document that the excellent performance of SAMS is robust to variations in key parameter values and to an extension of the model to stochastic interest rates. The application to a general problem is also outlined.

Our method applies the idea of artificial financial markets introduced by Karatzas et al. (1991) and Cvitanić and Karatzas (1992). A constrained, incomplete-market consumption–investment problem can be embedded in a family of consumption–investment problems in artificially unconstrained, complete-market problems. For our specific problem each artificial market corresponds to a given choice of (i) the Sharpe ratio of an artificial asset allowing perfect hedging of income risk and (ii) a certain perturbation of the risk-free rate and stock drift. Both the Sharpe ratio and the perturbation are generally stochastic processes. The optimal solution in any of the artificial markets generates at least as high an expected utility as the unknown optimal solution in the true market. For the case without stochastic income, Karatzas et al. (1991) and Cvitanić and Karatzas (1992) showed—using a convex duality/martingale approach—that the optimal consumption and investment strategy under

incomplete markets and portfolio constraints is identical to the strategy that is optimal in the worst of all the artificial markets. Cvitanić et al. (2001) showed that the result also holds in the presence of stochastic income if the domain of the dual problem is enlarged appropriately. However, these papers provide no practical procedure for finding the worst artificial market and thus the optimal strategy. We focus on the subfamily of “simple” artificial markets where both the Sharpe ratio and the perturbation are simple functions of time characterized by a low number of constant parameters, because the optimal strategies in those markets are then known in closed form due to Liu (2007) and others. By minimizing the value function over these parameters we find the worst of the simple artificial markets, which gives an upper bound on the utility that can be obtained in the true market.

The optimal strategy in any of the simple artificial markets is generally infeasible in the true market because it involves the artificial asset and may violate the portfolio constraints. We transform it into a feasible strategy by ignoring the investment in the artificial asset and by “pruning” the remaining part of the strategy to make sure constraints are respected. This generates a family of feasible strategies parameterized by a low number of constants. We compute the expected utility associated with each strategy by Monte Carlo simulation, and we embed this in a standard optimization routine, leading to a feasible and near-optimal consumption and investment strategy in the true constrained and incomplete market. Comparing the expected utility of this strategy with the upper utility bound, we find an upper bound on the welfare loss—the utility loss stated in terms of total wealth—associated with our strategy. We find small upper bounds in our quantitative examples, and our comparison with two well-established alternative numerical methods indicates that the actual welfare loss is significantly smaller than the upper bound suggests.

Let us compare our method to the existing alternatives. Grid-based dynamic programming, the finite difference solution of the Hamilton–Jacobi–Bellman (HJB) equation, and Markov chain approximations (MCAs) are closely related and frequently applied methods for numerically solving low-dimensional consumption–investment problems related to the one we study (Brennan et al. 1997, Munk 2000, Cocco et al. 2005, Yao and Zhang 2005, Van Hemert 2010, Munk and Sørensen 2010). However, with the current technology, these methods cannot be efficiently implemented with four or more state variables and are computationally intensive even in lower dimensions. Hence, coarse grids have to be used despite the implied reduced precision. Moreover, relevant state variables such as wealth, income, or the wealth–income ratio tend to fluctuate considerably over the

life cycle so that an age-dependent scaling must be implemented to keep the state variables within the grid with high probability. The appropriate scaling has to be determined experimentally and depends on the specific setting and parameter values. In contrast, our SAMS approach is based on closed-form solutions, needs no scaling, can handle higher dimensions, and provides an upper bound on the error.

Various Monte Carlo simulation-based approaches that can potentially handle higher-dimensional problems have been proposed. The approach of Detemple et al. (2003) is based on Malliavin calculus and was originally formulated for complete markets, but can be applied in incomplete markets if the dual problem can be solved in explicit (or approximate) form. Cvitanić et al. (2003) suggested another (simpler and slower) method for complete markets and unconstrained portfolios. Garlappi and Skoulakis (2010) introduced a method for discrete-time problems based on Monte Carlo simulation, a Taylor expansion of the value function, and a certain decomposition of the state variables. An alternative Monte Carlo approach to solving portfolio problems builds on the American option pricing method of Carriere (1996) and Longstaff and Schwartz (2001). This approach was developed by Brandt et al. (2005), Bouchard et al. (2004), van Binsbergen and Brandt (2007), and Kojien et al. (2007, 2010), among others. For comparison with our SAMS approach, we implement the method of Kojien et al. (2007, 2010) because this was formulated directly for a lifetime consumption–investment problem. In contrast to our approach, this method does not deliver explicit solutions for the optimal consumption–investment choice. Furthermore, to handle portfolio constraints, their method must rely on high-dimensional *constrained* optimization algorithms, whereas simpler *unconstrained* optimization techniques are sufficient for our approach.

A different approach to solve portfolio problems numerically was introduced by Fahim et al. (2011). They studied semilinear partial differential equations that arise in portfolio problems after substituting the first-order conditions back into the Bellman equation. This allowed them to calculate the agent’s indirect utility at a given point in time, but did not directly provide a way to solve for the optimal consumption–investment choice. Furthermore, their theoretical convergence results assume that the indirect utility function grows exponentially, which is not satisfied in a standard lifetime consumption–investment problem with power utility. A variety of other methods have been proposed; see, for example, Kogan and Uppal (2000), Viceira (2001), and Das and Sundaram (2002).

By applying the idea of artificially unconstrained and complete markets, as we do, and the associated duality technique, Haugh et al. (2006) explained how to compute an upper bound on the expected utility from any given feasible consumption and investment strategy. A comparison of the expected utility delivered by the given strategy and the upper bound—both computed by Monte Carlo simulation—provides a measure of the performance of the strategy, an idea that we also apply. In contrast to their work, we search for the best possible strategy among a parameterized family of promising candidates motivated by simple artificial markets, and we also search over a parameterized family of upper utility bounds to find the tightest possible bound. We exploit the fact that the optimal strategies in some artificial markets are known in closed form.

2. The Problem

We first implement the approach for a specific problem that has been studied frequently in the literature and allows us to illustrate the power of our approach in a transparent way. The application of our approach in a more general setting is outlined in §9.

The individual has access throughout life to trade in an instantaneously risk-free asset (a bank account) and a risky asset (a stock or stock index). We assume a constant annualized risk-free rate given by r using continuous compounding; see §8 for an extension to stochastic interest rates. We let S_t denote the price of the stock at time t , and the price dynamics is assumed to be

$$dS_t = S_t[(r + \sigma_S \lambda_S) dt + \sigma_S dW_t],$$

where $W = (W_t)$ is a standard Brownian motion. Hence, σ_S is the volatility of the stock, and λ_S is the Sharpe ratio of the stock, both assumed constant and positive.

The individual earns a stochastic labor income rate Y_t until a predetermined retirement date \tilde{T} , after which the individual lives on until time $T > \tilde{T}$. We assume that

$$dY_t = Y_t[\alpha dt + \beta(\rho dW_t + \sqrt{1 - \rho^2} dW_{Yt})], \quad 0 \leq t \leq \tilde{T}, \quad (1)$$

where $W_Y = (W_{Yt})$ is another standard Brownian motion, independent of W . Here α is the expected growth rate of labor income, β is the income volatility, and ρ is the instantaneous correlation between stock returns and income growth. We assume that α , β , and ρ are all constants, but our analysis goes through with the deterministic age-related variations in α and β documented by Cocco et al. (2005). Unless $\beta = 0$ or $|\rho| = 1$, the investor faces an incomplete market,

because he cannot fully hedge against unfavorable

including the constraint that the investor cannot trade a specific risk, i.e., that the market is incomplete. For our problem, each artificial market is characterized by two real-valued stochastic processes, $\nu = (\nu_t)$ and $\lambda_I = (\lambda_{It})$. The process ν adjusts the drift of the stock and the risk-free rate because of the constraint $\pi_{St} \in [0, 1]$. Define $\nu_t^- = \max(-\nu_t, 0)$ and $\nu_t^+ = \max(\nu_t, 0)$. In the artificially unconstrained market corresponding to any given ν , the risk-free rate is assumed to be $\tilde{r}_t = r + \nu_t^-$ (instead of just r as in the true market), and the drift of the stock is assumed to be $r + \sigma_S \lambda_S + \nu_t^+ = \tilde{r}_t + \sigma_S \lambda_S + \nu_t$. If the unconstrained π_{St} is above 1, which often happens with substantial human wealth and low or even negative ρ , we increase the risk-free rate and keep the drift of the stock fixed, which makes the stock less attractive relative to the bank account. This corresponds to a negative value of ν_t . If the unconstrained π_{St} is below 0, which may happen when ρ , β , and γ are relatively high and λ_S relatively low, we increase the drift of the stock and keep the risk-free rate fixed, making the stock more attractive relative to the bank account. This corresponds to a positive value of ν_t .

The process λ_I relates to the unspanned income risk that the individual faces until retirement if $\beta > 0$ and $|\rho| < 1$. The artificially unconstrained markets allow for investment in an “income contract” characterized by the market price of risk λ_{It} associated with the standard Brownian motion W_Y . The time t price is I_t and evolves according to

$$dI_t = I_t [(\tilde{r}_t + \lambda_{It}) dt + dW_{Yt}]. \quad (8)$$

Here, λ_I can be positive or negative. We let π_{It} be the fraction of wealth invested in the income contract.

Every pair of processes (ν, λ_I) , satisfying certain technical conditions, defines an artificial market. There are no constraints on the consumption and investment strategy in the artificial markets except for the standard integrability conditions and the constraint that consumption and terminal wealth have to be nonnegative. Because labor income is perfectly hedgeable in the artificial markets, we do not need $X_t \geq 0$ for all t to ensure $X_T \geq 0$, as we do in the true market.

Let $J(t, x, y; \nu, \lambda_I)$ denote the indirect utility in the artificial market corresponding to (ν, λ_I) . A strategy (c, π_S) that is feasible in the true market will, together with a zero investment in the income contract, lead to at least the same expected utility in any of the artificial markets as in the true market. The reason is that the risk-free rate and the return on the risky investment is at least as big in the artificial markets, and hence terminal wealth will also be at least as big. Many other strategies are feasible in the artificial markets, so the indirect utility in each artificial market

is greater than or equal to the indirect utility in the true market, $J(t, x, y; \nu, \lambda_I) \geq J(t, x, y)$. Karatzas et al. (1991), Cvitanic and Karatzas (1992), and Cvitanic et al. (2001) showed that the minimum of the indirect utility $J(t, x, y; \nu, \lambda_I)$ over all the processes (ν, λ_I) satisfying certain technical conditions is equal to the indirect utility in the true constrained market, i.e., the solution in the true constrained market is equal to the solution in the worst of all the artificially unconstrained markets. Alas, because we cannot compute $J(t, x, y; \nu, \lambda_I)$ for all (ν, λ_I) , we cannot minimize over (ν, λ_I) , so this result does not generally provide a way of finding the optimal constrained strategy.

However, we can compute $J(t, x, y; \nu, \lambda_I)$ in some artificial markets. To keep the solution tractable, we focus on the artificial markets with deterministic (ν, λ_I) and use the notation $\nu(t)$ instead of ν_t and similarly for λ_I . The pair (ν, λ_I) representing the worst artificial market will presumably depend on financial wealth and income (and age) and will thus be stochastic processes.⁴ Our method could be extended to certain exogenous stochastic processes ν and λ_I . As long as the price dynamics in the artificial market has an affine or quadratic structure (Liu 2007), closed-form solutions exist (in some cases one or more simple ordinary differential equations have to be solved numerically), but the solutions will be more complex with stochastic (ν, λ_I) . Apparently, we cannot allow ν or λ_I to depend explicitly on wealth and still obtain closed-form solutions.⁵ As we report below, the method is already very precise when restricted to simple deterministic (ν, λ_I) .

At retirement, income becomes risk free, which will presumably lead to a big shift in the allocation of the investment between the risk-free asset and the risky asset, so that a constraint that is binding just before retirement may not be binding immediately after and vice versa. Therefore, we allow for different $\nu(t)$ in retirement and in the active phase as represented by $\nu_R(t)$ and $\nu_A(t)$.

⁴ In the HJB equation corresponding to the problem (4), π_S has to maximize $\pi_S \lambda_S J_x + \frac{1}{2} \pi_S^2 \sigma_S^2 x J_{xx} + \beta \rho \pi_S y J_{xy}$. If we impose the constraint $\pi_S \leq 1$ and let m denote the associated nonnegative Lagrange multiplier, the Lagrangian consists of the terms listed before plus $m(1 - \pi_S)$. By maximizing with respect to π_S , we find

$$\pi_S = -\frac{J_x}{x J_{xx}} \frac{\lambda_S - m/J_x}{\sigma_S} - \frac{\beta \rho y J_{xy}}{\sigma_S x J_{xx}}.$$

This shows that the appropriate reduction of the Sharpe ratio is closely related to the Lagrangian multiplier associated with the constraint. For $\pi_S = 1$, we get $m/J_x = \lambda_S + \beta \rho y J_{xy}/J_x + \sigma_S x J_{xy}/J_x$, which depends on x and y as well as the age and risk aversion of the individual.

⁵ In the presence of labor income, it is generally difficult to describe the domain of dual minimizers (Cvitanic et al. 2001), and some nondeterministic dual controls (ν, λ_I) may not be admissible.

In retirement, there is no income risk, so λ_I is irrelevant and the artificial markets are just characterized by ν_R . For any function $\nu_R(t)$, the solution to the utility maximization in the corresponding artificial unconstrained market is stated below; see the appendix for proofs.

THEOREM 1. *The indirect utility during retirement in the artificial market characterized by $\nu_R(t)$ is given by*

$$J_R(t, x; \nu_R) = \frac{1}{1-\gamma} g_R(t; \nu_R)^\gamma (x + \Upsilon Y_{\tilde{T}} F_R(t; \nu_R))^{1-\gamma},$$

where

$$F_R(t; \nu_R) = \int_t^T e^{-\int_t^u (r + \nu_R(\tau)^-) d\tau} du,$$

$$g_R(t; \nu_R) = \varepsilon^{1/\gamma} e^{-\int_t^T h_R(\nu_R(\tau)) d\tau} + \int_t^T e^{-\int_t^u h_R(\nu_R(\tau)) d\tau} du,$$

$$h_R(\nu_R(\tau)) = \frac{\delta}{\gamma} + \frac{\gamma-1}{\gamma} (r + \nu_R(\tau)^-) + \frac{\gamma-1}{2\gamma^2} \left(\lambda_S + \frac{\nu_R(\tau)}{\sigma_S} \right)^2.$$

The corresponding optimal consumption and investment strategy is

$$c_t = \frac{X_t + \Upsilon Y_{\tilde{T}} F_R(t; \nu_R)}{g_R(t; \nu_R)}, \tag{9}$$

$$\pi_{St} = \left[1 + \frac{\Upsilon Y_{\tilde{T}} F_R(t; \nu_R)}{X_t} \right] \frac{\sigma_S \lambda_S + \nu_R(t)}{\gamma \sigma_S^2}. \tag{10}$$

In retirement, income affects the optimal strategy only through the addition of human capital to current financial wealth and thus drives up consumption and the risky investment.

Note that during retirement, π_S will never be negative because λ_S is positive, so we can focus on the constraint $\pi_S \leq 1$. To avoid $\pi_S > 1$, we may need $\nu_R < 0$. Then F_R and the present value of future income are smaller, and thus π_S is indeed smaller. Because h_R can be smaller or bigger (than with $\nu_R = 0$), it is not clear whether g_R is smaller or bigger, so the effect on consumption is not obvious.

Let $J_A(t, x, y; \nu_A, \nu_R, \lambda_I)$ denote the indirect utility function in the active phase in this artificial market. We have the boundary condition

$$\begin{aligned} J_A(\tilde{T}, x, y; \nu_A, \nu_R, \lambda_I) \\ = J_R(\tilde{T}, x; \nu_R) = \frac{1}{1-\gamma} g_R(\tilde{T}; \nu_R)^\gamma (x + \Upsilon Y_{\tilde{T}} F_R(\tilde{T}; \nu_R))^{1-\gamma}. \end{aligned}$$

Via this boundary condition, the indirect utility in the active phase will also depend on the perturbation $\nu_R(t)$ of the expected returns on the risk-free asset and the stock in retirement.

THEOREM 2. *The indirect utility before retirement in the artificial market characterized by $\nu_A(t)$ and $\lambda_I(t)$ is given by*

$$\begin{aligned} J_A(t, x, y; \nu_A, \nu_R, \lambda_I) \\ = \frac{1}{1-\gamma} g_A(t; \nu_A, \nu_R, \lambda_I)^\gamma (x + y F_A(t; \nu_A, \nu_R, \lambda_I))^{1-\gamma}, \end{aligned}$$

where

$$\begin{aligned} F_A(t; \nu_A, \nu_R, \lambda_I) &= e^{-\int_t^{\tilde{T}} r_A(\nu_A(u), \lambda_I(u)) du} \Upsilon F_R(\tilde{T}; \nu_R) \\ &\quad + \int_t^{\tilde{T}} e^{-\int_t^u r_A(\nu_A(\tau), \lambda_I(\tau)) d\tau} du, \end{aligned}$$

$$\begin{aligned} g_A(t; \nu_A, \nu_R, \lambda_I) &= e^{-\int_t^{\tilde{T}} h_A(\nu_A(u), \lambda_I(u)) du} g_R(\tilde{T}; \nu_R) \\ &\quad + \int_t^{\tilde{T}} e^{-\int_t^u h_A(\nu_A(\tau), \lambda_I(\tau)) d\tau} du, \end{aligned}$$

$$\begin{aligned} r_A(\nu_A(t), \lambda_I(t)) &= r + \nu_A(t)^- - \alpha + \beta \rho \left(\lambda_S + \frac{\nu_A(t)}{\sigma_S} \right) \\ &\quad + \beta \sqrt{1-\rho^2} \lambda_I(t), \end{aligned} \tag{11}$$

$$\begin{aligned} h_A(\nu_A(t), \lambda_I(t)) &= \frac{\delta}{\gamma} + \frac{\gamma-1}{\gamma} (r + \nu_A(t)^-) \\ &\quad + \frac{\gamma-1}{2\gamma^2} \left[\left(\lambda_S + \frac{\nu_A(t)}{\sigma_S} \right)^2 + \lambda_I(t)^2 \right]. \end{aligned}$$

The corresponding optimal consumption and investment strategy is

$$c_t = \frac{X_t + Y_t F_A(t; \nu_A, \nu_R, \lambda_I)}{g_A(t; \nu_A, \nu_R, \lambda_I)}, \tag{12}$$

$$\begin{aligned} \pi_{St} &= \frac{\sigma_S \lambda_S + \nu_A(t)}{\gamma \sigma_S^2} \\ &\quad + \frac{Y_t F_A(t; \nu_A, \nu_R, \lambda_I)}{X_t} \left[\frac{\sigma_S \lambda_S + \nu_A(t)}{\gamma \sigma_S^2} - \frac{\beta \rho}{\sigma_S} \right], \end{aligned}$$

$$\begin{aligned} \pi_{It} &= \frac{\lambda_I(t)}{\gamma} + \frac{Y_t F_A(t; \nu_A, \nu_R, \lambda_I)}{X_t} \\ &\quad \cdot \left[\frac{\lambda_I(t)}{\gamma} - \beta \sqrt{1-\rho^2} \right]. \end{aligned} \tag{13}$$

Before retirement, income affects the optimal strategy via the addition of human wealth to financial wealth, just as in retirement. Moreover, the term $-(\beta \rho / \sigma_S)(Y_t F_A(t) / X_t)$ adjusts the explicit investment in the stock by the implicit stock investment in human wealth through the income–stock correlation ρ . This follows the intuition of Bodie et al. (1992).

We want to minimize the indirect utility over the selected artificial markets because that provides an upper bound for the indirect utility in the true constrained market. To perform the minimization, we need to parameterize the functions $\nu_R(t)$, $\nu_A(t)$, and $\lambda_I(t)$.

First, consider $\nu_R(t)$. If $\lambda_S < \gamma\sigma_S$, the constraint will not be active without income, i.e., just before the terminal date. We can let $\nu_R(t) = -v_R(\hat{T} - t)^+$ for some $\hat{T} \leq T$ and some (presumably nonnegative) constant v_R (the integrals in the expressions for F_R and g_R are then easily computed).

Next, consider the choice of $\nu_A(t)$ and $\lambda_I(t)$. How binding the constraints are will depend on wealth. Because expected wealth in most models increases up to retirement, we try affine functions

$$\lambda_I(t) = \Lambda_0 + \Lambda_1 t, \quad \nu_A(t) = v_0 + v_1 t.$$

The integrals in the expressions for F_A and g_A can then be computed by standard numerical integration techniques. With these specifications, the strategies and the indirect utility are parameterized by the six constants $\Phi = (v_0, v_1, v_R, \hat{T}, \Lambda_0, \Lambda_1)$, and we denote the associated indirect utility by $J_A(t, x, y; \Phi)$ and the corresponding optimal strategy by $(c(\Phi), \pi_S(\Phi), \pi_I(\Phi))$.

We can now compute an upper bound on the indirect utility in the true constrained market by a minimization over the parameterized artificial markets,

$$\bar{J}(t, x, y) \equiv J_A(t, x, y; \bar{\Phi}) \equiv \min_{\Phi} J_A(t, x, y; \Phi), \quad (14)$$

which is implemented using a standard unconstrained numerical optimization algorithm.

4. A Near-Optimal Strategy in the True Market

We derive a promising candidate for a good consumption–investment strategy in the true constrained market from the optimal strategies in the parameterized family of artificial markets in the following way. For each $\Phi = (v_0, v_1, v_R, \hat{T}, \Lambda_0, \Lambda_1)$, we take the optimal strategy $(c(\Phi), \pi_S(\Phi))$ in the corresponding artificial market—the strategy given by (9) and (10) in retirement and (12) and (13) before retirement—and transform it into a strategy that is feasible in the true market.

We need to make sure that financial wealth stays nonnegative. Intuitively, this liquidity constraint implies that future income has a smaller present value when current financial wealth is small. A parsimonious way to capture this effect is by multiplying the human capital $Y_t F_A(t)$ by a factor $(1 - e^{-\eta X_t})$, where η is a positive constant to be determined. For large financial wealth, the factor is close to one so that human capital is not significantly reduced. When financial wealth approaches zero, the factor approaches zero so that human capital is reduced to zero.⁶ Furthermore,

we have to make sure $\pi_{St} \in [0, 1]$. For strictly positive financial wealth, an obvious candidate for a good strategy in retirement, i.e., for $t \in (\tilde{T}, T]$, is

$$c_t = \frac{X_t + \Upsilon Y_{\tilde{T}} F_R(t)(1 - e^{-\eta X_t})}{g_R(t)}, \quad (15)$$

$$\pi_{St} = \left(\min \left\{ 1, \left[1 + \frac{\Upsilon Y_{\tilde{T}} F_R(t)(1 - e^{-\eta X_t})}{X_t} \cdot \frac{\sigma_S \lambda_S + \nu_R(t)}{\gamma \sigma_S^2} \right]^+ \right\}, \quad (16)$$

where we have suppressed the dependence of F_R and g_R on ν_R . In the active phase, we disregard the investment in the artificial income contract and ensure that constraints are satisfied just as in retirement. For strictly positive financial wealth, the modified strategy for $t \in [0, \tilde{T}]$ is thus

$$c_t = \frac{X_t + Y_t F_A(t)(1 - e^{-\eta X_t})}{g_A(t)}, \quad (17)$$

$$\pi_{St} = \left(\min \left\{ 1, \frac{\sigma_S \lambda_S + \nu_A(t)}{\gamma \sigma_S^2} + \frac{Y_t F_A(t)(1 - e^{-\eta X_t})}{X_t} \cdot \left[\frac{\sigma_S \lambda_S + \nu_A(t)}{\gamma \sigma_S^2} - \frac{\beta \rho}{\sigma_S} \right]^+ \right\}, \quad (18)$$

where we have suppressed the dependence of F_A and g_A on ν_A, ν_R, λ_I . If financial wealth equals zero at any point in time, the investment in the risky asset is restricted to zero, and consumption is a set to fraction of current income, $c_t = k Y_t$, where $k \in (0, 1)$. This ensures that the liquidity constraint is respected.

For any set of the (seven) constants $\Psi = (\Phi, \eta)$, the Equations (15)–(18) define a feasible strategy $(c(\Psi), \pi_S(\Psi))$ in the true market. For any Ψ , we can approximate the expected utility $J(t, x, y; \Psi)$ generated with $(c(\Psi), \pi_S(\Psi))$ by Monte Carlo simulation of the income dynamics (1) and the wealth dynamics (3) substituting in $(c(\Psi), \pi_S(\Psi))$. Searching over different Ψ , we find the best of the feasible strategies $(c(\Psi^*), \pi_S(\Psi^*))$ with an associated expected utility of $J(t, x, y; \Psi^*)$. Again, this can be implemented by a standard unconstrained numerical optimization algorithm.

We can evaluate the performance of any admissible consumption and investment strategy (c, π_S) —including our candidate $(c(\Psi^*), \pi_S(\Psi^*))$ defined above—in the following way. We compare the expected utility generated by the strategy, $J^{c, \pi_S}(t, x, y)$, to the upper bound $\bar{J}(t, x, y)$ on the maximum utility. If the distance is small, the strategy is near optimal. More precisely, we can compute an upper bound $L = L^{c, \pi_S}(t, x, y)$ on the welfare loss suffered when following the specific strategy (c, π_S) by solving the equation

$$J^{c, \pi_S}(t, x, y) = \bar{J}(t, x[1 - L], y[1 - L]).$$

⁶ Other specifications of the factor gave similar results.

Hence, $L^{c, \pi_s}(t, x, y)$ is interpreted as an upper bound on the fraction of total wealth (current wealth as well as current and future income) that the individual would be willing to throw away to get access to the unknown optimal strategy, instead of following the strategy (c, π_s) . If we focus on the active phase, it follows from Theorem 2 that

$$\begin{aligned}\bar{J}(t, x[1-L], y[1-L]) &= J_A(t, x[1-L], y[1-L]; \bar{\Phi}) \\ &= (1-L)^{1-\gamma} J_A(t, x, y; \bar{\Phi}),\end{aligned}$$

so the upper bound on the welfare loss becomes

$$L^{c, \pi_s}(t, x, y) = 1 - \left(\frac{J^{c, \pi_s}(t, x, y)}{J_A(t, x, y; \bar{\Phi})} \right)^{1/(1-\gamma)}.$$

5. Numerical Results

The results presented below are based on simulations using 10,000 paths. Along each path, the consumption and investment strategy is reset 20 times a year (more frequent resetting does not change our results significantly). To reduce any simulation bias in the loss, we also compute the upper bound $J_A(t, x, y; \bar{\Phi})$ by simulation—here it is the wealth dynamics in the artificial market which is simulated—applying the same sequence of random numbers as used in the computation of the utility $J(t, x, y; \psi^*)$ for our best feasible strategy. Table 1 summarizes the benchmark values for the model parameters, which are similar to those used in the existing literature; see Cocco et al. (2005) and Kraft and Munk (2011) and the references therein. The individual has a relative risk aversion of 4, has 30 years until retirement, and subsequently lives for another 20 years. The initial time is $t = 0$, unless mentioned otherwise. Whenever we need to use levels of current wealth, labor income, etc., we use a unit of \$10,000 scaled by one plus the inflation rate in the perishable consumption good. As

Table 1 Benchmark Parameter Values

δ	Time preference rate	0.03
γ	Relative risk aversion	4
\bar{T}	Retirement date	30
T	Terminal date	50
x	Financial wealth	2
r	Risk-free rate	0.02
λ_s	Stock Sharpe ratio	0.25
σ_s	Stock volatility	0.2
y	Annual income	2
α	Expected income growth	0.01
β	Income volatility	0.1
Υ	Replacement ratio	0.6

Notes. This table shows the values of the model parameters used in the numerical computations unless mentioned otherwise. Time is measured in years. The initial wealth $x = 2$ and annual income $y = 2$ are interpreted as \$20,000.

Table 2 Upper Bound on the Welfare Loss as a Percentage of Total Wealth for Different Values of the Stock–Income Correlation and the Wealth–Income Ratio

Wealth–income ratio x/y	Stock–income correlation ρ (%)					
	0	0.2	0.4	0.6	0.8	1
0.1	0.686	0.426	0.327	0.045	0.084	0.135
0.25	0.714	0.435	0.305	0.046	0.078	0.144
1	0.842	0.477	0.182	0.038	0.082	0.100
4	1.110	0.478	0.153	0.039	0.035	0.093
10	0.815	0.351	0.149	0.057	0.019	0.033

the benchmark we put $x = 2$ and $y = 2$, representing an initial financial wealth of \$20,000 and an initial annual income of \$20,000, which are in line with the median net worth and before-tax income statistics derived from the 2007 Survey of Consumer Finances for individuals of age 30–40 (see Bucks et al. 2009, pp. A5, A11); the survey also reveals a large variation in wealth and income across individuals.

5.1. Main Results

First, we consider the size of the upper bound L on the welfare loss from following the strategy derived by our method instead of the unknown optimal strategy throughout the entire life. Consumption and investment strategies are known to depend on the ratio between financial wealth and income, as well as the correlation between stock returns and labor income, so we focus first on the sensitivity of the loss with respect to these quantities. Table 2 shows that for a wide range of values for the initial wealth–income ratio, the welfare loss bound is below 0.5% of current total wealth for an income–stock correlation of 0.2 or higher; in fact, in many cases the welfare loss bound is much lower than 0.5%. The welfare loss bound is somewhat higher for a zero income–stock correlation, but at most 1.1%. These results confirm that our proposed strategy is indeed near optimal.

It is intuitively reasonable that the loss bound is largest for a zero correlation, because in this case the optimal unconstrained strategy will be a highly leveraged position in the stock for many years. Moreover, with zero correlation, the labor income is “far” from being spanned so the true market is very different from the artificial, unconstrained markets. For intermediate correlations, the portfolio constraint on the stock is rarely binding, and the loss bound is very small. For very high correlations, the optimal unconstrained strategy would involve shorting of the stock in the early years, so the loss bound is slightly higher than for intermediate correlations.

Table 3 shows how the upper bound on the welfare loss varies with the initial date and thus with the time to retirement assuming that the initial value of the wealth–income ratio is fixed at 1. As the investment

Table 3 Upper Bound on the Percentage Welfare Loss for Different Time Horizons and Stock–Income Correlations

Initial time t	Stock–income correlation ρ (%)					
	0	0.2	0.4	0.6	0.8	1
0	0.842	0.477	0.182	0.038	0.082	0.100
10	0.645	0.415	0.191	0.029	0.085	0.120
20	0.467	0.377	0.265	0.228	0.247	0.289
25	0.801	0.795	0.781	0.749	0.804	0.844

Note. The wealth–income ratio at time t is fixed at 1.

horizon decreases, human wealth decreases. This reduces the wealth effect of labor income on the stock investment as well as the adjustment for stock-like income risk. Because the latter reduction depends on the correlation, the net effect of the decrease in the investment horizon also depends on correlation. For low correlations, this implies that the no borrowing constraint is less tight and thus the loss bound tends to decrease. For higher correlations, constraints may become more frequently active over shorter horizons, and the loss bound may increase. As the initial date is moved close to the retirement date, the loss bound begins to increase for any correlation value.

5.2. A Comparison with Two Alternative Numerical Methods

We have also solved the utility maximization problem (4) with the so-called MCA method, which is a well-studied and frequently applied numerical approach (Kushner and Dupuis 2001; Munk 2000, 2003). The homogeneity property of power utility implies that the indirect utility function can be written as $J(t, x, y) = y^{1-\gamma}H(t, x/y)$. From the HJB equation for J , a nonlinear second-order partial differential equation (PDE) for H can be derived, and this PDE is the HJB equation for another stochastic control problem where the controls are simple scalings of

the original consumption and portfolio plans. The MCA method discretizes this control problem. The dynamics of the wealth–income ratio are approximated by a Markov chain on a grid defined by N equidistant time points and I equidistant values of the wealth–income ratio. In the continuous-time model, the wealth–income ratio is unbounded from above, but the MCA method has to impose an upper bound. The optimization problem is solved by backward recursion starting at the terminal date T . In each time step the value function for each state in the grid is maximized by policy iterations. The entire procedure is roughly equivalent to solving the HJB equation for H by a (specific) finite difference approach, similar to the one used by Brennan et al. (1997) and others. The precision of the method depends heavily on the number of grid points and the size of the imposed upper bound. Ideally, the bound should be set so high that it is very unlikely that the wealth–income ratio would exceed that bound when the optimal strategies are followed. This can be checked by simulating the wealth–income ratio using the strategies obtained with the method for a given upper bound. If sufficiently many paths exceed the bound, the MCA method must be rerun with a higher imposed bound. This complicates the application of the MCA method as well as other grid-based methods.

We have solved our problem with the MCA method both for a very fine grid ($N = 4,000$, $I = 12,000$) and a coarser—but still quite fine—grid ($N = 2,000$, $I = 4,000$). We evaluate the expected utility of the consumption and investment strategy derived with the MCA method by simulation using the same random numbers as in the valuation of our suggested strategy, to avoid any bias stemming from the simulations. Table 4 presents the upper bounds on the welfare loss for the MCA method for both grids and compares with the upper bound for our SAMS method.

Table 4 Upper Bounds on Welfare Loss for Two Different Methods

Method	Run time (in sec.)	Stock–income correlation ρ (%)					
		0	0.2	0.4	0.6	0.8	1
SAMS	600	0.842	0.477	0.182	0.038	0.082	0.100
MCA, fine grid	42	0.728	0.454	0.185	0.038	0.080	0.068
Loss difference		0.115	0.023	−0.003	−0.000	0.003	0.032
MCA, coarse grid	8	0.879	0.621	0.212	0.076	0.105	0.246
Loss difference		−0.037	−0.144	−0.030	−0.038	−0.023	−0.147
Koijen et al. (2007, 2010)	3,200	0.991	0.606	0.382	0.302	0.302	0.173
Loss difference		−0.148	−0.128	−0.200	−0.264	−0.220	−0.074
SAMS parsimon	<1	1.143	0.608	0.186	0.039	0.101	0.141
Loss difference		−0.301	−0.131	−0.004	−0.000	−0.019	−0.042

Notes. This table shows upper bounds on the percentage welfare loss for our SAMS method and two alternative methods as well as the increase in the upper bound (the loss difference) by applying our SAMS method instead of the alternative method. We assume the benchmark parameter values as well as $t = 0$ and $x/y = 1$. All computations were run on an Intel Quad Core with 2.33 GHz and 4 GB of random access memory. The run time is slightly dependent on the assumed correlation ρ , and the reported run time is a rough average over the run times for the different correlations.

The loss difference is simply the loss bound for our method minus the loss bound for the MCA method, i.e., the increase in the upper bound by applying our SAMS method instead of the MCA method. The table assumes an initial wealth–income ratio of 1, but the results are similar for other values. The table shows that our method is roughly as good as the MCA method with the very fine grid in the sense that the additional loss from following our strategy instead of the MCA-based strategy is at most 0.115% of wealth. With the very fine grid, the derived strategy can be expected to be very close to the truly optimal strategy, so the results suggest that the upper bound on the welfare loss is not very tight. Furthermore, our method beats the MCA method with the coarser grid.

For our problem with one state variable, the computation time for the MCA is less than for our SAMS method (excluding time used to search for a reasonable upper grid bound for MCA). However, for problems with two or three state variables, fine grids are intractable. For example, a grid with two state variables and 12,000 gridpoints per variable, as we have used, would have 144 million gridpoints at each point in time considered, requiring a lot of computer memory and leading to long computation times. The above results suggest that our SAMS approach would outperform tractable implementations of the MCA method for problems with two or three state variables. For problems with more than three state variables (after any homogeneity is exploited), all grid-based methods seem computationally infeasible. In contrast, our SAMS method is still applicable.

The SAMS approach involves a simulation-based maximization over the parameters of the family of feasible strategies. However, SAMS delivers a very good feasible strategy without the time-consuming maximization: The minimization in (14) that delivers the lowest upper bound on utility is obtained for a parameter set $\bar{\phi}$. The corresponding strategy $(c(\bar{\phi}), \pi_s(\bar{\phi}))$ can be pruned as described in §4 and then evaluated by simulation. As can be seen in the row labeled “SAMS parsimon” in Table 4, this strategy can be evaluated extremely fast and it performs very well.

We have also implemented the method by Kojien et al. (2007, 2010) on a discrete-time version of our benchmark problem. Their method involves the construction of an endogenous grid of the wealth–income ratio (after consumption), which is the only state variable for our problem because of homogeneity of the utility function. Paths of asset returns and income are simulated. The optimal strategy on the grid is computed using a backward dynamic programming algorithm. Our implementation has 20 time steps per year and a grid with 101 values of the state variable, and 10,000 paths are simulated. Table 4 shows

that their method is slightly less precise than our method, and the computational time is significantly higher. The optimal strategies derived by our method and their method are very close, whereas the strategy derived by the MCA approach deviates somewhat for high values of the wealth–income ratio. For models with additional state variables, the complexity and computation time of the Kojien et al. (2007, 2010) method grow considerably because it then involves simulation-based regressions for the approximation of conditional expectations. Furthermore, we emphasize again that our method relies on closed-form strategies unlike the alternative methods.

5.3. Detailed Results from Our Method

Next, we investigate the auxiliary parameters Ψ^* underlying the best of the strategies of the form in (15)–(18). Table 5 shows the optimal auxiliary parameters, corresponding to Ψ^* , for different stock–income correlations and for an initial wealth–income ratio of 1. As for any multidimensional numerical optimization, some experimentation with starting values, possible sequential optimization over different subsets of parameters, and so on is recommended. With the parameter values and initial state variables listed in Table 1, our experiments show that best results are obtained with $v_R = 0$ (then \hat{T} is meaningless) and $\eta = 30$, and then running a simulated annealing optimization routine to find the optimal remaining parameters displayed in Table 5.⁷ The values of Λ_0 and Λ_1 indicate that the risk premium that the individual associates with the unspanned income risk is positive and decreasing over life, and the risk premium is higher for low correlations. A high risk premium translates into a low value of the income multiplier $F_A(t)$, i.e., a low human wealth.

For a zero correlation, the optimized $\nu_{At} = v_0 + v_1 t$ is negative over most of the working life. This indicates that the risk-free rate is artificially increased to make the stock less attractive, because the unrestricted fraction of wealth invested in the stock would exceed one. For very high correlations, the optimized ν_{At} is positive, artificially increasing the expected return on the stock, which makes good sense early in life where the unrestricted stock investment would be negative in this case. However, Λ_0 and Λ_1 also affect the stock investment via the income multiplier $F_A(t)$, so that it is difficult to interpret the values of v_0 and v_1 .

In general, we would optimize over the auxiliary parameters in Ψ for each value of the wealth–income ratio x/y of interest. However, the optimal

⁷ With our parameters, financial wealth will be sufficiently big at retirement that the portfolio constraints are not binding, so it is natural to have $\nu_R(t) = 0$. The large value of η indicates that downward scaling of human wealth imposed in (17) and (18) is only significant for very small levels of financial wealth.

Table 5 Optimal Parameters for the Incomplete Market for $x/y = 1$ and $t = 0$

Parameter	Stock–income correlation ρ					
	0	0.2	0.4	0.6	0.8	1
Λ_0	0.43578	0.40760	0.40764	0.28137	0.28139	0.00799
Λ_1	−0.00261	−0.00143	−0.00141	−0.00067	−0.00060	−0.00027
v_0	−0.01559	0.00070	0.00074	0.00008	0.00014	0.00009
v_1	0.00064	−0.00002	0.00006	0.00002	0.00005	0.00015

Note. In addition, $v_R = 0$ (so that \hat{T} is meaningless) and $\eta = 30$.

parameters vary only relatively little with x/y and, for a fixed x/y , the expected utility $J(t, x, y; \psi)$ is relatively insensitive to ψ around the optimal value ψ^* . We have performed the following experiment. For any given correlation, we find the optimal parameters for $x/y = 1$ and then apply the same parameters for the other values of x/y considered. Obviously, applying the strategy based on the nonoptimal parameters leads to a higher welfare loss bound. Table 6 documents that the increase in the percentage welfare loss caused by not reoptimizing over the parameters when x/y is different from 1 is very small. In particular, for a local sensitivity analysis of the near-optimal strategy with respect to changes in x or y , it is fair to keep the parameters fixed.

5.4. A Comparison with Ad Hoc Strategies

We briefly compare with some ad hoc strategies also evaluated by Monte Carlo simulation. As before, the individual starts at age 30 with a wealth–income ratio of 1. One popular recommendation is to let the percentage of financial wealth invested in the stock market be given by “100 minus age.” If we combine this with the consumption strategy derived by the SAMS method (using the wealth level generated by the ad hoc portfolio rule), the welfare loss is 2.90% (0.68%) higher than for the fully SAMS-generated strategy if the stock–income correlation is 0.0 (0.6). The extra loss caused by the ad hoc strategy is due both to the deviation from the optimal

average strategy (see Figure 1) and the lack of state dependence. If we further replace the consumption strategy before retirement by the ad hoc rule that consumption equals 90% of current income, then the additional welfare losses increase to 9.17% (correlation 0.0) and 3.53% (correlation 0.6). We have also implemented the more aggressive strategy of having “130 minus age” in stocks, which seems to gain popularity. Combined with our near-optimal consumption strategy, this portfolio rule generates welfare losses that are 0.91% (correlation 0.0) and 1.26% (correlation 0.6) higher than our strategy. With the same ad hoc consumption rule as above, the additional losses increase to 8.89% and 4.21%, respectively.

6. Some Properties of the Solution

Although the focus of this paper is on the solution technique, we take a brief look at some key properties of our solution. First, we derived our near-optimal consumption and investment strategy as a function of age, financial wealth, and labor income. Then we simulated 10,000 paths of financial wealth using this strategy and labor income over the life of the individual. Figure 1 shows a random sample path of the fraction of wealth invested in the stock as well as the average, the 5th percentile, and the 95th percentile across the 10,000 paths, for income–stock correlations ρ of 0.0 (left panel) and 0.6 (right panel). For a zero or low stock–income correlation, the individual will typically invest the entire financial wealth in the stock in the early years and then gradually, as human wealth decreases, replace a fraction of the stock investment by a risk-free investment. For the medium–high correlation of 0.6, the income resembles a stock more so the optimal portfolio is initially balanced between the stock and the risk-free asset. In fact, for a very high correlation the individual would initially like to short the stock, so the optimal constrained investment strategy is to invest the entire financial wealth in the risk-free asset. Later, as human wealth decreases, the stock is included in the portfolio.

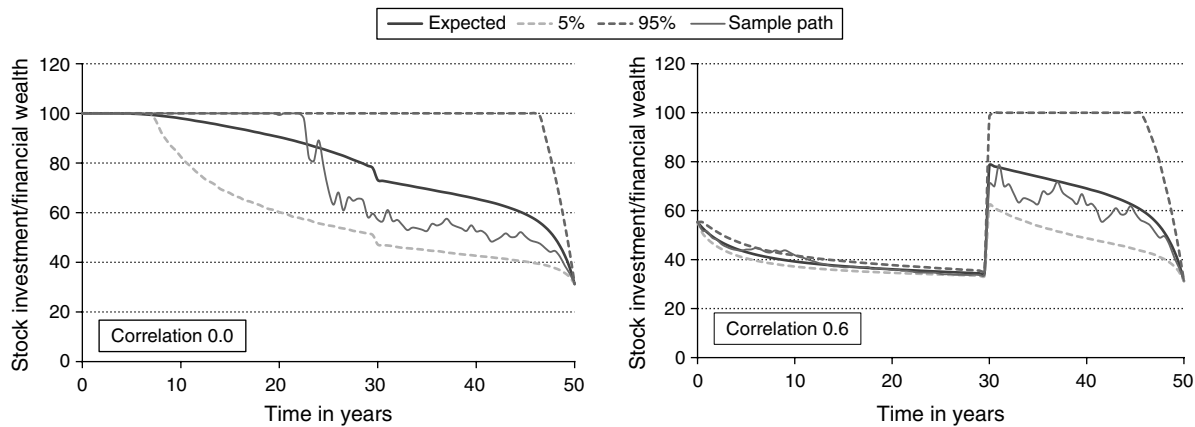
At retirement, the optimal asset allocation changes dramatically because the income risk is suddenly resolved. After retirement, both human wealth and

Table 6 Increase in Percentage Welfare Loss Bound from Applying the Strategy Based on the Parameter ψ^* Optimal for $x/y = 1$ for Other Values of x/y

Wealth–income ratio x/y	Stock–income correlation ρ (%)					
	0	0.2	0.4	0.6	0.8	1
0.1	0.08	0.04	0.00	0.00	0.00	0.01
0.25	0.04	0.03	0.01	0.01	0.00	0.01
1	0	0	0	0	0	0
4	0.15	0.02	0.00	0.01	0.00	0.01
10	0.37	0.05	0.01	0.00	0.00	0.02

Notes. For instance, the 0.08% reported for $x/y = 0.1$ and $\rho = 0$ means that the upper bound on the welfare loss increases by 0.08 percentage points when using the auxiliary parameters found optimal for $x/y = 1$ in the consumption and investment strategy for $x/y = 0.1$ instead of the auxiliary parameters that are indeed optimal for $x/y = 0.1$.

Figure 1 Stock Weight over the Life Cycle



Notes. For each point in time the graph shows the average, the 5th percentile, and the 95th percentile of the fraction of financial wealth invested in the stock across all simulated paths as well as a randomly chosen sample path. The stock–income correlation is 0.0 in the graph to the left and 0.6 in the graph to the right. The initial wealth–income ratio is assumed to be $x/y = 1$.

the optimal fraction of total wealth invested in the stock are independent of what the stock–income correlation was before retirement. However, the financial wealth built up during the active phase will depend on the correlation. On average, individuals with a high stock–income correlation enter retirement with a low financial wealth, because they have been investing little in the stock compared to individuals with a low correlation. For an individual with a high correlation, financial wealth constitutes a lower fraction of total wealth at retirement, and to obtain the desired risk exposure of total wealth, this individual will have to invest a higher fraction of financial wealth in the stock, as shown in the graph.

Figure 2 depicts how the optimal fraction of financial wealth invested in the stock at time $t = 0$ depends on the initial ratio of financial wealth to annual income. As the wealth–income ratio approaches infinity, income becomes irrelevant so the investor should optimally invest the fraction $\lambda_S/(\gamma\sigma_S) = 0.3125$ of

financial wealth in the stock, as in the no-income setting of Merton (1971). For a zero or low stock–income correlation, labor income mainly influences investments through the addition of human wealth to financial wealth. When the wealth–income ratio is small, the entire (but small) financial wealth is therefore invested in the stock in order to obtain the best possible overall risk exposure. As the wealth–income ratio is increased, the optimal fraction of financial wealth invested in the stock will eventually fall below 1 and gradually decrease toward the asymptotic value of 0.3125. Conversely, for a high correlation, labor income is much like an implicit stock investment. The optimal fraction of financial wealth invested in the stock will therefore be zero for low wealth–income ratios, but become positive for a high enough wealth–income ratio, and eventually approach the asymptotic 0.3125. These results demonstrate that the sensitivity of the optimal stock investment to the initial wealth and income is highly dependent on the risk characteristics of labor income.

Because of unspanned income risk (unless $|\rho| = 1$ or $\beta = 0$) and portfolio constraints, there is no unique, market-set value of the future income stream. The human wealth will depend on the stock–income correlation and on the risk aversion and wealth of the individual. We define human wealth $H = H(t, x, y)$ as the minimum extra financial wealth that the individual would need as compensation if the entire income stream is taken away, i.e.,

$$J(t, x, y) = J(t, x + H, 0),$$

where the left-hand side is the indirect utility with the income stream, and the right-hand side is the indirect utility without the income stream but a higher financial wealth. Given our benchmark parameter values,

Figure 2 The Initial Fraction of Financial Wealth Invested in the Stock as a Function of the Initial Wealth–Income Ratio for Different Values of the Stock–Income Correlation

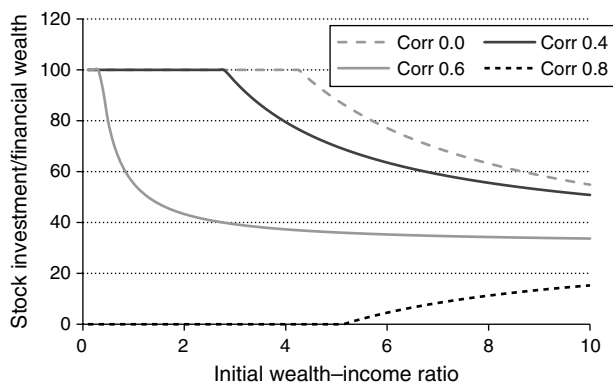


Table 7 Income Multipliers at Time $t = 0$ for Different Wealth–Income Ratios x/y and Different Stock–Income Correlations ρ

Wealth–income ratio x/y	Stock–income correlation ρ					
	0	0.2	0.4	0.6	0.8	1
0.1	23.06	22.49	21.96	21.57	21.48	21.53
0.25	23.16	22.56	22.00	21.59	21.50	21.54
1	23.60	22.87	22.19	21.72	21.59	21.64
4	24.82	23.70	22.74	22.13	21.86	21.86
10	26.12	24.65	23.52	22.70	22.19	22.02

the portfolio constraints are not binding in the case without income, so the right-hand side equals

$$J(t, x + H, 0) = \frac{1}{1 - \gamma} (g^{\text{com}}(t))^{\gamma} (x + H)^{1-\gamma},$$

where g^{com} is defined in (7). Hence, we can compute H as

$$H = [(1 - \gamma)J(t, x, y)]^{1/(1-\gamma)} (g^{\text{com}}(t))^{\gamma/(\gamma-1)} - x.$$

We can interpret H/y as an income multiplier, because this is the factor that current income has to be multiplied by to get the human wealth. We replace the unknown indirect utility $J(t, x, y)$ with the expected utility generated by our near-optimal strategy.

Table 7 reports the income multiplier H/y for different combinations of the stock–income correlation and the initial wealth–income ratio. The income multiplier is increasing in the wealth–income ratio, because with relatively high financial wealth the individual is less concerned with unspanned income risk, and the portfolio constraints will rarely bind, especially for low correlations. For example, with zero correlation, the multiplier is 13.3% larger starting with a high wealth ($x/y = 10$) than a low wealth ($x/y = 0.1$). The income

multiplier is decreasing in the stock–income correlation (except close to perfect correlation, which is a very special case). The lower the correlation, the better the inherent income risk hedging properties of a positive investment in the risky asset and, thus, the more valuable the income stream.

7. Comparative Statics

As a robustness check we vary selected parameters one by one. We focus on the relative risk aversion and the parameters driving the income process, and for each parameter we consider a value below and a value above the benchmark value. Table 8 reports both the upper loss bounds for our method and the increase in the loss bound relative to the MCA approach implemented with the fine grid. Overall, the welfare loss bound remains small for all the considered parameter values and the two methods provide very similar results. For 7 of the 24 parameter combinations considered in the table, our method outperforms the fine grid MCA approach. Compared to the MCA with the coarser grid, our method does better for 16 of the 24 parameter combinations considered in the table (results available upon request).

The loss bound is somewhat higher for a low risk aversion than for a high risk aversion. For a lower risk aversion, the unrestricted speculative stock demand will be more sensitive to the human capital and will stay above the imposed maximum of 100% for a longer period of time, so the imposed portfolio constraints are more binding for $\gamma = 3$ than for $\gamma = 5$ and a zero or moderate correlation. It is then not surprising that the welfare loss bound is higher for $\gamma = 3$ than for $\gamma = 5$.

The welfare loss bound tends to increase with the riskiness of the income stream measured by its volatility β , which makes sense as the unspanned income risk is then bigger. Note that for the case

Table 8 Parameter Sensitivity of the Welfare Loss Bound

Parameter	Stock–income correlation ρ (%)					
	0		0.4		0.8	
	Bound	MCA diff.	Bound	MCA diff.	Bound	MCA diff.
$\gamma = 3$	0.938	0.168	0.414	0.038	0.102	0.006
$\gamma = 5$	0.813	0.061	0.044	−0.022	0.062	0.020
$\alpha = 0.005$	0.887	0.097	0.155	−0.004	0.098	−0.000
$\alpha = 0.015$	0.783	0.125	0.242	−0.000	0.067	0.000
$\beta = 0.05$	0.624	0.156	0.383	0.068	0.224	0.027
$\beta = 0.15$	1.500	−0.357	0.266	−0.023	0.457	0.042
$\Upsilon = 0.5$	0.809	0.093	0.151	−0.005	0.086	0.001
$\Upsilon = 0.7$	0.896	0.139	0.245	0.006	0.102	0.003

Notes. For each correlation value, the left column shows the upper bound on percentage welfare losses when the parameter value is changed from its benchmark as indicated. The right column shows the additional welfare loss from applying the consumption and investment strategy suggested by our method instead of the strategy suggested by the Markov chain approximation method with a fine grid ($N = 4,000, I = 12,000$).

of high income volatility and zero income–stock correlation, in which the upper loss bound is highest (1.5%), our method performs significantly better than the MCA approach.

The expected income growth rate α enters the optimal strategies only via the F_A -function, which is increasing in α . As indicated by (11), variations in α that make constraints more or less binding are easily mitigated by varying ν_A and λ_I . Consequently, after optimizing over these parameters, the welfare loss bound is relatively insensitive to α .

Finally, the welfare loss is slightly increasing in the income replacement ratio Υ , because this increases human capital and thus tends to make portfolio constraints more binding early in life.

8. Extension to Stochastic Interest Rates

Until now we have assumed a simple Black–Scholes-type financial market, but our approach applies to more general settings. As an example we consider the case where interest rates are stochastic as described by the Vasicek (1977) model so that the short-term interest rate r_t has dynamics

$$dr_t = \kappa[\bar{r} - r_t]dt - \sigma_r dW_{rt},$$

where \bar{r} , κ , and σ_r are constants, and W_r is a standard Brownian motion. The price B_t of any bond has dynamics of the form

$$dB_t = B_t[(r_t + \lambda_B \sigma_B(r_t, t))dt + \sigma_B(r_t, t)dW_{rt}],$$

where λ_B is a constant market price of interest rate risk. For a zero-coupon bond with a time to maturity of τ , the price is of the form $B_t = \exp\{-\mathcal{A}(\tau) - \mathcal{B}_\kappa(\tau)r_t\}$, so that $\sigma_B(r_t, t) = \sigma_r \mathcal{B}_\kappa(\tau)$. Here $\mathcal{B}_m(\tau) = (1 - e^{-m\tau})/m$ for any constant m , and \mathcal{A} is a deterministic function of minor importance for what follows. The dynamics of the stock price and the labor income is now assumed to be

$$dS_t = S_t[(r_t + \lambda_S \sigma_S)dt + \sigma_S(\rho_{SB} dW_{rt} + \hat{\rho}_S dW_{St})],$$

$$dY_t = Y_t[\alpha dt + \beta(\rho_{YB} dW_{rt} + \hat{\rho}_{YS} dW_{St} + \hat{\rho}_Y dW_{Yt})], \quad t < \tilde{T},$$

where W_r , W_S , and W_Y are independent standard Brownian motions and

$$\hat{\rho}_S = \sqrt{1 - \rho_{SB}^2}, \quad \hat{\rho}_{YS} = \frac{\rho_{YS} - \rho_{SB}\rho_{YB}}{\sqrt{1 - \rho_{SB}^2}},$$

$$\hat{\rho}_Y = \sqrt{1 - \rho_{YB}^2 - \hat{\rho}_{YS}^2},$$

where ρ_{SB} , ρ_{YB} , and ρ_{YS} are the pairwise stock–bond, income–bond, and income–stock correlations. In retirement, the income is again given by (2).⁸

The individual can trade in the stock, the instantaneously risk-free bank account, and a single bond

index. The bond index is continuously rebalanced so that at any point in time it corresponds to a zero-coupon bond having a time to maturity of $\bar{\tau}$ and, consequently, a constant volatility of $\sigma_B = \sigma_r \mathcal{B}_\kappa(\bar{\tau})$. Let π_{St} and π_{Bt} denote the fractions of financial wealth invested in the stock and the bond index, respectively, at time t . The remaining financial wealth $W_t(1 - \pi_{St} - \pi_{Bt})$ is invested in the bank account. We impose the constraints that $\pi_{St}, \pi_{Bt} \in [0, 1]$ and $\pi_{St} + \pi_{Bt} \leq 1$ (borrowing prohibited), and before retirement we also have to make sure that financial wealth stays nonnegative as in the problem studied in the preceding sections.

An artificial market corresponding to this constrained, incomplete market is characterized by a triple $(\nu_S, \nu_B, \lambda_I)$ of stochastic processes such that (Cvitanic and Karatzas 1992)

1. the short-term interest rate is $\tilde{r}_t = r_t + \max(\nu_{Bt}^-, \nu_{St}^-)$,
2. the drift of the stock price is $\tilde{r}_t + \sigma_S \lambda_S + \nu_{St}$,
3. the drift of the bond price is $\tilde{r}_t + \sigma_B \lambda_B + \nu_{Bt}$, and
4. until retirement the individual can trade in an income contract with price dynamics (8), that is, with Sharpe ratio λ_{It} .

For artificial markets associated with *deterministic* processes $(\nu_S, \nu_B, \lambda_I)$, the unconstrained utility maximization problem is solved in closed form by extending results of Liu (2007) and Munk and Sørensen (2010) (see the online appendix at <https://sites.google.com/site/munkfinance/publications> for details). We specialize again to the simple functions

$$\nu_S(t) = \begin{cases} v_{S0} + v_{S1}t & \text{for } t \in [0, \tilde{T}], \\ -v_{SR}(\hat{T}_S - t)^+ & \text{for } t \in [\tilde{T}, T], \end{cases}$$

$$\nu_B(t) = \begin{cases} v_{B0} + v_{B1}t & \text{for } t \in [0, \tilde{T}], \\ -v_{BR}(\hat{T}_B - t)^+ & \text{for } t \in [\tilde{T}, T], \end{cases}$$

where $\hat{T}_S < T$ and $\hat{T}_B < T$, and $\lambda_I(t) = \Lambda_0 + \Lambda_1 t$. We minimize $J_A(t, x, y, r; \nu_S, \nu_B, \lambda_I)$ over all the simple artificial markets to find the upper bound on the obtainable utility in the true market. The optimal strategy in any simple artificial market is transformed into a feasible strategy in the true market by disregarding the income contract, pruning the investments in the bond and the stock to comply with constraints,⁹ and multiplying human wealth by $(1 - e^{-\eta X_t})$ to ensure nonnegative financial wealth before retirement. We can evaluate each of these feasible strategies by Monte Carlo simulation and build this evaluation into a standard optimization routine.

and Viceira (2001) without labor income and Kojien et al. (2010), Van Hemert (2010), Munk and Sørensen (2010), and Kraft and Munk (2011) with labor income.

⁸ Dynamic portfolio choice with Vasicek-type interest rates has been studied by Sørensen (1999), Brennan and Xia (2000), and Campbell

⁹ If π_S and π_B are both positive and their sum is above 1, we divide both of them by the sum.

Table 9 Additional Parameter Values with Stochastic Interest Rates

κ	Mean reversion speed	0.2
\bar{r}	Long-run short rate	0.02
σ_r	Short rate volatility	0.01
λ_B	Bond Sharpe ratio	0.1
$\bar{\tau}$	Bond maturity	50
ρ_{SB}	Stock–bond correlation	0.1
ρ_{VB}	Income–bond correlation	0.25

Note. This table shows the values of the additional parameters in the model with stochastic interest rates.

We assume the values of the interest rate parameters listed in Table 9, whereas the other parameters are still the same as in Table 1. In particular, we consider a long-duration bond index with a volatility of $\sigma_B \approx \sigma_r/\kappa = 5\%$ and an excess expected return of $\sigma_B \lambda_B \approx 0.5\%$. We assume that the initial value of the short-term interest rate equals the long-run average of 2%.

The upper bound on the welfare loss associated with the strategy derived by our method is shown in Table 10 for different values of the stock–income correlation ρ_{YS} and the ratio x/y between initial financial wealth and initial annual income. The loss bound is at most 1.4% and often much lower (results are similar for other parameter combinations). The results for the model with constant interest rates suggest that the bound is not very tight so that actual losses are much smaller.

To indicate that our solution makes economical sense, Figure 3 shows the average (over 10,000 paths) optimal allocation to the stock, the bond index, and the bank account over the life cycle. The left panel is for a zero income–stock correlation, whereas the right panel is for a correlation of 0.8. With zero correlation, almost the entire financial wealth is invested in the stock early in life. The long-term bond would be useful to hedge interest rate risk, but, on the other hand, the bond has a much smaller Sharpe ratio than

Table 10 Upper Bound on the Percentage Welfare Loss in the Case with Stochastic Interest Rates for Different Values of the Wealth–Income Ratio and the Stock–Income Correlation

Wealth–income ratio x/y	Stock–income correlation ρ_{YS} (%)				
	0	0.2	0.4	0.6	0.8
0.1	1.094	0.958	0.679	0.400	0.175
0.25	1.113	0.963	0.629	0.420	0.190
1	1.261	1.151	0.739	0.350	0.236
4	1.431	0.652	0.428	0.425	0.451
10	0.715	0.327	0.180	0.149	0.525

the stock, and the bond is positively correlated with income. In retirement, the portfolio is still dominated by the stock, but now a significant fraction of financial wealth should be invested in the bond because of its hedging property. When the income is highly correlated with the stock, all financial wealth is invested in the bond through most of working life. In retirement, income is risk-free so the stock becomes attractive. Note again that the income–stock correlation before retirement affects the optimal portfolio in retirement. With a high income–stock correlation, financial wealth at retirement tends to be lower so that a bigger share of that wealth has to be invested in the stock to obtain the desired overall risk-return balance in retirement.


9. Applications to More General Problems

In this section we briefly outline how to apply our method in a general affine market setting; we ignore some technical conditions for brevity. Suppose that the individual potentially can trade in a bank account with rate $r(z_t, t)$ and up to d risky assets with price vector P_t satisfying

$$dP_t = \text{diag}(P_t)[(r(z_t, t)\mathbf{1}_d + \sigma_t \lambda(z_t, t)) dt + \sigma_t dW_t],$$

where W is a d -dimensional standard Brownian motion, and σ_t is a nonsingular $d \times d$ matrix. The

Figure 3 Optimal Portfolios over the Life Cycle



Notes. The graphs show the average portfolio over the life cycle when the income–stock correlation ρ_{YS} is 0 (left panel) or 0.8 (right panel). The optimal strategies are computed with our numerical method explained in the text. Ten thousand paths of income, interest rates, and wealth (applying those strategies) are then simulated over the 50-year period considered. The graphs show averages over the paths.

process z represents a state variable of dimension $k \leq d$ with dynamics

$$dz_t = m(z_t, t) dt + \Gamma(z_t, t) dW_t.$$

We take an affine market structure¹⁰ so that $r(z, t) = r_0(t) + r_1(t)^\top z$, $m(z, t) = m_0(t) + m_1(t)z$, $\|\lambda(z, t)\|^2 = \Lambda_0(t) + \Lambda_1(t)z$, and $\Gamma(z, t) = D\sqrt{v(z, t)}$ where D is a $k \times d$ constant matrix, $v(z, t)$ is a diagonal $d \times d$ matrix with $[v(z, t)]_{ii} = v_{0i}(t) + v_{1i}(t)^\top z$, and moreover, $v(z, t)\lambda(z, t) = K_0(t) + K_1(t)z$.

Let π_t be the vector of fractions of financial wealth invested in the d risky assets. Portfolios are constrained to $\pi_t \in \mathcal{K}$, where \mathcal{K} is a closed, convex subset of \mathbb{R}^d . As shown by Cvitanić and Karatzas (1992), this captures typical portfolio constraints including (i) nontraded assets where $\mathcal{K} = \mathbb{R}^n \times \{0\}^{d-n}$, (ii) short-sale constraints where $\mathcal{K} = \mathbb{R}^n \times \mathbb{R}_+^{d-n}$, and (iii) short-sale and borrowing constraint where $\mathcal{K} = \{\pi \in \mathbb{R}_+^d : \pi^\top \mathbf{1}_d \leq 1\}$, as well as combinations of different constraints.

The individual has an income rate Y_t , which before retirement evolves as

$$dY_t = Y_t[\alpha(t) dt + \beta(t)(\rho^\top dW_t + \sqrt{1 - \|\rho\|^2} dW_{Y_t})], \quad 0 \leq t \leq \tilde{T},$$

where W_Y is a standard Brownian motion independent of W so that ρ is the vector of correlations between the income and the risky asset prices. We assume α and β are deterministic as in Cocco et al. (2005), although some state dependence in α could be handled as in Munk and Sørensen (2010).

An artificial, unconstrained market is characterized by a pair of processes (ν, λ_I) . Here ν is valued in $\tilde{\mathcal{K}}$, which is the barrier cone of $-\mathcal{K}$, that is, the set of $\zeta \in \mathbb{R}^d$ for which the support function

$$\Delta(\zeta) = \sup_{\pi \in \mathcal{K}} (-\pi^\top \zeta)$$

is finite. For the three examples of \mathcal{K} above, we have (i) $\tilde{\mathcal{K}} = \{0\}^n \times \mathbb{R}^{d-n}$ and $\Delta(\zeta) = 0$ on $\tilde{\mathcal{K}}$, (ii) $\tilde{\mathcal{K}} = \{0\}^n \times \mathbb{R}_+^{d-n}$ and $\Delta(\zeta) = 0$ on $\tilde{\mathcal{K}}$, and (iii) $\tilde{\mathcal{K}} = \mathbb{R}^d$ and $\Delta(\zeta) = \max\{\zeta_1^-, \dots, \zeta_d^-\}$ on $\tilde{\mathcal{K}}$.

In the artificial market corresponding to (ν, λ_I) , the risk-free rate is $\tilde{r}_t = r_t + \Delta(\nu_t)$, and the expected rates of return on the d risky assets are $\tilde{r}_t + \sigma_t \lambda_t + \nu_t$. As before, let π_{It} be the fraction of wealth invested in an “income contract” with Sharpe ratio λ_I as in (8). The

wealth dynamics of a consumption and investment strategy (c, π, π_I) in this artificial market is

$$\begin{aligned} dX_t &= X_t[(\tilde{r}_t + \pi_t^\top[\sigma_t \lambda_t + \nu_t] + \pi_{It} \lambda_{It}) dt \\ &\quad + \pi_t^\top \sigma_t dW_t + \pi_{It} dW_{Y_t}] + (y_t - c_t) dt \\ &= X_t[(r_t + \pi_t^\top \sigma_t \lambda_t) dt + \pi_t^\top \sigma_t dW_t] + (y_t - c_t) dt \\ &\quad + \pi_{It}[\lambda_{It} dt + dW_{Y_t}] + X_t[\Delta(\nu_t) + \pi_t^\top \nu_t] dt. \end{aligned}$$

Note that $\Delta(\nu_t) + \pi_t^\top \nu_t \geq 0$ when $\pi_t \in \mathcal{K}$ and $\nu_t \in \tilde{\mathcal{K}}$. Hence, any strategy that is feasible in the true, constrained market will, for the same consumption strategy, lead to at least the same terminal wealth in each of the artificial markets as in the true market, and thus at least the same expected utility. Because many other strategies are feasible in the artificial markets, the indirect utility in each artificial market will be as at least as large as in the true market.

In an artificial market where ν, λ_I are simple deterministic functions, the indirect utility is

$$J(t, x, y, z; \nu, \lambda_I) = \frac{1}{1 - \gamma} (e^{G_0(t; \nu, \lambda_I) + G_1(t; \nu, \lambda_I)^\top z})^\gamma \cdot (x + yF(t, z; \nu, \lambda_I))^{1 - \gamma},$$

where the functions G_0, G_1 solve a system of differential equations, which in some models can be solved explicitly and in others by standard numerical methods. Furthermore, F takes the form

$$F_R(t, z; \nu) = \int_t^T e^{-A(t,s) - B(t,s)^\top z - \int_t^s \Delta(\nu(u)) du} ds$$

in retirement, whereas before retirement

$$\begin{aligned} F_A(t, z; \nu, \lambda_I) &= \Upsilon \int_{\tilde{T}}^T e^{-\tilde{A}(t,s) - \tilde{B}(t,s)^\top z - \int_t^s \Delta(\nu(u)) du} ds \\ &\quad + \int_t^T e^{-\tilde{A}(t,s) - \tilde{B}(t,s)^\top z - \int_t^s \Delta(\nu(u)) du} ds. \end{aligned}$$

The functions A, B , etc., also solve systems of differential equations that can be solved numerically and sometimes explicitly. Before retirement, the optimal strategy in this artificial market is

$$\begin{aligned} c_t &= (e^{-G_0(t; \nu, \lambda_I) - G_1(t; \nu, \lambda_I)^\top z}) (X_t + Y_t F(t, z_t; \nu, \lambda_I)), \\ \pi_t &= (\sigma_t^\top)^{-1} \left[\frac{1}{\gamma} (\lambda(z_t, t) + \sigma_t^{-1} \nu(t)) \right. \\ &\quad \left. - \frac{\gamma - 1}{\gamma} \sqrt{v(z_t, t)} D^\top G_1(t; \nu, \lambda_I) \right] \\ &\quad \cdot \frac{X_t + Y_t F_A(t, z_t)}{X_t} - \sigma_t^{-1} \rho \beta(t) \frac{Y_t F_A(t, z_t; \nu, \lambda_I)}{X_t} \end{aligned}$$

plus a position in the income contract. After retirement, $\beta(t) = 0$, and F_R replaces F_A .

¹⁰ Our method also works in quadratic settings (Liu 2007). Our method can be implemented outside the affine–quadratic settings, but then we lose the benefits of building on closed-form solutions and it may be more difficult to see how to transform the numerically computed, consumption–investment strategy in an artificial market into a feasible strategy in the true market.

We find an upper bound on the value function in the true market by minimizing $J(t, x, y, z; \nu, \lambda_I)$ over all simple (ν, λ_I) via a standard numerical optimization routine. Each optimal strategy in a simple artificial market is transformed into a feasible strategy in the true market by disregarding the investment in the income contract, scaling human wealth, e.g., by a factor $(1 - e^{-\eta x_t})$ as in our main example, and by replacing the unconstrained π_i by some $\hat{\pi}_i \in \mathcal{K}$, which is close to π_i in some norm. For example, $\hat{\pi}_i$ could be the projection¹¹ of π_i on \mathcal{K} . The appropriate choice of $\hat{\pi}_i$ may depend on the precise setting. The resulting feasible strategy $(\hat{c}, \hat{\pi})$ depends on ν, λ_I, η . Each of these strategies can be evaluated by Monte Carlo simulation, and the evaluation can be built into a standard optimization routine to find the best of these feasible strategies. This will be the strategy suggested by our method. By comparing the expected utility generated by this strategy to the upper utility bound, an upper bound on the error (in wealth terms) can be computed.¹²

10. Conclusion

This paper has suggested and tested an easy procedure for finding a simple, near-optimal consumption and investment strategy of a power-utility investor receiving an unspanned labor income stream. This procedure is valuable because the truly optimal solution is not known in closed form and is very difficult to approximate precisely using standard numerical solution techniques. For illustrative purposes we have focused on standard models of the price dynamics of traded assets. However, as explained above, the procedure applies to models of the affine or quadratic classes considered in many recent papers on portfolio choice in the absence of labor income, because in those settings (i) we would still be able to find explicit solutions in the artificially completed markets, and (ii) we can still evaluate the performance of a specific strategy by Monte Carlo simulations.

Acknowledgments

The authors appreciate comments from anonymous referees, Nicole Branger, Jakša Cvitanić, Bjarne Astrup Jensen,

¹¹ Several projections may be possible. For example, assume two risky assets and constraints (a) $\pi_1, \pi_2 \in [0, 1]$ and (b) $\pi_1 + \pi_2 \leq 1$. If the unconstrained artificial market portfolio satisfies (a), but $\pi_1 + \pi_2 > 1$, we can perform a “proportional projection” to $\hat{\pi}_i = \pi_i / (\pi_1 + \pi_2)$, or a “one-sided projection” so that $\hat{\pi}_i = \pi_i, \hat{\pi}_j = 1 - \pi_i$, or any combination hereof. Different projections may have different consequences for expected utility, and some experimentation is recommended.

¹² As for other numerical methods, the performance of our method may vary with the specifics of the problem. For certain problems, it may be important to consider a broader class of parameterized consumption–investment strategies, for example, generated from a broader class of artificial markets than those considered here.

and participants at presentations at the CREATES Symposium on Dynamic Asset Allocation (Aarhus 2009), the German Finance Association Meeting (Frankfurt 2009), the Swiss Society for Financial Market Research Conference (Zürich 2010), the Bachelier Finance Society Congress (Toronto 2010), the European Finance Association Meeting (Frankfurt 2010), and the Danish Center for Accounting and Finance (Nyborg 2011). Bick and Kraft gratefully acknowledge financial support by Deutsche Forschungsgemeinschaft. Munk gratefully acknowledges financial support by the Danish Center for Accounting and Finance. An earlier version of this paper was entitled “Investment, Income, and Incompleteness.”

Appendix. Proofs for Constant Interest Rates

PROOF OF THEOREM 1. The dynamics of financial wealth in retirement in the $\nu_R(t)$ artificial market is

$$dX_t = X_t \left[(\tilde{r}(t) + \pi_{St} [\sigma_S \lambda_S + \nu_R(t)]) dt + \pi_{St} \sigma_S dW_t \right] + (\Upsilon Y_{\tilde{T}} - c_t) dt,$$

where $\tilde{r}(t) = r + \nu_R(t)^-$. The HJB equation for the indirect utility function $J = J_R(t, x; \nu_R)$ is

$$\delta J = \sup_{c, \pi_S} \left\{ U(c) + J_t + (\Upsilon Y_{\tilde{T}} - c) J_x + (\tilde{r}(t) + \pi_S [\sigma_S \lambda_S + \nu_R(t)]) x J_x + \frac{1}{2} \pi_S^2 \sigma_S^2 x^2 J_{xx} \right\},$$

where subscripts on J denote partial derivatives, and the terminal condition is $J(T, x) = \varepsilon(1/(1-\gamma))x^{1-\gamma}$. The first-order conditions for c and π_S lead to

$$c = J_x^{-1/\gamma}, \quad \pi_S = - \frac{\sigma_S \lambda_S + \nu_R(t)}{\sigma_S^2} \frac{J_x}{x J_{xx}}. \quad (19)$$

After substitution of these controls, the HJB equation reduces to

$$\delta J = \frac{\gamma}{1-\gamma} J_x^{1-1/\gamma} + J_t + \Upsilon Y_{\tilde{T}} J_x + \tilde{r}(t) x J_x - \frac{1}{2} \frac{(\sigma_S \lambda_S + \nu_R(t))^2}{\sigma_S^2} \frac{J_x^2}{J_{xx}}. \quad (20)$$

Conjecture a solution of the form

$$J(t, x) = \frac{1}{1-\gamma} g(t)^\gamma (x + \Upsilon Y_{\tilde{T}} F(t))^{1-\gamma}, \quad (21)$$

where $g(T) = \varepsilon^{1/\gamma}$, $F(T) = 0$ to satisfy the terminal condition. After substituting (21) into (20), we collect terms involving $(x + \Upsilon Y_{\tilde{T}} F(t))^{1-\gamma}$ and the remaining terms that all involve $(x + \Upsilon Y_{\tilde{T}} F(t))^{-\gamma}$. This leads to the ordinary differential equations

$$F'(t) - \tilde{r}(t)F(t) + 1 = 0, \quad g'(t) - h(t)g(t) + 1 = 0,$$

where

$$h(\tau) = \frac{\delta}{\gamma} + \frac{\gamma-1}{\gamma} (r + \nu_R(\tau)^-) + \frac{\gamma-1}{2\gamma^2} \frac{(\sigma_S \lambda_S + \nu_R(\tau))^2}{\sigma_S^2}.$$

The solutions that satisfy the above-mentioned terminal values are

$$F(t) = \int_t^T e^{-\int_t^\tau \tilde{r}(\tau) d\tau} d\tau,$$

$$g(t) = \varepsilon^{1/\gamma} e^{-\int_t^T h(\tau) d\tau} + \int_t^T e^{-\int_t^\tau h(\tau) d\tau} d\tau.$$

Inserting the conjectured indirect utility function into (19), we obtain the optimal controls

$$c = \frac{x + \Upsilon Y_{\tilde{T}} F(t)}{g(t)}, \quad \pi_S = \frac{1}{\gamma} \frac{x + \Upsilon Y_{\tilde{T}} F(t)}{x} \frac{\sigma_S \lambda_S + \nu_A(t)}{\sigma_S^2},$$

which define an admissible strategy in this artificial market. \square

PROOF OF THEOREM 2. The dynamics of financial wealth before retirement in the artificial market characterized by $\nu_A(t)$, $\nu_R(t)$, $\lambda_I(t)$ is

$$dX_t = X_t [\tilde{r}(t) + \pi_{St} [\sigma_S \lambda_S + \nu_A(t)] + \pi_{It} \lambda_I(t)] dt + \pi_{St} \sigma_S dW_t + \pi_{It} dW_{Yt} + (Y_t - c_t) dt,$$

where $\tilde{r}(t) = r + \nu_A(t)^-$. The HJB equation for the indirect utility function $J = J_A(t, x; \nu_A, \lambda_I)$ is

$$\begin{aligned} \delta J = \sup_{c, \pi_S, \pi_I} \{ & U(c) + J_t + (y - c) J_x \\ & + (\tilde{r}(t) + \pi_S [\sigma_S \lambda_S + \nu_A(t)] + \pi_I \lambda_I(t)) x J_x \\ & + \frac{1}{2} (\pi_S^2 \sigma_S^2 + \pi_I^2) x^2 J_{xx} + \alpha y J_y \\ & + \frac{1}{2} \beta^2 y^2 J_{yy} + (\rho \pi_S \sigma_S + \sqrt{1 - \rho^2} \pi_I) \beta x y J_{xy} \}, \end{aligned}$$

where subscripts on J denote partial derivatives, and the terminal condition is

$$J(\tilde{T}, x, y) = \frac{1}{1 - \gamma} g_R(\tilde{T}; \nu_R) (x + \Upsilon y F_R(\tilde{T}; \nu_R))^{1 - \gamma}.$$

The first-order conditions for c , π_S , and π_I lead to

$$\begin{aligned} c &= J_x^{-1/\gamma}, \quad \pi_S = -\frac{\sigma_S \lambda_S + \nu_A(t)}{\sigma_S^2} \frac{J_x}{x J_{xx}} - \frac{\beta \rho y J_{xy}}{\sigma_S x J_{xx}}, \\ \pi_I &= -\lambda_I(t) \frac{J_x}{x J_{xx}} - \beta \sqrt{1 - \rho^2} \frac{y J_{xy}}{x J_{xx}}. \end{aligned} \quad (22)$$

After substitution of these controls, the HJB equation becomes

$$\begin{aligned} \delta J &= \frac{\gamma}{1 - \gamma} J_x^{1 - 1/\gamma} + J_t + y J_x + \tilde{r}(t) x J_x + \alpha y J_y + \frac{1}{2} \beta^2 y^2 J_{yy} \\ &\quad - \frac{1}{2} \left(\frac{(\sigma_S \lambda_S + \nu_A(t))^2}{\sigma_S^2} + \lambda_I(t)^2 \right) \frac{J_x^2}{J_{xx}} - \frac{1}{2} \beta^2 y^2 \frac{J_{xy}^2}{J_{xx}} \\ &\quad - \beta \left(\rho \frac{\sigma_S \lambda_S + \nu_A(t)}{\sigma_S} + \sqrt{1 - \rho^2} \lambda_I(t) \right) y \frac{J_x J_{xy}}{J_{xx}}. \end{aligned} \quad (23)$$

Conjecture a solution of the form

$$J(t, x, y) = \frac{1}{1 - \gamma} g(t)^\gamma (x + y F(t))^{1 - \gamma}, \quad (24)$$

where $g(\tilde{T}) = g_R(\tilde{T}; \nu_R)$, $F(\tilde{T}) = \Upsilon F_R(\tilde{T}; \nu_R)$ to satisfy the terminal condition. After substituting (24) into (23), the terms involving $(x + y F(t))^{-\gamma - 1}$ cancel. We collect terms involving $(x + y F(t))^{1 - \gamma}$ and the remaining terms that all involve $(x + y F(t))^{-\gamma}$. This leads to the ordinary differential equations

$$F'(t) - r_A(t) F(t) + 1 = 0, \quad g'(t) - h_A(t) g(t) + 1 = 0,$$

where

$$\begin{aligned} r_A(\tau) &= \tilde{r}(\tau) - \alpha - \beta \left(\rho \frac{\sigma_S \lambda_S + \nu_A(\tau)}{\sigma_S} + \sqrt{1 - \rho^2} \lambda_I(\tau) \right), \\ h_A(\tau) &= \frac{\delta}{\gamma} + \frac{\gamma - 1}{\gamma} (r + \nu_A(\tau)^-) \\ &\quad + \frac{\gamma - 1}{2\gamma^2} \left[\frac{(\sigma_S \lambda_S + \nu_A(\tau))^2}{\sigma_S^2} + \lambda_I(\tau)^2 \right]. \end{aligned}$$

The solutions consistent with the above-mentioned values at time \tilde{T} are

$$\begin{aligned} F(t) &= e^{-\int_t^{\tilde{T}} r_A(u) du} \Upsilon F_R(\tilde{T}; \nu_R) + \int_t^{\tilde{T}} e^{-\int_t^u r_A(\tau) d\tau} du, \\ g(t) &= e^{-\int_t^{\tilde{T}} h_A(u) du} g_R(\tilde{T}; \nu_R) + \int_t^{\tilde{T}} e^{-\int_t^u h_A(\tau) d\tau} du. \end{aligned}$$

Inserting the conjectured indirect utility function into (22), we obtain the optimal controls

$$\begin{aligned} c &= \frac{x + y F(t)}{g(t)}, \quad \pi_S = \frac{1}{\gamma} \frac{x + y F(t)}{x} \frac{\sigma_S \lambda_S + \nu_A(t)}{\sigma_S^2} - \frac{\beta \rho y F(t)}{\sigma_S x}, \\ \pi_I &= \frac{\lambda_I(t)}{\gamma} \frac{x + y F(t)}{x} - \beta \sqrt{1 - \rho^2} \frac{y F(t)}{x}, \end{aligned}$$

which define an admissible strategy in this artificial market. \square

References

- Bodie Z, Merton RC, Samuelson WF (1992) Labor supply flexibility and portfolio choice in a life cycle model. *J. Econom. Dynam. Control* 16(3–4):427–449.
- Bouchard B, Ekeland I, Touzi N (2004) On the Malliavin approach to Monte Carlo approximation of conditional expectations. *Finance Stochastics* 8(1):45–71.
- Brandt MW, Goyal A, Santa-Clara P, Stroud JR (2005) A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Rev. Financial Stud.* 18(3):831–873.
- Brennan MJ, Xia Y (2000) Stochastic interest rates and the bond-stock mix. *Eur. Finance Rev.* 4(2):197–210.
- Brennan MJ, Schwartz ES, Lagnado R (1997) Strategic asset allocation. *J. Econom. Dynam. Control* 21(8–9):1377–1403.
- Bucks BK, Kennickell AB, Mach TL, Moore KB (2009) Changes in U.S. family finances from 2004 to 2007: Evidence from the Survey of Consumer Finances. *Federal Reserve Bull.* 95(February):A1–A56.
- Campbell JY, Viceira LM (2001) Who should buy long-term bonds? *Amer. Econom. Rev.* 91(1):99–127.
- Carriere JF (1996) Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Math. Econom.* 19(1):19–30.
- Christensen PO, Larsen K, Munk C (2012) Equilibrium in securities markets with heterogeneous investors and unspanned income risk. *J. Econom. Theory* 147(3):1035–1063.
- Cocco JF, Gomes FJ, Maenhout PJ (2005) Consumption and portfolio choice over the life cycle. *Rev. Financial Stud.* 18(2):491–533.
- Cvitanic J, Karatzas I (1992) Convex duality in constrained portfolio optimization. *Ann. Appl. Probab.* 2(4):767–818.
- Cvitanic J, Goukasian L, Zapatero F (2003) Monte Carlo computation of optimal portfolios in complete markets. *J. Econom. Dynam. Control* 27(6):971–986.
- Cvitanic J, Schachermayer W, Wang H (2001) Utility maximization in incomplete markets with random endowment. *Finance Stochastics* 5(2):259–272.

- Das SR, Sundaram RK (2002) An approximation algorithm for optimal consumption/investment problems. *Intelligent Systems Accounting, Finance Management* 11(2):55–69.
- Detemple J, Garcia R, Rindisbacher M (2003) A Monte-Carlo method for optimal portfolios. *J. Finance* 58(1):401–446.
- Fahim A, Touzi N, Warin X (2011) A probabilistic numerical method for fully nonlinear parabolic PDEs. *Ann. Appl. Probab.* 21(4):1322–1364.
- Garlappi L, Skoulakis G (2010) Solving consumption and portfolio choice problems: The state variable decomposition method. *Rev. Financial Stud.* 23(9):3346–3400.
- Hakansson NH (1970) Optimal investment and consumption strategies under risk for a class of utility functions. *Econometrica* 38(5):587–607.
- Haugh MB, Kogan L, Wang J (2006) Evaluating portfolio policies: A duality approach. *Oper. Res.* 54(3):405–418.
- Henderson V (2005) Explicit solutions to an optimal portfolio choice problem with stochastic income. *J. Econom. Dynam. Control* 29(7):1237–1266.
- Karatzas I, Lehoczky JP, Shreve SE, Xu G-L (1991) Martingale and duality methods for utility maximization in an incomplete market. *SIAM J. Control Optim.* 29(3):702–730.
- Kogan L, Uppal R (2000) Risk aversion and optimal portfolio policies in partial and general equilibrium economies. Working paper, Massachusetts Institute of Technology, Cambridge.
- Kojien RSJ, Nijman TE, Werker BJM (2007) Appendix describing the numerical method used in “When can life-cycle investors benefit from time-varying bond risk premia?” Working paper, University of Chicago, Chicago.
- Kojien RSJ, Nijman TE, Werker BJM (2010) When can life-cycle investors benefit from time-varying bond risk premia? *Rev. Financial Stud.* 23(2):741–780.
- Kraft H, Munk C (2011) Optimal housing, consumption, and investment decisions over the life-cycle. *Management Sci.* 57(6):1025–1041.
- Kushner HJ, Dupuis PG (2001) *Numerical Methods for Stochastic Control Problems in Continuous Time, Applications of Mathematics*, Vol. 24, 2nd ed. (Springer-Verlag, New York).
- Liu J (2007) Portfolio selection in stochastic environments. *Rev. Financial Stud.* 20(1):1–39.
- Longstaff FA, Schwartz ES (2001) Valuing American options by simulation: A simple least-squares approach. *Rev. Financial Stud.* 14(1):113–147.
- Merton RC (1971) Optimum consumption and portfolio rules in a continuous-time model. *J. Econom. Theory* 3(4):373–413.
- Munk C (2000) Optimal consumption-investment policies with undiversifiable income risk and liquidity constraints. *J. Econom. Dynam. Control* 24(9):1315–1343.
- Munk C (2003) Numerical methods for continuous-time, continuous-state stochastic control problems: Experiences from Merton’s problem. *Appl. Math. Comput.* 136(1):47–77.
- Munk C, Sørensen C (2010) Dynamic asset allocation with stochastic income and interest rates. *J. Financial Econom.* 96(3):433–462.
- Sørensen C (1999) Dynamic asset allocation and fixed income management. *J. Financial Quant. Anal.* 34(4):513–531.
- Svensson LEO, Werner IM (1993) Nontraded assets in incomplete markets. *Eur. Econom. Rev.* 37(5):1149–1168.
- van Binsbergen JH, Brandt MW (2007) Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function? *Computational Econom.* 29(3–4):355–367.
- Van Hemert O (2010) Household interest rate risk management. *Real Estate Econom.* 38(3):467–505.
- Vasicek O (1977) An equilibrium characterization of the term structure. *J. Financial Econom.* 5(2):177–188.
- Viceira LM (2001) Optimal portfolio choice for long-horizon investors with nontradable labor income. *J. Finance* 56(2):433–470.
- Yao R, Zhang HH (2005) Optimal consumption and portfolio choices with risky housing and borrowing constraints. *Rev. Financial Stud.* 18(1):197–239.