

Dynamic Mean-Variance Asset Allocation*

Suleyman Basak
London Business School and CEPR
Institute of Finance and Accounting
Regent's Park
London NW1 4SA
United Kingdom
Tel: (44) 20 7000 8256
Fax: (44) 20 7000 8201
sbasak@london.edu

Georgy Chabakauri
London Business School
Institute of Finance and Accounting
Regent's Park
London NW1 4SA
United Kingdom
Tel: (44) 20 7000 8241
Fax: (44) 20 7000 8201
gchabakauri.phd2004@london.edu

This revision: June 2008

*We are grateful to Ravi Bansal, Tomas Bjork, Peter Bossaerts, Darrell Duffie, Bernard Dumas, Francisco Gomes, Hong Liu, Marcel Rindisbacher, Raman Uppal and the seminar participants at Duke-UNC Asset Pricing Conference, European Finance Association Meetings, Swiss Finance Institute Annual Meeting, SFI Conference on Derivatives in Portfolio Management, Duke University, Goethe University of Frankfurt, Instituto Empresa, London Business School, Stockholm School of Economics, University of Lausanne, University of Mannheim, University of Toronto and University of Warwick for helpful comments. All errors are our responsibility.

Dynamic Mean-Variance Asset Allocation

Abstract

Mean-variance criteria remain prevalent in multi-period problems, and yet not much is known about their dynamically optimal policies. We provide a fully analytical characterization of the optimal dynamic mean-variance portfolios within a general incomplete-market economy, and recover a simple structure that also inherits several conventional properties of static models. We also identify a probability measure that incorporates intertemporal hedging demands and facilitates much tractability in the explicit computation of portfolios. We solve the problem by explicitly recognizing the time-inconsistency of the mean-variance criterion and deriving a recursive representation for it, which makes dynamic programming applicable. We further show that our time-consistent solution is generically different from the pre-commitment solutions in the extant literature, which maximize the mean-variance criterion at an initial date and which the investor commits to follow despite incentives to deviate. We illustrate the usefulness of our analysis by explicitly computing dynamic mean-variance portfolios under various stochastic investment opportunities in a straightforward way, which does not involve solving a Hamilton-Jacobi-Bellman differential equation. A calibration exercise shows that the mean-variance hedging demands may comprise a significant fraction of the investor's total risky asset demand.

Journal of Economic Literature Classification Numbers: G11, D81, C61.

Keywords: Mean-Variance Analysis, Multi-Period Portfolio Choice, Stochastic Investment Opportunities, Time-Consistency, Dynamic Programming, Incomplete Markets.

1. Introduction

The mean-variance analysis of Markowitz (1952) has long been recognized as the cornerstone of modern portfolio theory. Its simplicity and intuitive appeal have led to its widespread use in both academia and industry. Originally cast in a single-period framework, the mean-variance paradigm has no doubt also inspired the development of the multi-period portfolio choice literature. To this day, the mean-variance criteria are employed in many multi-period problems by financial economists, but typically for a myopic investor, who in each period maximizes her next-period objective (e.g., among others, Ait-Sahalia and Brandt, 2001; Campbell and Viceira, 2002; Jagannathan and Ma, 2003; Bansal, Dahlquist and Harvey, 2004; Brandt, 2004; Acharya and Pedersen, 2005; Hong, Scheinkman and Xiong, 2006; Campbell, Serfaty-de Medeiros and Viceira, 2007).¹ While the myopic assumption allows analytical tractability and abstracts away from dynamic hedging considerations, there is growing evidence that intertemporal hedging demands may comprise a significant part of the total risky asset demand (e.g., Campbell and Viceira, 1999; Brandt, 1999). However, solving the dynamic asset-allocation problem with mean-variance criteria has had mixed success to date. A major obstacle has been the inability to directly apply the traditional dynamic programming approach due to the failure of the iterated-expectations property for mean-variance objectives. A growing recent literature tackles this by just characterizing the optimal policy chosen at an initial date, by either employing martingale methods or tractable auxiliary problems in complete market settings (as discussed below). However, due to the time-inconsistency of the mean-variance criteria, the investor may find it optimal to deviate from this policy unless she is able to pre-commit, and henceforth we refer to it as the pre-commitment policy. Many decades have passed since the original Markowitz analysis, and yet we still lack a comprehensive treatment of dynamically optimal policies consistent with the mean-variance criteria.

In this paper, we solve the dynamic asset allocation problem of a mean-variance optimizer in an incomplete-market setting, and provide a simple, tractable solution for the risky stock holdings. To our knowledge, ours is the first to obtain within a general environment a fully analytical characterization of the dynamically optimal mean-variance policies, from which the investor has no incentive to deviate, namely, the time-consistent policies. Towards this, we consider the familiar multi-period asset allocation problem of an investor, who has preferences over terminal wealth and dynamically allocates wealth between a risky stock and a riskless bond. The investor is guided by the mean-variance criterion, linearly trading-off mean and variance of terminal wealth. Our setting is a continuous-time Markovian economy with stochastic

¹We acknowledge the well-known theoretical objections to the mean-variance criteria if interpreted as investors' preferences, namely admitting potentially negative terminal wealth, increasing absolute risk aversion, and potentially non-monotonicity of preferences. Despite the theoretical limitations, the mean-variance criteria remain relatively popular in practice and academia due to their simplicity and tractability, which we will also demonstrate in our analysis. Interestingly, recent evidence in neuroscience, as provided and discussed by Bossaerts, Preuschoff, and Quartz (2006, 2008), suggests that the human brain appears to analyze risky gambles by considering variance and expectation separately, consistent with the mean-variance criteria.

investment opportunities, allowing for a potentially incomplete market. Our solution method for the determination of optimal dynamic mean-variance policies is based on the derivation of a recursive formulation so that dynamic programming can be employed. This recursive derivation is complicated by the fact that mean-variance criteria in a multi-period setting result in time-inconsistency of investment policies, in that the investor has an incentive to deviate from an initial policy at a later date. The intuition for this is that sitting at a point in time, the mean-variance investor perceives the variability of terminal wealth to be higher than the anticipated variability at a future date. To address this problem, we decompose the investor’s conditional objective function as her expected future objective plus a term accounting for the incentives to deviate, which then leads to the desired recursive formulation. This in turn allows us to employ dynamic programming, derive the Hamilton-Jacobi-Bellman (HJB) equation, and obtain an analytical solution to the problem. We also note that the same solution can alternatively be obtained as the Nash equilibrium outcome of an intra-personal game by the dynamic mean-variance investor, similarly to the literature on consumer choice under hyperbolic discounting (e.g., Harris and Laibson, 2001).

The optimal stock investment policy of a dynamic mean-variance optimizer has a simple structure, being comprised of familiar myopic and intertemporal hedging terms. The novel feature of our case is that we identify the hedging demand to be driven by the expected total gains or losses from the stock investments over the investment horizon, in contrast to being driven by the value function in the extant literature. This is because the mean-variance value function is linear in wealth. Since the conditional variance of terminal wealth equals that of future portfolio gains, the mean-variance hedging demands are determined by the anticipated portfolio gains. The economic role of the hedging demands in our setting is then straightforward: when the stock return is negatively related to the anticipated portfolio gains, the gains in one offset the losses in the other. This leads to a lower variability of wealth, making the stock more attractive, and hence inducing a positive hedging demand; and vice versa for a negative hedging demand. We then identify a unique probability measure, labeled a “hedge-neutral” measure, which absorbs the hedging demands so that the anticipated investment gains under this measure look as if the investor were myopic. This representation under the new measure facilitates considerable tractability, allowing one to easily determine the mean-variance portfolios explicitly or otherwise perform Monte-Carlo simulation straightforwardly.²

We also find the dynamic mean-variance policies to inherit a number of conventional properties of single-period models, such as, the higher the stock volatility, bond interest rate or investor risk aversion, the lower the stock investment (in absolute terms). However, these dynamic policies also generate rich implications related to the effects of investment horizon, market price of risk, and market incompleteness. For example, the variance of terminal wealth in incomplete markets is higher than that in complete markets, and consequently the mean-variance investor

²We also remark that given our dynamically optimal mean-variance policy, it is possible to recover time-consistent objective functions that would lead to the same policy. One such function is an increasing, concave, state-dependent criterion of CARA form.

with positive hedging demand is worse off in incomplete markets. We also compare our time-consistent solution to the mean-variance pre-commitment solution in a simple complete market settings. The mean-variance investor under pre-commitment maximizes her initial objective and pre-commits to that initial investment policy, not deviating at subsequent times. We demonstrate that the time-consistent investment policy, obtained via dynamic programming, is generically different from the pre-commitment policy, obtained via martingale methods. Although for very short investment horizons the pre-commitment solution approximates the time-consistent one up to second order terms, for plausible horizons, the two solutions can differ considerably. Of course, with standard utility functions, the two solutions are well-known to coincide (Karatzas and Shreve, 1998).

We illustrate the practical usefulness of our analysis by considering the dynamic mean-variance problem under several stochastic investment opportunities that have been studied in the literature for other preference specifications. In particular, we specialize our economic setting to the constant elasticity of variance model in a complete market (Cox and Ross, 1976; Schroder, 1989), a mean-reverting stochastic-volatility model in an incomplete market (Liu, 2001; Chacko and Viceira, 2005; Heston, 1993), and a time-varying Gaussian mean-returns model in an incomplete market (Kim and Omberg, 1998; Campbell and Viceira, 1999; Wachter, 2002). In all these applications, we explicitly derive the dynamic mean-variance portfolios as a straightforward exercise, by computing the anticipated gains process under the hedge-neutral measure, which amounts to evaluating the expectation of the squared market price of risk. We emphasize that our computations do not resort to solving an HJB PDE for the investor's value function, as would be the case for other popular objective specifications. In addition to providing further insights, our explicit solutions allow us to assess the economic significance of the intertemporal hedging demands of a mean-variance optimizer. Specifically, we compute the percentage hedging demand over total demand in our richer incomplete market settings for a range of plausible parameter values. We find our results to be in line with those in the literature and show that the percentage hedging demand can be considerable in some economic settings, ranging from 18% to 84%, supporting the findings of Brandt (1999) and Campbell and Viceira (1999).

Finally, we consider extensions of our baseline analysis to economic settings in discrete-time and settings with stochastic interest rates, multiple stocks and multiple sources of uncertainty. We demonstrate our main results to be valid also under these alternative environments. Moreover, we here provide fully-explicit closed-form solutions for optimal investment policies in several discrete-time settings with stochastic investment opportunities that, to our knowledge, are new in the literature. In contrast, the extant literature characterizes optimal policies in such settings by employing either numerical methods or various approximations (e.g., Ait-Sahalia and Brandt, 2001; Bansal and Kiku, 2007; Brandt, Goyal, Santa-Clara and Stroud, 2005; Brandt and Santa-Clara, 2006; Campbell and Viceira, 1999, 2002, among others).

There is a growing literature investigating the multi-period portfolio problem of a mean-variance investor. Bajeux-Besnainou and Portait (1998), Bielecki, Jin, Pliska and Zhou (2005),

Cvitanic, Lazrak and Wang (2007), Cvitanic and Zapatero (2004), Zhao and Ziemba (2002) consider continuous-time complete market settings and employ martingale methods to solve for the variance minimizing policy subject to the constraint that expected terminal wealth equals some given level, sitting at an initial date. Cochrane (2005) in an incomplete-market setting solves for the optimal investment policy that minimizes the “long-run” variance of portfolio returns subject to the constraint that the long-run mean of portfolio returns equals a pre-specified target level. However, the ensuing solution in these works is a pre-commitment investment policy chosen at an initial date since the investor may subsequently find it optimal to deviate from if the constraint is violated in the future. Duffie and Richardson (1991) study the futures hedging policy in a continuous-time incomplete market. They solve the hedging problem with a mean-variance objective sitting at an initial date, obtaining the pre-commitment solution, by observing that the optimal policy here also solves the hedging problem with a quadratic objective for some specific parameters. Recognizing the difficulty of applying dynamic programming, Li and Ng (2000), Leippold, Trojani and Vanini (2004) in discrete-time, Zhou and Li (2000), Lim and Zhou (2002) in continuous-time, use a similar approach to solve for mean-variance portfolios in complete market settings. Specifically, these authors show that the investment policy that solves the mean-variance problem sitting at an initial date also solves the one with a quadratic objective for some specific parameters. The solution to the quadratic auxiliary optimization is then derived, which gives the pre-commitment strategy for the mean-variance problem. Brandt (2004) considers portfolio choice with mean-variance criterion over portfolio returns. The solution is provided when the investor chooses portfolio weights for several periods ahead, implicitly assuming pre-commitment.

Our work also contributes to the multi-period portfolio choice literature that provides explicit closed-form solutions for optimal investment policies under various stochastic investment opportunities, all obtained in continuous-time settings. Kim and Omberg (1996) explicitly solve for the optimal portfolio of an investor with constant relative risk aversion (CRRA) preferences over terminal wealth when the market price of risk follows a mean-reverting Ornstein-Uhlenbeck process in an incomplete market setting. Merton (1971) and Wachter (2002) provide solutions to similar problems for constant absolute risk aversion (CARA) and CRRA investors, respectively, with intermediate consumption under complete markets. Maenhout (2006) extends the Kim-Omberg results by providing explicit solutions for an investor who worries about model specification, while Huang and Liu (2007) provide a generalization with incomplete information. Liu (2001, 2007) obtains explicit solutions for an investor with CRRA preferences over terminal wealth facing an incomplete market with stochastic volatility. In similar models, Chacko and Viceira (2005) provide the explicit solution for an investor having recursive preferences over intertemporal consumption with unit elasticity of intertemporal substitution, while Liu (2007) for a CRRA investor with intertemporal consumption in a complete market. In related problems, nearly-explicit closed-form solutions have additionally been obtained by Brennan and Xia (2002) and Sangvinatsos and Wachter (2005). In general, however, obtaining fully-explicit closed-form solutions to dynamic portfolio choice problems with stochastic investment opportunities is

a daunting task, and one would need to resort to numerical methods, such as those proposed by Detemple, Garcia and Rindisbacher (2003), Cvitanic, Goukasian and Zapatero (2003), and Brandt, Goyal, Santa-Clara and Stroud (2005).

The remainder of the paper is organized as follows. In Section 2, we present our methodology for the determination of optimal dynamic mean-variance policies. We then provide the time-consistent solution, discuss its properties, and compare it with the pre-commitment policy. In Section 3, we provide applications of our analysis to various stochastic investment opportunities, while in Section 4, we discuss the extensions to discrete-time, multiple-stock and stochastic interest rate settings. Section 5 concludes and the Appendix provides all proofs.

2. Asset Allocation with Mean-Variance Criteria

2.1. Economic Setup

We consider a continuous-time Markovian economy with a finite horizon $[0, T]$. Uncertainty is represented by a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$, on which are defined two correlated Brownian motions, w and w_X , with correlation ρ . All stochastic processes are assumed to be adapted to $\{\mathcal{F}_t, t \in [0, T]\}$, the augmented filtration generated by w and w_X . In what follows, given our focus, we assume all processes and moments introduced are well-defined, without explicitly stating the regularity conditions.

Trading may take place continuously in two securities, a riskless bond and a risky stock. The bond provides a constant interest rate r . The stock price, S , follows the dynamics

$$\frac{dS_t}{S_t} = \mu(S_t, X_t, t)dt + \sigma(S_t, X_t, t)dw_t, \quad (1)$$

where the stock mean return, μ , and volatility, σ , are deterministic functions of S and the state variable X , which satisfies

$$dX_t = m(X_t, t)dt + \nu(X_t, t)dw_{Xt}. \quad (2)$$

Under appropriate conditions, the stochastic differential equations (1)–(2) have a unique solution (S, X) , which is a joint Markov process. We will denote μ_t , σ_t , m_t and ν_t as shorthand for the coefficients in equations (1)–(2). We note that under this setup, the market is incomplete as trading in the stock and bond cannot perfectly hedge the changes in the stochastic investment opportunity set. However, in the special cases of perfect correlation between the stock return and state variable, $\rho = \pm 1$, dynamic market completeness obtains. For the case of zero correlation, there is no hedging demand for the state variable since trading in the stock cannot hedge the fluctuations in the state variable.

An investor in this economy is endowed at time zero with an initial wealth of W_0 . The investor chooses an investment policy, θ , where θ_t denotes the dollar amount invested in the stock at time

t . The investor's wealth process W then follows

$$dW_t = [rW_t + \theta_t(\mu_t - r)] dt + \theta_t \sigma_t dw_t. \quad (3)$$

We assume that the investor is guided by mean-variance objectives over horizon wealth W_T . In particular, the dynamic optimization problem of the investor is given by

$$\max_{\theta} E[W_T] - \frac{\gamma}{2} \text{var}[W_T], \quad (4)$$

subject to the dynamic budget constraint (3). In Section 2.2, we provide the *time-consistent* solution to this problem via a recursive formulation that employs dynamic programming, while in Section 2.4, we provide the *pre-commitment* solution via a static formulation that employs martingale methods. We demonstrate that the two solutions are generically different.

In order to keep our problem analytically tractable, we follow the related literature and make the simplifying assumptions of constant interest rate and lack of intermediate consumption. It is unlikely that our model, with stochastic investment opportunities and potentially incomplete markets, could be solved analytically if these assumptions were relaxed, as in the related works of Kim and Omberg (1996), Liu (2001), Maenhout (2006). However, with an appropriate choice of a numeraire, we provide an extension of our results for the case with stochastic interest rates in Section 4.3. We further note that even though the mean-variance criterion (4) is in many ways similar to the time-consistent quadratic utility function, to our best knowledge the latter does not admit tractable optimal policies in our economic setting. For example, Brandt and Santa-Clara (2006) investigate dynamic portfolio selection with a quadratic criterion in an incomplete market setting and develop an approach that leads to approximate solutions.

2.2. Determination of Optimal Dynamic Investment Policy

In this Section, we first present our solution method, based on dynamic programming, for the determination of optimal dynamic mean-variance policies. The dynamic programming approach, however, is complicated by the presence of the variance term in the mean-variance objective function: it cannot be represented as the expected utility over terminal wealth, such as $E[u(W_T)]$, for which dynamic programming is readily applicable due to the iterated-expectation property $E_t[E_{t+\tau}[u(W_T)]] = E_t[u(W_T)]$. The violation of this property for mean-variance preferences makes the application of dynamic programming problematic (see e.g., Zhou and Li, 2000). To our best knowledge, there are no works that apply dynamic programming to derive explicit solutions to the multi-period mean-variance portfolio choice. We tackle this problem by first obtaining a tractable recursive formulation for the mean-variance objective, expressed as its expected future value plus an adjustment term, given by the time- t variance of expected terminal wealth. This explicit identification allows us to employ dynamic programming, derive the HJB equation and obtain an analytical solution to the problem. The intuition for the adjustment term is based on the observation that for a mean-variance optimizer sitting at time $t + \tau$, the variability of

terminal wealth may be lower than that of sitting at time t . This induces her to revise her time- t optimal policy at subsequent dates, and hence the need for the adjustment in her objective function sitting at any point in time.

Formally, the variability of terminal wealth, by the law of total variance (e.g., Weiss, 2005), is given by

$$\text{var}_t[W_T] = E_t[\text{var}_{t+\tau}(W_T)] + \text{var}_t[E_{t+\tau}(W_T)], \quad \tau > 0. \quad (5)$$

Clearly, the time- t variance exceeds the expected variance at time $t+\tau$. As a result, the investment policy θ_τ , for $\tau \geq t$, chosen at time t , accounts not only for the expected time- $(t+\tau)$ variance of the terminal wealth, but also for the variance of time- $(t+\tau)$ expected terminal wealth. However, since the latter vanishes as time interval τ elapses, the investor may deviate from the time- t optimal policy at time $t+\tau$. We now account for these incentives to deviate in the time- t objective function of the investor, who for each $t \in [0, T]$ maximizes

$$U_t \equiv E_t[W_T] - \frac{\gamma}{2} \text{var}_t[W_T], \quad (6)$$

subject to the dynamic budget constraint (3). Substituting (5) into (6) and using the law of iterated expectations, we obtain the following recursive representation for the time- t objective function of the mean-variance optimizer:

$$U_t = E_t[U_{t+\tau}] - \frac{\gamma}{2} \text{var}_t[E_{t+\tau}(W_T)]. \quad (7)$$

This representation reveals that decision-making at time t involves maximizing the expected future objective function, plus an adjustment that quantifies the investor's incentives to deviate from the time- t optimal policy. This adjustment enables us to determine the investment policy by backward induction, namely the *time-consistent* policy in that the investor optimally chooses the policy taking into account that she will act optimally in the future, if she is not restricted from revising her policy at all times. We elaborate more on the issue of time-consistency in Section 2.4.

Our next step towards the derivation of the HJB equation is to determine a recursive relationship for the value function. Given the optimal time-consistent policy θ_s^* , $s \in [t, T]$, derived by backward induction, the value function, J , is defined as

$$J(W_t, S_t, X_t, t) \equiv E_t[W_T^*] - \frac{\gamma}{2} \text{var}_t[W_T^*],$$

where terminal wealth W_T^* is computed under the optimal policy θ_s^* , $s \geq t$. Now let $\tau > 0$ denote the decision-making interval such that the investor can reconsider her investment policy chosen at time t only after the time interval τ elapses. Suppose further that at time t , the investor anticipates to follow the optimal policy θ_s^* from time $t+\tau$ onwards. Then, from the recursive representation of the objective function (7) and the definition of $J_{t+\tau}$, shorthand for the value function at $t+\tau$, the investor's time- t problem would be to find an investment policy θ_s , for $s \in [t, t+\tau]$ that maximizes

$$E_t[J_{t+\tau}] - \frac{\gamma}{2} \text{var}_t[E_{t+\tau}(W_T)]. \quad (8)$$

Sitting at time t , the investor accounts for the fact that starting from $t + \tau$, she will follow the policy that is optimal sitting at time $t + \tau$. Note, however, that because of the time-consistency adjustment term in (8), the investment policy θ_s^* , $s \geq t + \tau$, under which $J_{t+\tau}$ is computed, will not necessarily be optimal, when sitting at time t . Moreover, J_t is not equal to the maximum of its expected future value, $J_{t+\tau}$, as it would be in the case of standard utility functions over terminal wealth that have the form $E_t[u(W_T)]$.

Problem (8) presented above, and the definition of the value function after some algebra lead to the following recursive equation for J :

$$J_t = \max_{\theta_s, s \in [t, t+\tau]} E_t[J_{t+\tau}] - \frac{\gamma}{2} \text{var}_t[f_{t+\tau} - f_t + W_{t+\tau}e^{r(T-t-\tau)} - W_t e^{r(T-t)}], \quad (9)$$

subject to the budget constraint (3) and the terminal condition $J_T = W_T$, where f_t is shorthand for $f(W_t, S_t, X_t, t)$ defined as

$$f(W_t, S_t, X_t, t) \equiv E_t[W_T^*] - W_t e^{r(T-t)}, \quad (10)$$

representing expected total gains or losses from the optimal stock investment over the horizon $T - t$, while W_T^* is terminal wealth under the optimal policy θ_s^* , $s \geq t$.³ The dynamic budget constraint (3) allows us to obtain the following representation for f_t in terms of the optimal stock investment policy θ_s^* :

$$f(W_t, S_t, X_t, t) = E_t \left[\int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds \right]. \quad (11)$$

Going back to (9), it is clear that $f_{t+\tau}$ is defined using the optimal policy. This observation enables us to formulate the following Lemma, which gives the *HJB equation* in differential form and establishes some properties of θ_t^* , f_t and J_t .

Lemma 1. *The value function $J(W_t, S_t, X_t, t)$ of a mean-variance optimizing investor satisfies the following recursive equation:*

$$0 = \max_{\theta_t} E_t[dJ_t] - \frac{\gamma}{2} \text{var}_t[df_t + d(W_t e^{r(T-t)})], \quad (12)$$

subject to $J_T = W_T$ and the budget constraint (3), where f_t is as in (11). Moreover, $J(W_t, S_t, X_t, t)$ is separable in wealth and admits the representation

$$J(W_t, S_t, X_t, t) = W_t e^{r(T-t)} + \tilde{J}(S_t, X_t, t), \quad (13)$$

while f_t and the optimal investment policy θ_t^* do not depend on time- t wealth W_t and are functions of S_t , X_t and t only.

³In deriving (9) we use the fact that $\text{var}_t[f_{t+\tau} + W_{t+\tau}e^{r(T-t-\tau)}] = \text{var}_t[f_{t+\tau} - f_t + W_{t+\tau}e^{r(T-t-\tau)} - W_t e^{r(T-t)}]$.

We note that df_t term in (12) is unaffected by the control θ_t since according to Lemma 1, f_t does not depend on W_t and by definition is evaluated at the optimal policy. So, θ_t affects the adjustment term $\text{var}_t[df_t + d(W_t e^{r(T-t)})]$ via $d(W_t e^{r(T-t)})$ only. Using the separability property of J in (13) and applying Itô's Lemma to \tilde{J}_t , f_t and $W_t e^{r(T-t)}$, from the HJB equation (12) we obtain

$$0 = \max_{\theta_t} \left\{ \mathcal{D}\tilde{J}_t dt + \theta_t(\mu_t - r)e^{r(T-t)} dt - \frac{\gamma}{2} \text{var}_t \left[\sigma_t S_t \frac{\partial f_t}{\partial S_t} dw_t + \nu_t \frac{\partial f_t}{\partial X_t} dw_{X_t} + \theta_t \sigma_t e^{r(T-t)} dw_t \right] \right\}, \quad (14)$$

where \mathcal{D} denotes the Dynkin operator.⁴ Computation of the variance term in (14) yields the following PDE for the function \tilde{J}_t :

$$0 = \max_{\theta_t} \left\{ \mathcal{D}\tilde{J}_t + \theta_t(\mu_t - r)e^{r(T-t)} - \frac{\gamma}{2} \left[\sigma_t^2 S_t^2 \left(\frac{\partial f_t}{\partial S_t} \right)^2 + \nu_t^2 \left(\frac{\partial f_t}{\partial X_t} \right)^2 + 2\rho\nu_t \sigma_t S_t \frac{\partial f_t}{\partial S_t} \frac{\partial f_t}{\partial X_t} + \theta_t^2 \sigma_t^2 e^{2r(T-t)} + 2\theta_t \sigma_t \left(\sigma_t S_t \frac{\partial f_t}{\partial S_t} + \rho\nu_t \frac{\partial f_t}{\partial X_t} \right) e^{r(T-t)} \right] \right\}, \quad (15)$$

subject to $\tilde{J}_T = 0$. The HJB equation (15) is nonstandard in that in addition to the conventional term, $\mathcal{D}\tilde{J}_t + \theta_t(\mu_t - r)e^{r(T-t)}$, there is an adjustment component that is explicitly characterized in terms of anticipated investment gains, f_t , and the investment policy, θ_t . An attractive feature of the HJB equation (15) is that the maximized expression is a quadratic function of θ_t . We use this property to derive the following Proposition that provides a recursive representation for the optimal investment policy θ_t^* .

Proposition 1. *The optimal stock investment policy of a dynamic mean-variance optimizer is given by*

$$\theta_t^* = \frac{\mu_t - r}{\gamma\sigma_t^2} e^{-r(T-t)} - \left(S_t \frac{\partial f_t}{\partial S_t} + \frac{\rho\nu_t}{\sigma_t} \frac{\partial f_t}{\partial X_t} \right) e^{-r(T-t)}, \quad (16)$$

where the process f_t represents the expected total gains or losses from the stock investment and is given by

$$f(S_t, X_t, t) = E_t \left[\int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds \right]. \quad (17)$$

The optimal investment policy has a simple, familiar structure, and is given by myopic and intertemporal hedging terms. The myopic demand, $(\mu_t - r)/\gamma\sigma_t^2$, would be the investment policy for an investor who optimized over the next instant not accounting for her future investments, or the optimal policy if the investment opportunity set were constant. The intertemporal hedging demands, then, arise due to the need to hedge against the fluctuations in the investment opportunities, as in the related portfolio choice literature, following Merton (1971). What is different in our case is that we explicitly identify the hedging demands to be given by the sensitivities

⁴The Dynkin operator transforms an arbitrary twice continuously differentiable function $F(S_t, X_t, t)$ as follows:

$$\mathcal{D}F(S_t, X_t, t) = \frac{\partial F_t}{\partial t} + \mu_t S_t \frac{\partial F_t}{\partial S_t} + m_t \frac{\partial F_t}{\partial X_t} + \frac{1}{2} \left(\sigma_t^2 S_t^2 \frac{\partial^2 F_t}{\partial S_t^2} + \nu_t^2 \frac{\partial^2 F_t}{\partial X_t^2} + 2\rho\nu_t \sigma_t S_t \frac{\partial^2 F_t}{\partial X_t \partial S_t} \right).$$

of anticipated portfolio gains (f) to the stock price and state variable fluctuations, whereas in other works these sensitivities are in terms of the investor's value function. The reason is that the mean-variance conditional expected terminal wealth, and hence the value function, are linear in time- t wealth, and as a result, no hedging demand arises due to marginal utility fluctuations. Consequently, since the conditional variance of terminal wealth equals the conditional variance of future portfolio gains or losses, the anticipated portfolio gains or losses drive the hedging demands. This, in turn, enables us to provide more direct intuition on the implications of the hedging terms.

To see the role of the hedging demand, θ_H , we observe that

$$\theta_{Ht} \equiv - \left(S_t \frac{\partial f_t}{\partial S_t} + \frac{\rho \nu_t}{\sigma_t} \frac{\partial f_t}{\partial X_t} \right) e^{-r(T-t)} = - \frac{\text{cov}_t(dS_t/S_t, df_t)}{\sigma_t^2 dt} e^{-r(T-t)}. \quad (18)$$

The hedging demand is positive when the instantaneous stock return is negatively correlated with instantaneous portfolio gains. The reason for this is that when the stock return and anticipated portfolio gains move in opposite directions, losses in one are offset by the gains in the other. This leads to a lower variability of wealth, making the stock more attractive, and hence induces a positive hedging demand.

Even though the optimal stock investment expression is fairly intuitive, it is not characterized in terms of the exogenous parameters of the model since it relies on knowing the future optimal policy. To address this, we next recover an explicit representation for the anticipated portfolio gains, f . Substituting (16) into (17), we obtain the following representation for f under the original measure P :

$$f(S_t, X_t, t) = E_t \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] - E_t \left[\int_t^T \left(S_s \frac{\partial f_s}{\partial S_s} + \frac{\rho \nu_s}{\sigma_s} \frac{\partial f_s}{\partial X_s} \right) (\mu_s - r) ds \right]. \quad (19)$$

The first component in (19) comes from the myopic demand, while the second comes from the hedging demand. To facilitate tractability, we next look for a new probability measure under which the representation of f does not have the hedging related component. Since f_t is represented as a conditional expectation, by the Feynman-Kac theorem (Karatzas and Shreve, 1991), we obtain the following PDE after some manipulation:

$$\begin{aligned} \frac{\partial f_t}{\partial t} + r S_t \frac{\partial f_t}{\partial S_t} + \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} \right) \frac{\partial f_t}{\partial X_t} &+ \frac{1}{2} \left(\sigma_t^2 S_t^2 \frac{\partial^2 f_t}{\partial S_t^2} + \nu_t^2 \frac{\partial^2 f_t}{\partial X_t^2} + 2 \rho \nu_t \sigma_t S_t \frac{\partial^2 f_t}{\partial X_t \partial S_t} \right) \\ &+ \frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma_t} \right)^2 = 0, \end{aligned} \quad (20)$$

with $f_T = 0$. Again, by the Feynman-Kac theorem, (20) admits a unique solution with the following representation:

$$f(S_t, X_t, t) = E_t^* \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right], \quad (21)$$

where $E_t^*[\cdot]$ denotes the expectation under a new probability measure P^* such that the stock and state variable now follow dynamics with modified drifts⁵

$$\frac{dS_t}{S_t} = rdt + \sigma_t dw_t^*, \quad dX_t = \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} \right) dt + \nu_t dw_{Xt}^*, \quad (22)$$

and where w_t^* and w_{Xt}^* are Brownian motions under P^* with correlation ρ . Comparing (21) with (19), we see that measure P^* absorbs the hedging demand so that f represents the anticipated gains from the myopic portfolio only. We henceforth label P^* as the *hedge-neutral measure*. Note that this measure is also a risk-neutral measure since it modifies the drift of S to equal to rS . However, in our setting the risk-neutral measure is not unique due to market incompleteness;⁶ in the special case of a complete market, the hedge-neutral and risk-neutral measures coincide. Proposition 2 summarizes the results above.

Proposition 2. *The anticipated portfolio gains, f , can be expressed as*

$$f(S_t, X_t, t) = E_t^* \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right], \quad (23)$$

where $E_t^*[\cdot]$ denotes the expectation under the unique hedge-neutral measure P^* on which are defined two Brownian motions w^* and w_X^* with correlation ρ , given by

$$dw_t^* = dw_t + \frac{\mu_t - r}{\sigma_t} dt, \quad dw_{Xt}^* = dw_{Xt} + \rho \frac{\mu_t - r}{\sigma_t} dt, \quad (24)$$

and measure P^* is defined by the Radon-Nikodym derivative

$$\frac{dP^*}{dP} = e^{-\frac{1}{2} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds - \int_0^T \frac{\mu_s - r}{\sigma_s} dw_s}. \quad (25)$$

Consequently, the optimal investment policy is given by

$$\theta_t^* = \frac{\mu_t - r}{\gamma \sigma_t^2} e^{-r(T-t)} - \frac{1}{\gamma} \left(S_t \frac{\partial E_t^* \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right]}{\partial S_t} + \frac{\rho \nu_t}{\sigma_t} \frac{\partial E_t^* \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right]}{\partial X_t} \right) e^{-r(T-t)}. \quad (26)$$

⁵Since the coefficients assigned to partial derivatives $\partial f_t / \partial S_t$ and $\partial f_t / \partial X_t$ in the PDE (20) represent the drifts of stochastic processes for S and X , it follows that measure P^* modifies the drifts so that S and X satisfy (22).

⁶To see this, observe that dw_{Xt} can be decomposed as $dw_{Xt} = \rho dw_t + \sqrt{1 - \rho^2} d\tilde{w}_t$, where w_t and \tilde{w}_t are uncorrelated Brownian motions under P . Hence, any measure under which $dw_t^* = dw_t + (\mu_t - r)/\sigma_t dt$ and $d\tilde{w}_t^* = d\tilde{w}_t + g_t dt$ will be a risk-neutral measure irrespective of the process g_t . We further note that in an incomplete market, there is generally not a unique no-arbitrage price for a given payoff as it is impossible to hedge perfectly. Towards this, a common approach for pricing and hedging with market incompleteness is to choose a specific risk-neutral measure according to some criterion. Related to minimizing a quadratic loss function, a large literature in mathematical finance has developed which employs: the “minimal martingale measure” (Follmer and Sondermann, 1986; Schweizer, 1999) solving $\min_Q E[-\ln(dQ/dP)]$, the “variance optimal measure” (Schweizer, 1992) solving $\min_Q E[(dQ/dP)^2]$, and the “minimal entropy measure” (Miyahara, 1996) solving $\min_Q E[dQ/dP \ln(dQ/dP)]$, where dQ/dP denotes the Radon-Nikodym derivative of a risk-neutral measure Q with respect to the original measure P . Interestingly, our measure P^* , employed in a somewhat different context, turns out to coincide with the minimal martingale measure.

Proposition 2 provides a fully analytical characterization of the optimal investment policy in terms of the model parameters.⁷ The characterization identifies a unique measure P^* that incorporates intertemporal hedging demands so that only the expected gains or losses from the myopic portfolio need to be considered explicitly. This in turn allows us to explicitly compute the optimal dynamic mean-variance portfolios in a straightforward manner, as will be demonstrated in Section 3. For economic environments in which explicit computations are not possible, the optimal investment expression (26) can easily be computed numerically by standard Monte Carlo simulation methods, where the simulation would be performed under measure P^* . Additionally, the partial derivatives can be written in terms of Malliavin derivatives, leading to a more refined representation, which can then be computed by Monte Carlo simulation following the method of Detemple, Garcia and Rindisbacher (2003).

The optimal investment expression (26) also allows some simple comparative statics to be carried out. First, the optimal dynamic investment displays a number of appealing, conventional properties that are present in simple single-period or myopic models. Looking at the risk aversion parameter γ , we see that the more risk averse the investor, the lower her optimal investment in the risky stock (in absolute terms $|\theta^*|$), with the investment tending to zero for an extremely risk averse investor. Similar conclusions can be drawn on the effects of the stock volatility σ and bond interest rate r on investment behavior. As is commonly assumed in the literature (and also in the applications of Sections 3.2–3.3), suppose that the market price of risk, $(\mu_t - r)/\sigma_t$, is driven by the state variable X_t only, and not by σ_t or r . Under this scenario, higher the stock volatility or bond interest rate, lower the stock investment (in absolute terms), since the stock now becomes less attractive, with the investment monotonically tending to zero for higher levels of volatility or interest rate.

The correlation parameter ρ captures the extent of market incompleteness in the economy. When the market is incomplete, hedging against fluctuations in the investment opportunities is complex since it may affect the variability of terminal wealth. The implications of this effect are addressed in Section 2.3. The correlation parameter also affects the joint probability distribution of the stock and state variable under which the expressions in (26) are evaluated. This indirect correlation effect will be assessed in the applications studied in Section 3. Finally, note that the quantitative effect of the hedging demand due to the state variable is directly driven by the correlation parameter. Clearly, this effect is higher for the case of complete markets, $\rho = \pm 1$, and disappears for zero correlation, $\rho = 0$. However, with zero correlation, an intertemporal hedging term still arises (second term in (26)) due to the market price of risk possibly being dependent

⁷In particular, the Proposition proves the existence of the optimal policy θ_t^* satisfying the recursive equations (21)–(22) (assuming the expectations and integrals in (32) are well-defined). Moreover, from the Feynman-Kac theorem, the policy is unique in the class of policies such that $\theta_t(\mu_t - r)$ has polynomial growth in the stock price S_t and state variable X_t . This polynomial growth can directly be checked in specific applications that provide explicit closed-form expressions for θ_t^* . Verifying sufficient conditions for optimality are, in general, technically involved (e.g., Korn and Kraft, 2004) and are beyond the scope of this paper. Specifically, for the mean-variance framework, verification does not amount to comparing the value functions of the time-consistent and arbitrary policies, as would be for the standard framework. The reason is that by construction, the value function under the time-consistent mean-variance policy is lower than that under the pre-commitment policy.

on the stock price, consistent with such a term arising when perfectly replicating a payoff by no-arbitrage in complete markets.

Turning to the time-horizon parameter $T-t$, we see that longer time-horizons unambiguously decreases the myopic demand in (26) (in absolute terms). This is because longer horizons imply higher variability of terminal wealth, and hence the investor decreases the risky stock investment. The impact of the time-horizon on the hedging demand is, however, ambiguous. To illustrate this, suppose that both the myopic and hedging demands are positive. When the time-horizon decreases, while the investor's myopic demand increases, her expected portfolio gains are lower, which may decrease the hedging demand. For short time-horizons, the latter effect dominates, and the hedging demand vanishes as the horizon T is reached.

Finally, the optimal investment expression highlights the importance of the market price of risk process, $(\mu_t - r)/\sigma_t$. The myopic demand is increasing in the price of risk, while the effect on the hedging demand is ambiguous. However, its impact on the hedging component becomes less pronounced with shorter time-horizons since the integrals in (26) shrink as the horizon T is approached. The effect on the hedging demand also depends on whether anticipated portfolio gains become more or less sensitive to the stock and state variable as the market price of risk increases. However, this effect can be disentangled in some applications for which the expectation under measure P^* can be explicitly computed. For constant market price of risk, the optimal mean-variance policy reduces to the myopic demand expression, which is identical to the policy that would be obtained under CARA preferences.

Remark 1 (Recovering time-consistent objective functions). It is of interest to see whether there are time-consistent objective functions leading to our dynamically optimal investment policy (26). In our Markovian economy, it turns out to be possible to recover a time-consistent, increasing, concave, state-dependent objective function that implies the same optimal portfolio policy as our dynamically optimal one. In particular, we consider the following dynamic optimization problem involving a state-dependent objective function of CARA form:

$$\max_{\theta_t} E_t \left[-\varepsilon_T e^{-\gamma W_T} \right] \quad (27)$$

with ε following the process

$$d\varepsilon_t = -\frac{\gamma^2}{2} \left(\left(\frac{\mu_t - r}{\gamma \sigma_t} \right)^2 + (1 - \rho^2) \nu_t \left(\frac{\partial f_t}{\partial X_t} \right)^2 \right) \varepsilon_t dt$$

and f_t given by (23), subject to the budget constraint (3). Applying dynamic programming to this problem one can derive an HJB equation and verify that the value function is given by $J_t = \exp(-\gamma(W_t e^{r(T-t)} + f_t))$, and that the optimal investment policy coincides with (26). To understand the intuition behind the process ε_t we observe (from the optimal wealth (28)) that $d\varepsilon_t = -(\gamma^2/2)\varepsilon_t \text{var}_t[dW_t^* e^{r(T-t)}]$, and hence, an investor with the state-dependent utility (27) puts higher weight on those states of the economy in which the optimal wealth process is less volatile along its path. We finally note that there are other time-consistent, state-dependent

objective functions leading to the optimal policy (26).⁸

Remark 2 (Game-theoretic interpretation of optimal policies). Our methodology until now has employed the traditional dynamic programming approach to portfolio choice. However, the problem of finding the time-consistent mean-variance investment policy has an intra-personal game-theoretic interpretation, similar to that in the literature on consumer behavior under hyperbolic discounting (e.g., Harris and Laibson, 2001). In particular, the investor, unable to precommit, takes the investment policy of her future selves as given and reacts to them in an optimal way. Thus, her investment policy emerges as the outcome of a pure-strategy Nash equilibrium in this game.

In particular, consider a game with a continuum of players (selves) $[0, T]$. Each player $t \in [0, T]$ at time t is guided by the mean-variance criterion (6) over terminal wealth, and chooses a time- t Markovian investment strategy $\theta(W_t, S_t, X_t, t)$ subject to the budget constraint (3). Thus, the players impose an externality on each other by affecting the terminal wealth. Denote by $J(W_t, S_t, X_t, t)$ player t 's value function when all players $s \geq t$ follow the equilibrium strategies $\theta^*(W_t, S_t, X_t, t)$. Then, a pure-strategy Nash equilibrium of the game is defined as follows.

Definition: The set of strategies $\{\theta_t^*, t \in [0, T]\}$ constitutes a pure-strategy Nash equilibrium in the intra-personal game with the mean-variance objective if θ_t^* is an optimal response of player t to the strategies θ_s^* of players $s > t$ – that is, taking θ_s^* as given, θ_t^* solves the dynamic optimization problem (12).

It is straightforward to see that the set of strategies $\{\theta_t^*, t \in [0, T]\}$ remains an equilibrium in any subgame of this game, thus comprising a subgame-perfect pure-strategy Nash equilibrium. Moreover, the equilibrium strategy θ_t^* is characterized by the recursive equation for the optimal policy (16), which is now interpreted as the optimal response function of player t to the actions θ_s^* of other players. The equilibrium strategy then coincides with the closed-form expression for the optimal investment policy (26).

2.3. Further Properties of Optimal Policy

In this Section, we discuss further properties of the mean-variance optimizer's optimal behavior by providing explicit expressions for her terminal wealth, its moments, and her value function. We particularly focus on the implications of market incompleteness. Towards this, it is convenient to employ the decomposition $w_{Xt} = \rho w_t + \sqrt{1 - \rho^2} \tilde{w}_t$, where \tilde{w}_t is a Brownian motion independent of w_t , and so \tilde{w}_t represents the unhedgeable source of risk in the economy. In the sequel, the

⁸In particular, using the results of Lemma 1 and Proposition 1, it is straightforward to see that the dynamically optimal policy (26) can be obtained by solving a time-consistent instantaneous problem of the form

$$\max_{\theta_t} E_t[d(W_t e^{r(T-t)})] - \frac{\gamma}{2} \text{var}_t[d(W_t e^{r(T-t)}) + df_t]$$

with f_t given by (23), subject to budget constraint (3). Interestingly, this objective becomes myopic in the restrictive special case of deterministic anticipated portfolio gains f_t .

effect of market incompleteness on terminal wealth is identified via the \tilde{w} terms.

Proposition 3. *The optimal terminal wealth, its mean, variance and the value function of a dynamic mean-variance optimizer are given by*

$$W_T^* = W_t e^{r(T-t)} + f_t + \frac{1}{\gamma} \int_t^T \frac{\mu_s - r}{\sigma_s} dw_s + \sqrt{1 - \rho^2} \int_t^T \nu_s \frac{\partial f_s}{\partial X_s} d\tilde{w}_s, \quad (28)$$

$$\text{var}_t[W_T^*] = \frac{1}{\gamma^2} E_t \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] + (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 \left(\frac{\partial f_s}{\partial X_s} \right)^2 ds \right], \quad (29)$$

$$E_t[W_T^*] = W_t e^{r(T-t)} + f_t, \quad (30)$$

$$J_t = W_t e^{r(T-t)} + f_t - \frac{1}{2\gamma} E_t \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] - \frac{\gamma}{2} (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 \left(\frac{\partial f_s}{\partial X_s} \right)^2 ds \right], \quad (31)$$

where $f_t = E_t^* \left[\int_t^T (\mu_s - r)^2 / \gamma \sigma_s^2 ds \right]$ and $d\tilde{w}_t = (dw_{Xt} - \rho dw_t) / \sqrt{1 - \rho^2}$. Consequently, under the assumption that the market price of risk $(\mu_t - r) / \sigma_t$ depends only on X_t ,

- (i) *The variance of terminal wealth in incomplete markets is higher than that in complete markets;*
- (ii) *The mean of terminal wealth is increasing (decreasing) in the level of market incompleteness, ρ^2 , when the hedging demand is positive (negative) for all $s \in [t, T]$;*
- (iii) *The value function in incomplete markets is lower than that in complete markets when the hedging demand is positive for all $s \in [t, T]$. The effect is ambiguous when the hedging demand is negative.*

Optimal terminal wealth is given by conditionally riskless terms (first and second in (28)), capturing anticipated bond and stock gains, and risky terms driven by the hedgeable stock uncertainty (third term in (28)) and unhedgeable uncertainty (fourth term in (28)). The effect of market incompleteness on terminal wealth enters through the unhedgeable risk and the joint probability distribution of stock return and state variable, both under the original and new measures. The unhedgeable risk component vanishes in complete markets with $\rho^2 = 1$.

The variance of optimal terminal wealth is determined by the variances of hedgeable (first term in (29)) and unhedgeable uncertainties (second term in (29)). When the market price of risk depends only on the state variable, market incompleteness does not affect the hedgeable uncertainty variance. In that case, the terminal wealth variance is always higher in incomplete markets than in complete markets by the presence of unhedgeable risk (Proposition 3(i)). Naturally, this effect is more pronounced for higher state variable volatility or sensitivity of anticipated gains to the state variable. However, the anticipated gains process, f , itself depends on the correlation ρ , which convolutes the exact dependence of wealth variance on correlation, and hence market

completeness.⁹ This indirect effect can be disentangled in the applications, where the expectation under measure P^* can explicitly be computed.

The effect of market incompleteness on the mean of terminal wealth enters via the anticipated gains, f . Proposition 3(ii) states that the direction of this effect is determined by the sign of the hedging demand. In particular, the expected terminal wealth is lower for higher levels of market incompleteness (i.e., lower ρ^2) when the hedging demand is positive till the horizon and the market price of risk depends only on the state variable. The reason is that lower correlation ρ decreases the hedging demand, which vanishes for zero correlation, as discussed in Section 2.2. So, the investor's positive hedging demand will be lower for higher levels of market incompleteness, leading to lower expected terminal wealth. Clearly, the converse is true when the hedging demand is negative.

Turning to the value function, we find that when the hedging demand is positive until the horizon, the mean-variance optimizer is worse off in incomplete markets due to higher variance and lower expectation of terminal wealth. However, the welfare effect is ambiguous in the case of negative hedging demand, for which the expected wealth is higher in incomplete markets, offsetting the effect of higher variance. As will be shown in Section 3, the sign of the hedging demand can readily be identified in particular applications, simplifying the analysis in incomplete markets.

2.4. Optimal Pre-commitment Policy

In Section 2.2, we have already demonstrated that the mean-variance objective in a dynamic setting results in time-inconsistency of the investment policy, in that an investor has an incentive to deviate from an initial policy at a later date. We have so far focused on the time-consistent investment policy in which the investor chooses an investment in each period that maximizes her objective at that period, taking into account the re-adjustments that she will make in the future. We now analyze the alternative way of dealing with this issue and look at the *pre-commitment* investment policy in which the investor initially chooses a policy to maximize her objective function at time 0, and thereafter does not deviate from that policy. Of course, with standard utility functions and absent market imperfections, the solutions to the time-consistent and pre-commitment formulations are well-known to coincide (Cox and Huang, 1989; Karatzas, Lehoczky and Shreve, 1987). The pre-commitment solution, in our view, serves as a useful benchmark against which to compare our time-consistent solution, especially because the explicit analytical solutions to the dynamic mean-variance problem so far have been obtained only in the pre-commitment case. Moreover, if there were a credible mechanism for the investor to commit to her initial policy, she would be better off to follow her initial policy than the time-consistent policy, since the dynamic time-consistency requirement restricts her to consider only policies that she would not be willing to deviate from.

⁹This is due to the fact that the drift of the state variable is affected by the correlation ρ under measure P^* , as revealed by equation (22).

The pre-commitment mean-variance problem and its variations have been analyzed in the literature, amongst others, by Bajoux-Besnainou and Portrait (1998), Bielecki, Jin, Pliska and Zhou (2005), Cvitanic and Zapatero (2004), Zhao and Li (2000), Zhao and Ziemba (2002). These works have primarily employed martingale methods in a complete market setting. For completeness, we here provide the pre-commitment solution for our setting, and follow the literature by specializing to a complete-market setting, $\rho = \pm 1$. Portfolio choice problems that employ martingale methods in incomplete markets are well-known to be a daunting task. However, we can illustrate our main points in the simple complete market setting.

Dynamic market completeness allows the construction of a unique state price density process, ξ , consistent with no-arbitrage, and given by

$$\xi_t = \xi_0 e^{-rt - \frac{1}{2} \int_0^t \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds - \int_0^t \frac{\mu_s - r}{\sigma_s} dw_s}. \quad (32)$$

The quantity $\xi_T(\omega)$ can be interpreted as the Arrow-Debreu price per unit probability P of one unit of wealth in state $\omega \in \Omega$ at time T , and without loss of generality, we set $\xi_0 = 1$. The dynamic investment problem of an investor can be restated as a static variational problem using the martingale representation approach (Cox and Huang, 1989; Karatzas, Lehoczky and Shreve, 1987). Accordingly, a mean-variance optimizer under pre-commitment solves the following problem at time 0:

$$\max_{W_T} E_0[W_T] - \frac{\gamma}{2} \text{var}_0[W_T], \quad (33)$$

$$\text{subject to } E_0[\xi_T W_T] \leq W_0. \quad (34)$$

Proposition 4 presents the optimal solution to this problem in terms of the state price density.

Proposition 4. *The optimal terminal wealth of a mean-variance optimizer under pre-commitment is given by*

$$\hat{W}_T = W_0 e^{rT} + \frac{1}{\gamma} E_0[\xi_T^2] e^{2rT} - \frac{1}{\gamma} \xi_T e^{rT}. \quad (35)$$

Furthermore, under the assumption of a constant market price of risk $(\mu_t - r)/\sigma_t \equiv (\mu - r)/\sigma$, the pre-committed investor's optimal terminal wealth and investment policy are given by

$$\hat{W}_T = W_0 e^{rT} + \frac{1}{\gamma} e^{(\frac{\mu-r}{\sigma})^2 T} - \frac{1}{\gamma} \xi_T e^{rT}, \quad (36)$$

$$\hat{\theta}_t = \frac{\mu - r}{\gamma \sigma^2} e^{-r(T-t)} \xi_t e^{(\frac{\mu-r}{\sigma})^2 (T-t) + rt}. \quad (37)$$

To facilitate comparisons with the pre-commitment solution above, we also provide the time-consistent solution (Propositions 2–3). In the special case of a complete market, the time-consistent optimal terminal wealth, expressed in terms of the state price density, is given by

$$W_T^* = W_0 e^{rT} + \frac{1}{\gamma} E_0 \left[\xi_T e^{rT} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] - \frac{1}{\gamma} \left[\ln \xi_T + rT + \frac{1}{2} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right]. \quad (38)$$

Under the additional assumption of constant market price of risk, the time-consistent optimal terminal wealth and investment policy are

$$W_T^* = W_0 e^{rT} + \frac{1}{\gamma} \left(\frac{\mu - r}{\sigma} \right)^2 T - \frac{1}{\gamma} \left[\ln \xi_T + rT + \frac{1}{2} \left(\frac{\mu - r}{\sigma} \right)^2 T \right], \quad (39)$$

$$\theta_t^* = \frac{\mu - r}{\gamma \sigma^2} e^{-r(T-t)}. \quad (40)$$

Clearly, the pre-commitment solution ((35)–(37)) and the time-consistent solution ((38)–(40)) are generically different. The two solutions coincide only in the knife-edge case of a zero market price of risk, in which case both the pre-commitment and time-consistent policies entail investing nothing in the stock and putting all wealth in the bond. We observe that for short investment horizons T , the pre-commitment solution approximates the time-consistent one up to second-order terms. However, for plausible horizons, the two solutions can differ considerably. In particular, for constant market price of risk case, the pre-commitment expected terminal wealth is higher than the time-consistent one for sufficiently long investment horizons.¹⁰ This is because the inability to pre-commit destroys investors welfare. While for short time-horizons the effect of time-inconsistency can be negligible, it is amplified at longer time-horizons. Moreover, since the state price density is positive, it can easily be observed from (36) that the terminal wealth under the pre-commitment policy is bounded from above. In contrast, the time-consistent policy retains the intuitive property that the terminal wealth can become arbitrarily large for sufficiently small state prices when the cost of wealth is low.¹¹

The pre-commitment policy is stochastic, driven by the state-price density, even under the assumed constant investment opportunity set, while the time-consistent investment is deterministic. Being stochastic, the investment policy under pre-commitment induces a hedging demand component, which is amplified at longer horizons. As a result, with longer horizons, the pre-committed investor tends to invest more in the risky stock than the time-consistent investor does. Finally, in bad states (high ξ) the pre-committed mean-variance optimizer increases her risky investment, and in good states decreases investments. This is because bad, costly states reduce her expected terminal wealth. To offset this, the investor takes on more risk by increasing her risky investment.

¹⁰To see this, observe that the expected wealth under pre-commitment (36) grows exponentially with the horizon, while the expected wealth under time-consistency (39) grows linearly. Even though the variance is also higher in the pre-commitment case, it can be verified that the time-0 probability $\text{Prob}_0(\tilde{W}_T > W_T^*)$ approaches unity with long horizons.

¹¹This property holds for any utility function satisfying the condition $\lim_{W_T \rightarrow \infty} u'(W_T) = 0$ since the marginal utility $u'(W_T)$ is proportional to the state price density ξ_T at the optimum. In particular, one can easily demonstrate that for CARA utility with absolute risk aversion parameter γ the optimal terminal wealth is unbounded and is given by $W_T = -(1/\gamma) \ln \xi_T - (1/\gamma) \ln(\lambda/\gamma)$, where λ is a constant, similarly to (39).

3. Applications

This Section provides several applications that illustrate the simplicity and the usefulness of the methodology developed in Section 2. In Sections 3.1–3.3, we consider the portfolio choice problem of a mean-variance optimizer for different stochastic investment opportunity sets. We obtain explicit solutions to these problems and provide further insights, disentangling some effects that cannot be analyzed in the general framework. We also assess the economic significance of the intertemporal hedging demands of a mean-variance optimizer by quantitatively comparing them with the total demand in the richer economic settings of Sections 3.2–3.3.

3.1. Constant Elasticity of Variance

In this Section, we specialize our setting to a complete market and the constant elasticity of variance (CEV) model for the stock price:

$$\frac{dS_t}{S_t} = \mu dt + \bar{\sigma} S_t^{\alpha/2} dw_t, \quad (41)$$

where α is the elasticity of instantaneous stock return variance, $\sigma_t^2 = \bar{\sigma}^2 S_t^\alpha$, with respect to the stock price. This process is a generalization of geometric Brownian motion, which corresponds to $\alpha = 0$, and has been successfully employed in the option pricing literature (e.g, Cox and Ross, 1976; Schroder, 1989; Cox, 1996) to model the empirically observed pattern of stock prices with heavy tails. Moreover, the CEV process with $\alpha < 0$ generates the finding that the volatility increases when the stock price falls (Black, 1976; Beckers, 1980). When $\alpha < 0$, the distribution of stock prices has the left tail heavier than the right one, while the converse is true for $\alpha > 0$. The CEV model also helps explain volatility smiles (Cox, 1996).

The mean-variance investor's optimal policy under the CEV setting can be computed explicitly by a straightforward application of Proposition 2. It amounts to computing the anticipated gains process under measure P^* , which coincides with the familiar risk-neutral one due to market completeness. We then derive the anticipated gains by computing the expectation of the squared market price of risk, which after some manipulation, reduces to solving an ordinary linear differential equation for which we obtain the unique explicit solution. We emphasize that this computation does not resort to solving an HJB PDE, as it would be the case under other popular objective functions, such as CRRA or CARA preferences. The following Corollary to Proposition 2 presents the optimal investment policy, as well as some of its properties.

Corollary 1. *The optimal stock investment policy for the CEV model (41) is given by:*

$$\theta_t^* = \frac{\mu - r}{\gamma \bar{\sigma}^2 S_t^\alpha} e^{-r(T-t)} - \frac{1}{\gamma} \left(\frac{\mu - r}{\bar{\sigma} S_t^{\alpha/2}} \right)^2 \frac{e^{-\alpha r(T-t)} - 1}{r} e^{-r(T-t)}. \quad (42)$$

Consequently,

- (i) The hedging demand is positive (negative) for $\alpha > 0$ ($\alpha < 0$), and vanishes for $\alpha = 0$;
- (ii) The optimal investment policy θ_t^* is a quadratic function of the market price of risk, $(\mu - r)/\bar{\sigma}S_t^{\alpha/2}$, and may become negative for large values of market price of risk when $\alpha < 0$;
- (iii) The optimal investment policy tends to 0 for $\alpha > -1$ and to $-\infty$ for $\alpha < -1$, as the time-horizon, $T - t$, increases.

Corollary 1(i) reveals that the sign of the hedging demand (the second term in (42)) depends on the sign of the elasticity α . Positive elasticity implies that the market price of risk decreases in the stock price. This induces a negative correlation between the stock returns and anticipated portfolio gains (given by (23)) since the latter are positively related to the market price of risk. As discussed in Section 2.2, this gives rise to a positive hedging demand. Analogously, the hedging demand is negative for negative elasticity.

Property (ii) of Corollary 1 sheds light on the impact of the market price of risk on the optimal investment policy. The optimal investment policy is a quadratic function of the market price of risk for a given stock volatility. Moreover, with negative hedging demand the investor may short the stock despite a high market price of risk or risk premium. In such a case, an increase in the market price of risk leads to a proportionally larger increase in anticipated gains. This then implies a larger covariance between stock returns and portfolio gains, making the hedging demand larger than the myopic demand in absolute terms, and hence the negative stock investment.

Turning to the horizon effect, property (iii) reveals that the optimal investment tends to either zero or negative infinity as the time-horizon increases. There are two effects working in opposite directions. On one hand, the investment is perceived riskier at longer horizons which induces the investor to invest less in the stock. On the other hand, the anticipated gains are higher with longer horizons which makes the hedging demand larger. Anticipated gains are determined by the expected squared market price of risk under measure P^* , which can be verified to stay bounded for positive elasticities as the horizon increases and explodes otherwise (Appendix, proof of Corollary 1). As a result, the first effect dominates for relatively high elasticities ($\alpha > -1$), while the second dominates for relatively low elasticities ($\alpha < -1$). The two effects exactly offset each other for the knife-edge case of $\alpha = -1$, for which the policy tends to a constant.

3.2. Stochastic Volatility

We now consider an incomplete market setting in which the stock price follows the stochastic-volatility model of Liu (2001):

$$\frac{dS_t}{S_t} = (r + \delta X_t^{(1+\beta)/2\beta})dt + X_t^{1/2\beta}dw_t, \quad (43)$$

where the state variable, X , follows a mean-reverting square-root process

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}\sqrt{X_t}dw_{xt}, \quad (44)$$

and where $\beta \neq 0$ is the elasticity of the market price of risk, $\delta\sqrt{X_t}$, with respect to instantaneous stock return volatility, $\sigma_t = X_t^{1/2\beta}$, and $\lambda > 0$ (to exclude explosive processes).

In this setting, Liu derives an explicit solution to the portfolio choice problem for an investor with CRRA preferences over terminal wealth. The case of $\beta = -1$ corresponds to the stochastic-volatility model employed by Chacko and Viceira (2005), who study the intertemporal consumption and portfolio choice problem for an investor with recursive preferences over intermediate consumption. They obtain an exact solution to the problem for investors with unit elasticity of intertemporal substitution of consumption. The case of $\beta = 1$ reduces to the stochastic-volatility model of Heston (1993), popular in option pricing.

Our mean-variance investor's dynamic optimal policy is again a straightforward, simple application of Proposition 2. Since the squared market price of risk equals $\delta^2 X_t$, explicitly finding the solution amounts to computing the conditional expectation of the state variable under measure P^* , which is easily seen (second equation in (22)) to also follow a mean-reverting, square-root process as in (44). The conditional expectation of such a process is well-known (e.g., Cox, Ingersoll and Ross, 1985). In contrast, the solution method of Liu is based on the derivation of the HJB equation for the investor's value function. However, in the case of CRRA preferences this approach is cumbersome for two reasons. First, it involves guessing the value function and reducing the HJB to a system of ODE, one of which is a Riccati equation. Second, this system of equations itself is notorious for complexity. Corollary 2 reports our solution and some of its properties.

Corollary 2. *The optimal stock investment policy for the stochastic-volatility model (43)–(44) is given by:*

$$\theta_t^* = \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} e^{-r(T-t)} - \rho\bar{\nu}\delta \left(\frac{1 - e^{-(\lambda+\rho\bar{\nu}\delta)(T-t)}}{\lambda + \rho\bar{\nu}\delta} \right) \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} e^{-r(T-t)}. \quad (45)$$

Consequently,

- (i) *The hedging demand is positive (negative) for $\rho < 0$ ($\rho > 0$) and vanishes for $\rho = 0$;*
- (ii) *The optimal investment policy θ_t^* is positive (negative) for positive (negative) stock risk premium;*
- (iii) *The optimal investment policy is increasing (decreasing) in the market price of risk, $\delta\sqrt{X_t}$, for $\beta < 0$ or $\beta > 1$ ($0 < \beta < 1$) when the stock risk premium is positive, and the converse is true when the stock risk premium is negative;*
- (iv) *The optimal investment policy tends to 0 for $\lambda + \rho\bar{\nu}\delta > -r$ and to ∞ for $\lambda + \rho\bar{\nu}\delta < -r$, as the time-horizon, $T - t$, increases;*

(v) *The expected terminal wealth, $E_t[W_T^*]$, is decreasing in the correlation ρ . The variance of terminal wealth, $\text{var}_t[W_T^*]$, attains its minimum when the market is complete, $\rho^2 = 1$, and its maximum for some $\rho^* < 0$. The value function, J_t , is decreasing in ρ on the interval $[-1, \rho^*]$ and ambiguous otherwise.*

Corollary 2(i) shows that the sign of the hedging demand (second term in (45)) is determined by the sign of the correlation between the stock and state variable. When this correlation is negative, the instantaneous stock returns are negatively correlated with anticipated portfolio gains since the latter are positively related to the squared market price of risk, $\delta^2 X_t$. As discussed in Section 2.2, such a negative correlation with anticipated gains induces a positive hedging demand. Analogously, a positive correlation ρ gives rise to a negative hedging demand.

Property (ii) of Corollary 2 reveals that the mean-variance optimizer always holds a long position in a risky stock with positive risk premium, as in static or myopic portfolio choice problems.¹² In contrast, Liu (2001) finds that a CRRA investor with low risk aversion may short the risky stock even for a high positive risk premium. Moreover, the mean-variance investment policy is increasing in the market price of risk for negative ($\beta < 0$) or relatively high ($\beta > 1$) elasticities of market price of risk with respect to stock volatility when the stock risk premium is positive (property (iii)). With a negative elasticity, the market price of risk is high when the stock volatility is low that makes the stock attractive. For high elasticities, high market price of risk is associated with a high volatility. However, since the elasticity is high, an increase in the market price of risk offsets an increase in the stock volatility making the stock attractive. Conversely, for intermediate elasticities ($0 < \beta < 1$), the optimal investment decreases in the market price of risk.

Property (iv) also shows that the optimal investment either vanishes or explodes as the time-horizon increases. This horizon effect depends on the covariance between the stock returns and state variable per unit of stock volatility, $\rho\bar{v}$, amplified by the risk premium scale parameter, δ . With positive correlation and high state variable volatility, hedging demand is small and vanishes with long horizons. Otherwise, increasing the stock investment would lead to higher anticipated gains and higher variability of terminal wealth amplified by longer time-horizon. Conversely for sufficiently negative correlation.

Corollary 2(v) also sheds further light on the effect of market incompleteness on wealth and welfare. First, the expected terminal wealth is decreasing in the correlation between the stock and state variable. For negative correlation, it decreases since the hedging demand is positive, and becomes smaller as the correlation approaches zero. For positive correlation, expected terminal wealth declines since the hedging demand is negative, and becomes larger in absolute terms as the correlation approaches unity. In congruence with Proposition 3, the variance of terminal wealth is lowest in the complete market case ($\rho^2 = 1$), in which perfect hedging is possible,

¹²The optimal investment, however, may become negative for negative speed of mean-reversion λ , which corresponds to an explosive process for the state variable.

Table 1
Percentage Hedging Demand over Total Demand
for the Stochastic-Volatility Model with elasticity $\beta = -1$

The table reports the percentage hedging demand over total demand for different levels of correlation ρ , speed of mean-reversion λ and time-horizon $T - t$. The other pertinent parameters are fixed at their estimated values. The relevant model parameter values are taken from Chacko and Viceira (2005, Table 1) who estimate the stochastic-volatility model with elasticity $\beta = -1$ using U.S. stock market data based on monthly returns from 1928 to 2000 and annual returns from 1871 to 2000. Panel A reports our results for (annualized) parameter values $\rho = 0.5241$, $\bar{\nu} = 0.6503$, $\delta = 0.0811$ and $\lambda = 0.3374$ based on estimates from the monthly data of 1926–2000. Panel B reports the results for parameter values $\rho = 0.3688$, $\bar{\nu} = 1.1703$, $\delta = 0.0848$ and $\lambda = 0.0438$ based on annual data of 1871–2000. Both tables also report results for varying levels of ρ and λ , with bolded ratios corresponding to the estimated parameter values of ρ and λ .

Panel A: Monthly Data Parameter Estimates						Panel B: Annual Data Parameter Estimates				
ρ	Horizon					6-month	1-year	5-year	10-year	20-year
	6-month	1-year	5-year	10-year	20-year					
-1.00	2.4	4.39	12.3	14.9	15.6	4.8	9.3	36.4	57.0	78.4
-0.50	1.2	2.2	6.23	7.5	7.8	2.4	4.7	20.1	33.8	51.3
0.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.37	-1.0	-1.7	-4.8	-5.6	-5.8	-1.8	-3.6	-17.7	-33.6	-57.3
0.50	-1.2	-2.3	-6.5	-7.6	-7.8	-2.5	-5.0	-24.7	-47.6	-81.5
0.52	-1.3	-2.4	-6.8	-8.0	-8.2	-2.6	-5.2	-26.0	-50.3	-86.2
1.00	-2.5	-4.6	-13.1	-15.3	-15.6	-5.0	-10.2	-54.9	-111.8	-189.1
λ	6-month	1-year	5-year	10-year	20-year	6-month	1-year	5-year	10-year	20-year
0.00	-1.4	-2.8	-14.8	-31.8	-73.8	-1.9	-3.7	-20.1	-44.2	-107.9
0.04	-1.4	-2.7	-13.2	-24.6	-41.7	-1.8	-3.6	-17.7	-33.6	-57.3
0.30	-1.3	-2.4	-7.3	-8.8	-9.2	-1.7	-3.2	-9.7	-11.7	-12.2
0.34	-1.3	-2.4	-6.8	-8.0	-8.2	-1.7	-3.2	-9.0	-10.6	-10.8
0.60	-1.2	-2.1	-4.4	-4.6	-4.6	-1.6	-2.8	-5.8	-6.1	-6.1
0.90	-1.1	-1.8	-3.0	-3.1	-3.1	-1.5	-2.4	-4.0	-4.1	-4.1

and attains a maximum at some intermediate correlation level $\rho^* < 0$. Thus, for relatively low correlation ($\rho < \rho^*$) the expected wealth is decreasing in correlation while the variance is increasing, which leads the welfare to decrease in correlation. The welfare effect is ambiguous for relatively high correlation ($\rho > \rho^*$) since lower expected wealth is counterbalanced by decreased variance. However, for plausible parameter values (Table 1, Panel A: $\rho = 0.5241$, $\bar{\nu} = 0.6503$, $\delta = 0.0811$, $\lambda = 0.3374$), it can be shown that the loss in expected wealth dominates, and hence the welfare decreases in correlation.

Finally, we investigate the economic significance of the mean-variance intertemporal hedging demands induced by the stochastic volatility setting. To this end, we compute the ratio of the hedging demand to total optimal demand, θ_{Ht}/θ_t^* , for a range of plausible parameter values. Conveniently, this ratio is deterministic and depends only on the correlation ρ , the state variable speed of mean-reversion λ and volatility parameter $\bar{\nu}$ for a given time-horizon. Table 1 presents the percentage hedging demand for varying levels of correlation, speed of mean-reversion and the

investor's horizon.¹³ The relevant parameter values are taken from Chacko and Viceira (2005), who estimate the stochastic-volatility model with elasticity $\beta = -1$ using U.S. stock market data based on monthly returns from 1926 to 2000 and annual returns from 1871 to 2000.

Inspection of the results in Panel A of Table 1, based on monthly data parameter estimates, reveals a relatively small ratio of the hedging demand over total demand, ranging from -1.3% to -8.2% for the parameter estimates $\rho = 0.52$ and $\lambda = 0.34$ (in bold). This small magnitude of the hedging demand is due to the relatively low correlation and high speed of mean-reversion estimates. In contrast, Panel B, based on annual data parameter estimates, reveals a considerably larger percentage hedging demand, ranging from -1.8% to -57.3% for the parameter estimates $\rho = 0.37$ and $\lambda = 0.04$. Our results are in line with the findings of Chacko and Viceira, although in absolute terms they are large. Chacko and Viceira find the percentage hedging demand to range from -1.5% to -3.6% for the monthly data and from -5.2% to -18.4% for the yearly data for an infinitely-lived recursive-utility investor with relative risk aversion and elasticity of intertemporal substitution ranging $[1.5, 40]$ and $[1/0.8, 1/40]$, respectively.¹⁴ Thus, the hedging demand in our setting is larger in absolute terms than in Chacko and Viceira.

3.3. Time-Varying Gaussian Mean Returns

In this Section, we consider the mean-variance optimizer's problem in an incomplete market in which the stock price dynamics are specialized to follow:

$$\frac{dS_t}{S_t} = (r + \sigma X_t)dt + \sigma dw_t, \quad (46)$$

where the market price of risk, X_t , follows a mean-reverting Ornstein-Uhlenbeck process

$$dX_t = \lambda(\bar{X} - X_t)dt + \nu dw_{xt}, \quad (47)$$

with $\lambda > 0$. Kim and Omberg (1996) explicitly solves the portfolio choice problem of an investor with CRRA preferences over terminal wealth in this incomplete market setting. Merton (1971) studies the consumption and portfolio choice problem of an agent with CARA preferences in this Gaussian mean-reverting setting for the special complete-market case of positive perfect correlation, $\rho = +1$. Wachter (2002) provides an explicit solution to the consumption and portfolio choice problem of an investor with CRRA preferences under this setting with negative perfect correlation, $\rho = -1$. Campbell and Viceira (1999) study the infinite-horizon discrete-time consumption and portfolio choice of an investor with recursive utility and under discrete-time

¹³We do not consider varying the levels of $\bar{\nu}$ and δ since they always appear multiplicatively with the correlation ρ in the hedging and total demand expressions. So, separately varying the levels of $\bar{\nu}$ and δ would lead to a range of percentage hedging demand similar to that generated by different levels of ρ .

¹⁴Chacko and Viceira compute the ratio of hedging demand to myopic demand, which we then convert into the ratio of hedging demand to total demand. Moreover, they also consider the case of the relative risk aversion being less than unity, which is less empirically plausible for the average investor. In that case, they find that the ratio of hedging demand to myopic demand is positive.

versions of the dynamics (46)–(47), where the state variable is taken to be the dividend-price ratio.

For the mean-variance optimizer, finding the optimal investment policy again reduces to computing the expectation of the squared market price of risk, X_t^2 , under measure P^* . It follows from (22) that under this measure, the market price of risk follows a simple mean-reverting process as in (47) for which the first and second moments can easily be derived (e.g., Vasicek, 1977). This approach avoids solving the HJB equation which is a tedious task in the case of CRRA preferences and incomplete markets since it amounts to solving a system of nonlinear ordinary differential equations (e.g., Kim and Omberg, 1996). The following Corollary to Proposition 2 reports our mean-variance solution and some of its unambiguous properties.

Corollary 3. *The optimal stock investment policy for the time-varying Gaussian mean returns model (46)–(47) is given by:*

$$\theta_t^* = \frac{X_t}{\gamma\sigma} e^{-r(T-t)} - \frac{\rho\nu}{\gamma\sigma} \left(\lambda \left(\frac{1 - e^{-(\lambda+\rho\nu)(T-t)}}{\lambda + \rho\nu} \right)^2 \bar{X} + \frac{1 - e^{-2(\lambda+\rho\nu)(T-t)}}{\lambda + \rho\nu} X_t \right) e^{-r(T-t)}. \quad (48)$$

Consequently,

- (i) *The mean hedging demand is positive (negative) for $\rho < 0$ ($\rho > 0$) and vanishes for $\rho = 0$ when $\bar{X} > 0$, and the converse is true when $\bar{X} < 0$;*
- (ii) *The optimal stock investment, θ_t^* , is increasing in the market price of risk, X_t .*

The hedging demand (second term in (48)) in general may become positive or negative for any combination of model parameters depending on the sign and magnitude of the market price of risk X_t , which is Gaussian and can possibly take on negative values. The mean hedging demand, the unconditional expectation of the hedging demand, however, is positive for negative correlation between the stock and state variable, and negative for positive correlation. The intuition for this is as in the stochastic-volatility model of the previous Section (Corollary 2(i)).

Corollary 3(ii) reveals the optimal investment policy to be increasing in the market price of risk. Thus, our dynamic mean-variance optimizer under the mean-reverting Gaussian setting retains this familiar property of the myopic or static portfolio choices despite a potentially large hedging demand, as demonstrated below. However, the welfare implications of the market incompleteness for this setting are complicated due to the fact that the hedging demand may change signs over time depending on the behavior of the market price of risk, but can explicitly be analyzed for a given set of model parameters.

We here assess the significance of the intertemporal hedging demands by computing the ratio of the mean hedging demand over the mean total demand, as in Campbell and Viceira (1999).¹⁵

¹⁵The ratio of the hedging demand over the total demand is stochastic and depends on the state variable X_t . Therefore, as a tractable quantitative assessment of the percentage hedging demand, we follow Campbell and Viceira and consider the mean hedging demand over the mean total demand, which is deterministic.

Table 2
Percentage Mean Hedging Demand over Mean Total Demand
for the Mean-Reverting Gaussian Returns Model

The table reports the percentage mean hedging demand over mean total demand for different levels of correlation ρ , speed of mean-reversion λ and the investor's time-horizon $T - t$. The other pertinent parameters are fixed at their estimated values. The relevant parameter values are taken from the estimates provided in Wachter (2002, Table 1). These parameter estimates are based on their discrete-time analogues in Barberis (2000) and Campbell and Viceira (1999), and are: $\rho = -0.93$, $\nu = 0.065$ and $\lambda = 0.27$. The table also reports results for varying levels of ρ and λ , with bolded ratios corresponding to the estimated parameter values of ρ and λ .

ρ	Horizon				
	6-month	1-year	5-year	10-year	20-year
-1.00	19.1	32.5	71.0	81.6	87.4
-0.93	17.9	30.7	68.2	78.7	84.4
-0.50	10.0	17.9	45.5	54.2	58.1
0.00	0.0	0.0	0.0	0.0	0.0
0.50	-11.2	-22.2	-96.5	-135.1	-143.5
1.00	-23.8	-50.1	-427.7	-920.6	-1020.9
λ	6-month	1-year	5-year	10-year	20-year
0.00	19.0	34.4	87.9	98.5	100.0
0.27	17.9	30.7	68.2	78.7	84.4
0.30	17.7	30.3	66.1	75.8	80.6
0.60	16.6	26.9	48.7	51.6	52.0
0.90	15.5	23.6	37.3	38.0	38.0

This ratio depends only on the correlation ρ , the speed of mean-reversion λ and the instantaneous variance of the state variable ν for a given time-horizon $T - t$. Table 2 reports the percentage mean hedging demand for varying levels of correlation, speed of mean-reversion parameter and the investor's horizon. The parameter values are taken from the estimates provided by Wachter (2002) and are described in the caption of Table 2.

Inspection of Table 2 establishes the percentage mean hedging demand over mean total demand to be positive and fairly large, ranging from 17.9% to 84.4% for the parameter estimates $\rho = -0.93$ and $\lambda = 0.27$ (in bold). This result is primarily due to the large negative correlation ρ , which implies (on average) a positive and large hedging demand. Our finding is consistent with that reported in the literature under a similar economic setting but with different investor preferences. Campbell and Viceira (1999) find the percentage mean hedging demand to range from 22.9% to 65.5% for an infinitely lived recursive-utility investor with relative risk aversion and elasticity of intertemporal substitution ranging $[1.5, 40]$ and $[1/0.75, 1/40]$, respectively. Results in Brandt (1999) confirm the findings of Campbell and Viceira for the case of CRRA preferences with relative risk aversion 5. A large hedging demand in proportion to wealth is also reported in Wachter.

4. Extensions and Ramifications

In this Section, we demonstrate that the baseline analysis of Section 2 can easily be adopted to alternative or richer economic environments. Section 4.1 illustrates our methodology in a discrete-time framework, and provides an explicit solution to the stochastic-volatility model in discrete time. Sections 4.2 – 4.3 demonstrate that the results of Section 2 are readily extendable to more realistic environments with stochastic interest rates and with multiple stocks, state variables and sources of uncertainty.

4.1. Discrete-Time Formulation

We consider the mean-variance asset allocation problem in a discrete-time setting. The extant literature, to our best knowledge, lacks analytic expressions for multi-period discrete-time investment policies in rich stochastic environments and characterizes optimal policies by employing either numerical methods or various approximations (e.g., Ait-Sahalia and Brandt, 2001; Bansal and Kiku, 2007; Brandt, Goyal, Santa-Clara and Stroud, 2005; Brandt and Santa-Clara, 2006; Campbell and Viceira, 1999, 2002, among others). In contrast, we here derive a recursive representation for the optimal investment policy in discrete time and provide fully-explicit closed-form solutions for specific stochastic investment opportunity sets as in the continuous-time formulation. To our knowledge, these explicit solutions are new in the literature.¹⁶

We let the time increment denote $\Delta t \equiv T/M$, where M is an integer number, and index time by $t = 0, \Delta t, 2\Delta t, \dots, T$. The uncertainty is generated by two correlated discrete-time stochastic processes w and w_X , with correlation ρ . The increments of the processes, Δw_t and Δw_{Xt} , are serially uncorrelated and distributed according to some distribution with zero mean and variance $\Delta t, D(0, \Delta t)$. An investor trades in two securities, a riskless bond that provides a constant interest rate r over the interval Δt , and a risky stock that has price dynamics given by

$$\frac{\Delta S_t}{S_t} = \mu(S_t, X_t, t)\Delta t + \sigma(S_t, X_t, t)\Delta w_t,$$

where the state variable X follows the process

$$\Delta X_t = m(X_t, t)\Delta t + \nu(X_t, t)\Delta w_{Xt}.$$

An investor's wealth W then follows

$$\Delta W_t = [rW_t + \theta_t(\mu_t - r)]\Delta t + \theta_t\sigma_t\Delta w_t, \tag{49}$$

where θ_t again denotes the dollar stock investment. The investor maximizes the objective function (6) subject to the dynamic budget constraint (49) for each time $t = 0, \Delta t, \dots, T - \Delta t$. Proposition

¹⁶Since the purpose of this Section is to demonstrate the tractability of our analysis in discrete time, we employ a simple Euler discretization scheme and abstract away from potential issues of convergence of our discrete-time stochastic processes to their continuous-time analogues.

5 is the discrete-time analogue of Proposition 1 and provides a recursive representation for the optimal investment policy in terms of the anticipated portfolio gains, $f_t = E_t[W_T^*] - W_t R^{T-t}$. The proof is similarly based on deriving the Bellman equation in discrete-time. Not surprisingly though, since the anticipated gains process cannot be represented in differential form in discrete-time, the optimal policy is characterized not in terms of partial derivatives of f , but in terms of its time- t conditional covariance with one-period stock returns.

Proposition 5. *The optimal stock investment policy of a dynamic mean-variance optimizer in discrete-time is given by*

$$\theta_t^* = \frac{\mu_t - r}{\gamma \sigma_t^2} R^{-(T-\Delta t-t)} - \frac{\text{cov}_t(\Delta S_t/S_t, \Delta f_t)}{\sigma_t^2 \Delta t} R^{-(T-\Delta t-t)}, \quad (50)$$

where process f_t represents the expected total gains or losses from the stock investment and is given by

$$f(S_t, X_t, t) = E_t \left[\sum_{s=t}^{T-\Delta t} \theta_s^* (\mu_s - r) R^{(T-\Delta t-s)} \Delta t \right], \quad (51)$$

$R = (1 + r\Delta t)^{1/\Delta t}$ and $t = 0, \Delta t, \dots, T - \Delta t$.

The discrete-time optimal investment policy has the same structure as in Proposition 1 and is given by myopic and hedging demands. The absence of a discrete-time version of the Feynman-Kac formula, however, does not allow us to characterize the optimal policy entirely in terms of the exogenous model parameters, as in Proposition 2. Nevertheless, expression (50) can be used to obtain an explicit representation for the optimal policy for specific applications either by solving (50) backwards or by guessing the structure of the solution.

To illustrate an application of Proposition 5, we solve the discrete-time versions of the models of Sections 3.2–3.3. Specifically, the discrete-time dynamics of the stock price and state variable for the stochastic-volatility model are specified as follows:

$$\frac{\Delta S_t}{S_t} = (r + \delta X_t^{(1+\beta)/2\beta}) \Delta t + X_t^{1/2\beta} \Delta w_t, \quad (52)$$

$$\Delta X_t = \lambda(\bar{X} - X_t) \Delta t + \bar{\nu} \sqrt{X_t} \Delta w_{Xt}, \quad (53)$$

where $\beta \neq 0$ and $\lambda > 0$. In discrete time, there is a probability that X_t hits the zero-boundary even with a non-explosive process. To exclude this, we assume that either the interval Δt is so small that this event has negligible probability or the distribution function of Δw_t and Δw_{Xt} is truncated in such a way that it never happens. To obtain the optimal investment policy explicitly, we first conjecture that the solution has the form $\theta_t^* = g(t) X_t^{(\beta-1)/2\beta} R^{-(T-\Delta t-t)}$, where $g(t)$ is a deterministic function. Substituting this expression into the recursive representation (50) gives a recursive equation for the function $g(t)$ that can be solved explicitly.

The discrete-time dynamics of the stock price and state variable for the time-varying Gaussian mean returns model are given by:

$$\frac{\Delta S_t}{S_t} = (r + \sigma X_t) \Delta t + \sigma \Delta w_t, \quad (54)$$

$$\Delta X_t = \lambda(\bar{X} - X_t)\Delta t + \nu\Delta w_{Xt}, \quad (55)$$

where $\lambda > 0$. Campbell and Viceira (1999) consider a discrete-time version of these dynamics and derive optimal policies under recursive utility by employing log-linear approximations. To obtain an explicit solution we conjecture that it has the form $\theta_t^* = X_t/\gamma\sigma - (g_1(t) + g_2(t)X_t)/\gamma\sigma$, where $g_1(t)$ and $g_2(t)$ are deterministic functions. Substituting θ_t^* into representation (50), as in the previous case, we obtain recursive equations for $g_1(t)$ and $g_2(t)$ which we solve explicitly. Corollary 4 reports the results.

Corollary 4. *The optimal investment policy for the discrete-time stochastic-volatility model (52)–(53) is given by*

$$\theta_t^* = \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} R^{-(T-\Delta t-t)} - \rho\bar{\nu}\delta \frac{1 - (1 - (\lambda + \rho\bar{\nu}\delta)\Delta t)^{(T-\Delta t-t)/\Delta t}}{\lambda + \rho\bar{\nu}\delta} \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} R^{-(T-\Delta t-t)}, \quad (56)$$

and for the discrete-time model with Gaussian mean-returns (54)–(55) is given by

$$\theta_t^* = \frac{X_t}{\gamma\sigma} R^{-(T-\Delta t-t)} - \frac{g_1(t) + g_2(t)X_t}{\gamma\sigma} R^{-(T-\Delta t-t)}, \quad (57)$$

where

$$\begin{aligned} g_1(t) &= (A + B) \left(1 - [(1 - \lambda\Delta t)(1 - \rho\nu\Delta t)]^{(T-\Delta t-t)/\Delta t} \right) \\ &\quad - B \left(1 - [(1 - \lambda\Delta t)^2(1 - 2\rho\nu\Delta t)]^{(T-\Delta t-t)/\Delta t} \right), \\ g_2(t) &= \left(1 - \frac{(1 - (1 - \lambda\Delta t)^2)(1 + 2\rho\nu\lambda\Delta t)}{1 - (1 - \lambda\Delta t)^2(1 - 2\rho\nu\Delta t)} \right) \left(1 - [(1 - \lambda\Delta t)^2(1 - 2\rho\nu\Delta t)]^{(T-\Delta t-t)/\Delta t} \right), \end{aligned}$$

and A and B are constants, explicitly reported in the Appendix.

It can be verified that as time interval Δt approaches zero, the discrete-time policies converge to the continuous-time ones reported in Corollaries 2 and 3. As a result, the comparative statics for (56) and (57) are similar to those in the continuous-time case. We note that in deriving expressions (56) and (57), we do not assume normality of the stochastic processes w and w_X , as in continuous-time.

4.2. Multiple Stock Formulation

We now generalize the baseline analysis of Section 2 with a single stock and state variable to the case of multiple stocks and state variables. Specifically, uncertainty is generated by two multi-dimensional Brownian motions $w = (w_1, \dots, w_N)^\top$ and $w_X = (w_{X1}, \dots, w_{XK})^\top$ with $N \times K$ correlation matrix ρ , where each element of the matrix $\rho = \{\rho_{nm}\}$ represents the correlation between the Brownian motions w_n and w_{Xm} . An investor trades in a riskless bond with a

constant interest rate r and N risky stocks, and so the market is again potentially incomplete. The stock prices, $S = (S_1, \dots, S_N)^\top$, follow the dynamics

$$\frac{dS_{it}}{S_{it}} = \mu_i(S_t, X_t, t)dt + \sigma_i(S_t, X_t, t)^\top dw_t, \quad i = 1, \dots, N,$$

where μ_i and σ_i are deterministic functions of S and K state variables, $X = (X_1, \dots, X_K)^\top$, which satisfy

$$dX_{jt} = m_j(X_t, t)dt + \nu_j(X_t, t)^\top dw_{Xt}, \quad j = 1, \dots, K.$$

We let $\mu \equiv (\mu_1, \dots, \mu_N)^\top$ denote the vector of stock mean returns and $\sigma \equiv (\sigma_1, \dots, \sigma_N)^\top$ the volatility matrix, assumed invertible, with each component $\sigma = \{\sigma_{in}\}$ capturing the covariance between the stock return and Brownian motion w_n . Similarly, $m \equiv (m_1, \dots, m_K)^\top$ and $\nu \equiv (\nu_1, \dots, \nu_K)^\top$ will denote the mean growth and the volatility matrix of the state variables X , respectively. The investor's wealth follows

$$dW_t = [rW_t + \theta_t^\top (\mu_t - r)]dt + \theta_t^\top \sigma_t dw_t, \quad (58)$$

where $\theta_t = (\theta_{1t}, \dots, \theta_{Nt})^\top$ denotes the vector of dollar investments in the N stocks at time t .

The dynamic optimization problem of the investor is as in Section 2. For each time $t \in [0, T]$, she maximizes the time- t objective function (6) subject to the dynamic budget constraint (58). As in Section 2, the optimal policy is characterized in terms of the anticipated portfolio gains, f , and arises from the HJB equation adjusted for time-inconsistency. Proposition 6 generalizes Proposition 2 and reports the optimal investment and anticipated gains in terms of the model parameters and the hedge-neutral measure.

Proposition 6. *The optimal investment policy in the multiple-stock economy is given by*

$$\theta_t^* = \frac{1}{\gamma} (\sigma_t \sigma_t^\top)^{-1} (\mu_t - r) e^{-r(T-t)} - \left(I_{S_t} \frac{\partial f_t}{\partial S_t^\top} + (\nu_t \rho^\top \sigma_t^{-1})^\top \frac{\partial f_t}{\partial X_t^\top} \right) e^{-r(T-t)}, \quad (59)$$

where I_{S_t} is a diagonal $N \times N$ matrix with S_{1t}, \dots, S_{Nt} on the main diagonal, $\partial f_t / \partial S_t$ and $\partial f_t / \partial X_t$ denote the row-vectors of partial derivatives with respect to relevant variables. The anticipated portfolio gains, f , can be represented as

$$f(S_t, X_t, t) = E_t^* \left[\int_t^T \frac{1}{\gamma} (\mu_s - r)^\top (\sigma_s \sigma_s^\top)^{-1} (\mu_s - r) ds \right],$$

where $E_t^*[\cdot]$ denotes the expectation under the unique hedge-neutral measure P^* on which are defined N -dimensional Brownian motion w^* and K -dimensional Brownian motion w_X^* with correlation ρ , given by

$$dw_t^* = dw_t + \sigma_t^{-1} (\mu_t - r) dt, \quad dw_{Xt}^* = dw_{Xt} + \rho^\top \sigma_t^{-1} (\mu_t - r) dt,$$

and measure P^* is defined by the Radon-Nikodym derivative

$$\frac{dP^*}{dP} = e^{-\frac{1}{2} \int_0^T (\mu_s - r)^\top (\sigma_s \sigma_s^\top)^{-1} (\mu_s - r) ds - \int_0^T (\sigma_s^{-1} (\mu_s - r))^\top dw_s}.$$

The optimal investment policy (59) is given by myopic and intertemporal hedging terms, retaining the structure of the single-stock case. It can again be shown that the hedging demands can be expressed in terms of the covariance of stock returns and anticipated portfolio gains. Proposition 6 also identifies the effect of cross-correlations on the optimal investment and reveals that the hedging term for one stock depends on the correlations of other stocks with the state variables. The optimal investment expression also allows for some simple comparative statics with respect to the risk aversion parameter, interest rate and stock volatility matrix with similar implications to those in Section 2. We can also obtain expressions for optimal terminal wealth, its moments and the value function of the mean-variance optimizer and identify the effect of market incompleteness, as in Section 2.3.

4.3. Stochastic Interest Rates

In this Section, we incorporate stochastic interest rates into our analysis and demonstrate that the optimal policies can explicitly be computed as in the baseline model of Section 2. Specifically, we consider an incomplete-market economy with an additional source of uncertainty generated by a Brownian motion w_r that is correlated with Brownian motions w and w_X with correlations ρ_{rs} and ρ_{rX} , respectively. The locally riskless bond now has a stochastic interest rate r that follows the dynamics

$$dr_t = \mu_r(X_t, r_t, t)dt + \sigma_r(X_t, r_t, t)dw_{rt}, \quad (60)$$

where μ_r and σ_r are deterministic functions of X and r . Furthermore, we allow the stock price and state variable parameters μ , σ , m and ν to additionally depend on the interest rate r .

In our analysis we take the bond as the numeraire so that all relevant quantities are in terms of the bond price $B_t = B_0 e^{\int_0^t r_s ds}$, as is common in various problems in finance. To facilitate tractability, we employ the mean-variance criterion over terminal wealth in units of this numeraire, which allows us to adopt our earlier solution method and characterize the optimal policy in units of the numeraire, that is, $\tilde{\theta}_t \equiv \theta_t^*/B_t$.¹⁷ Proposition 7 reports our results.

Proposition 7. *The optimal investment policy in the economy with stochastic interest rates is given by*

$$\tilde{\theta}_t = \frac{\mu_t - r_t}{\gamma\sigma_t^2} - \left(S_t \frac{\partial f_t}{\partial S_t} + \frac{\rho\nu_t}{\sigma_t} \frac{\partial f_t}{\partial X_t} + \frac{\rho_{rs}\sigma_{rt}}{\sigma_t} \frac{\partial f_t}{\partial r_t} \right), \quad (61)$$

where f_t is as in Proposition 2, but with r following (60).

The optimal policy (61) has the same structure as the baseline case. The main difference is that the hedging term now additionally accounts for the interest rate fluctuations by incorporating the sensitivity of anticipated portfolio gains (f) to interest rates. As in Section 3, the optimal

¹⁷Otherwise, if the mean-variance criterion is over W_T and the interest rate is stochastic, in contrast to Lemma 1 the value function is not separable in W_t and the policy θ_t^* is no longer independent of W_t , which makes the problem intractable.

policies may explicitly be computed for various stochastic investment opportunities. We consider a simple application where all the fluctuations in the investment opportunities are driven by the stochastic interest rate r . In particular, the stock price follows a geometric Brownian motion with constant parameters μ and σ , while the interest rate follows a Vasicek model (Vasicek, 1977)

$$dr_t = \lambda_r(\bar{r} - r_t)dt + \sigma_r dw_{rt}. \quad (62)$$

Along the lines of Corollaries 1–3, it can be demonstrated that the optimal policy is given by

$$\tilde{\theta}_t = \frac{\mu - r_t}{\gamma\sigma^2} - \frac{\rho_{rS}\sigma_r}{\gamma\sigma} \left(\lambda_r \left(\frac{1 - e^{-(\lambda_r - \rho_{rS}\sigma_r/\sigma)(T-t)}}{\lambda_r - \rho_{rS}\sigma_r/\sigma} \right)^2 \frac{\mu - \bar{r}}{\sigma} + \frac{1 - e^{-2(\lambda_r - \rho_{rS}\sigma_r/\sigma)(T-t)}}{\lambda_r - \rho_{rS}\sigma_r/\sigma} \frac{\mu - r_t}{\sigma} \right).$$

This policy is comparable to that of the case of time-varying Gaussian mean-returns (48) in Section 3.3, but now additionally allows us to consider comparative statics with respect to the parameters of the interest rate dynamics (62).

5. Conclusion

Despite the popularity of the mean-variance criteria in multi-period problems in finance, little is known about the dynamically optimal mean-variance portfolio policies. This work makes a step in this direction by providing a fully analytical characterization of the optimal mean-variance policies within a familiar, dynamic, incomplete-market setting. The optimal mean-variance dynamic portfolios are shown to have a simple, intuitive and tractable structure. The solution is obtained via dynamic programming and is facilitated by deriving a recursive formulation for the mean-variance criteria, accounting for its time-inconsistency. We also identify a “hedge-neutral” measure that absorbs intertemporal hedging demands and allows explicit computation of optimal portfolios in a straightforward way for various stochastic environments. Given the tractability offered by our analysis, we believe that our results are well suited for various applications in financial economics. In concurrent work, we investigate the hedging strategies of non-replicable claims in incomplete markets according to the minimum-variance criterion – the related “mean-variance hedging” literature has had limited success in providing explicit solutions to this problem. We also foresee potential applications in security pricing with incomplete markets, for which the investor preferences are to be accounted for.

Appendix: Proofs

Proof of Lemma 1. The HJB equation in differential form (12) follows from equation (9) when the decision making interval, τ , tends to zero. To derive the terminal condition for J_T , we note that $\text{var}_T[W_T] = 0$ and $E_T[W_T] = W_T$. The definition of the value function, J_T , then implies $J_T = W_T$.

To show that W_t does not affect θ_t^* , using Itô's Lemma we rewrite the budget constraint (3) as

$$d(W_t e^{r(T-t)}) = \theta_t(\mu_t - r)e^{r(T-t)}dt + \theta_t\sigma_t e^{r(T-t)}dw_t, \quad (\text{A.1})$$

integrate from t to T and substitute W_T into the time- t objective function:

$$\begin{aligned} E_t[W_T] - \frac{\gamma}{2}\text{var}_t[W_T] &= W_t e^{r(T-t)} + E_t \left[\int_t^T \theta_s(\mu_s - r)e^{r(T-s)}ds \right] \\ &\quad - \frac{\gamma}{2}\text{var}_t \left[\int_t^T \theta_s(\mu_s - r)e^{r(T-s)}ds + \int_t^T \theta_s\sigma_s e^{r(T-s)}dw_s \right]. \end{aligned} \quad (\text{A.2})$$

It can be observed from (A.2) that the objective function is separable in $W_t e^{r(T-t)}$, and hence the optimal policy θ_s^* does not depend on W_t for $s \geq t$. Since the investor solves for the investment policy by backwards induction, θ_s^* also does not depend on W_s for $s > t$. Due to the Markovian nature of the economy, θ_t^* depends only on S_t , X_t and t . The fact that the function f_t depends only on S_t , X_t and t follows from the expression for f_t in terms of the optimal policy, given in (11). The separability of the value function J_t from $W_t e^{r(T-t)}$ follows from (A.2). *Q.E.D.*

Proof of Proposition 1. To prove Proposition 1, it remains to derive the first order condition for the problem (15). The objective function in (15) is quadratic and concave in θ_t , and so the unique optimal policy solves the first order condition:

$$(\mu_t - r)e^{r(T-t)} - \gamma\theta_t^*\sigma_t^2 e^{2r(T-t)} - \gamma\sigma_t \left(\sigma_t S_t \frac{\partial f_t}{\partial S_t} + \rho\nu_t \frac{\partial f_t}{\partial X_t} \right) e^{r(T-t)} = 0,$$

leading to the expression (16). *Q.E.D.*

Proof of Proposition 2. Under standard conditions, there exists a probability measure P^* under which the function f_t admits the Feynman-Kac representation (21) (Karatzas and Shreve, 1991) and under this measure, the processes S and X satisfy the stochastic differential equations (22). Comparing (22) with (1)–(2), we obtain that measure P^* transforms Brownian motions w_t and w_{Xt} into w_t^* and w_{Xt}^* satisfying (24).

We next find the Radon-Nikodym derivative dP^*/dP . To apply Girsanov's Theorem (Karatzas and Shreve, 1991), we first decompose the Brownian motion w_X as a sum of two uncorrelated Brownian motions: $dw_{Xt} = \rho dw_t + \sqrt{1 - \rho^2} d\tilde{w}_t$, where $\tilde{w}_t = (w_{Xt} - \rho w_t)/\sqrt{1 - \rho^2}$. We observe that in terms of $d\tilde{w}_t$, the representations (24) can be rewritten as follows:

$$dw_t^* = dw_t + \frac{\mu_t - r}{\sigma_t} dt, \quad dw_{Xt}^* = \rho \left(dw_t + \frac{\mu_t - r}{\sigma_t} dt \right) + \sqrt{1 - \rho^2} d\tilde{w}_t.$$

Since measure P^* affects only the first component of the two-dimensional Brownian motion $(w_t, \tilde{w}_t)^\top$, the Radon-Nikodym derivative (25) obtains by Girsanov's Theorem. Finally, substituting f_t given by (23) into the recursive representation (16), we obtain (26). *Q.E.D.*

Proof of Proposition 3. First, we derive the terminal wealth expression. Substituting the optimal policy θ_t^* from (16) into (A.1) and rearranging terms we obtain

$$d(W_t e^{r(T-t)}) = -df_t + \frac{\mu_t - r}{\gamma \sigma_t} dw_t + \sqrt{1 - \rho^2} \nu_t \frac{\partial f_t}{\partial X_t} d\tilde{w}_t, \quad (\text{A.3})$$

where \tilde{w}_t is defined in Proposition 3. Integrating (A.3) from t to T , we obtain (28). Since \tilde{w}_t and w_t are uncorrelated, the variance of terminal wealth is given by

$$\text{var}_t[W_T^*] = \text{var}_t \left[\frac{1}{\gamma} \int_t^T \frac{\mu_s - r}{\sigma_s} dw_s \right] + \text{var}_t \left[\sqrt{1 - \rho^2} \int_t^T \nu_s \frac{\partial f_s}{\partial X_s} d\tilde{w}_s \right],$$

which leads to expression (29). The expressions for $E_t[W_T^*]$ and J_t are immediate.

We next prove assertions (i)–(iii). Property (i) follows from the wealth variance expression (29). Since $(\mu_t - r)/\sigma_t$ is assumed to not depend on S_t , the first term in (29) does not depend on the correlation ρ . The second term is strictly positive in incomplete markets and vanishes in complete markets, $\rho^2 = 1$, and hence the assertion. To prove property (ii) we compute the derivative of f_t with respect to correlation ρ in terms of the hedging demand. Since $(\mu_t - r)/\sigma_t$ depends only on X_t , f_t in (23) also depends only on X_t . As a result, the PDE (20) for f_t becomes:

$$\frac{\partial f_t}{\partial t} + \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} \right) \frac{\partial f_t}{\partial X_t} + \frac{\nu_t^2}{2} \frac{\partial^2 f_t}{\partial X_t^2} + \frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma_t} \right)^2 = 0, \quad (\text{A.4})$$

with $f_T = 0$. Differentiating (A.4) with respect to ρ and denoting $\tilde{f}_t \equiv \partial f_t / \partial \rho$, we obtain the equation for \tilde{f}_t :

$$\frac{\partial \tilde{f}_t}{\partial t} + \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} \right) \frac{\partial \tilde{f}_t}{\partial X_t} + \frac{\nu_t^2}{2} \frac{\partial^2 \tilde{f}_t}{\partial X_t^2} - \nu_t \frac{\mu_t - r}{\sigma_t} \frac{\partial f_t}{\partial X_t} = 0, \quad (\text{A.5})$$

where $\tilde{f}_T = 0$. Applying the Feynman-Kac Theorem to equation (A.5), using the expression (18) for θ_{Ht} and the fact that f_t does not depend on S_t we obtain:

$$\frac{\partial f_t}{\partial \rho} = \frac{1}{\rho} E_t^* \left[\int_t^T \theta_{Hs} (\mu_s - r) e^{r(T-s)} ds \right].$$

As a result, if $\theta_{Hs} > 0$ for $s \in [t, T]$, function f_t is increasing (decreasing) in ρ when ρ is positive (negative). This is equivalent to saying that f_t is increasing in ρ^2 . Conversely, if $\theta_{Hs} < 0$ for $s \in [t, T]$, f_t is decreasing in ρ^2 .

The proof of Assertion (iii) follows from (i) and (ii). If the hedging demand is positive over the horizon, the expected terminal wealth is lower in incomplete markets, while the variance is higher. As a result, the value function is unambiguously lower in this case. *Q.E.D.*

Proof of Proposition 4. The optimal pre-commitment terminal wealth \hat{W}_T solves the first order condition

$$1 - \gamma\hat{W}_T + \gamma E_0[\hat{W}_T] - \psi\xi_T = 0, \quad (\text{A.6})$$

where ψ is the Lagrange multiplier of the static budget constraint (34). Taking time-zero expectation on both sides of (A.6) yields $\psi = 1/E_0[\xi_T]$, or $\psi = e^{rT}$, since $E_0[\xi_T e^{rT}] = 1$. Substituting ψ back into (A.6) we obtain

$$\hat{W}_T = \frac{1}{\gamma} \left(1 + \gamma E_0[\hat{W}_T] - \xi_T e^{rT} \right). \quad (\text{A.7})$$

(A.7) substituted into the static budget constraint (34) leads to $\gamma E_0[\hat{W}_T] = \gamma W_0 e^{rT} - 1 + E_0[\xi_T^2] e^{2rT}$, which along with (A.7) yields the optimal terminal wealth (35).

With a constant market price of risk, $(\mu - r)/\sigma$, $E_0[\xi_T^2] = e^{-2rT + (\frac{\mu-r}{\sigma})^2 T}$ leading to (36). To compute the pre-commitment investment policy, $\hat{\theta}_t$, we first consider the optimal time- t wealth:

$$\hat{W}_t = E_t \left[\frac{\xi_T}{\xi_t} \hat{W}_T \right] = a(t) - \frac{1}{\gamma} e^{-(2r - (\mu-r)^2/\sigma^2)(T-t)} e^{rT} \xi_t, \quad (\text{A.8})$$

where the second equality follows by substituting \hat{W}_T from (36) and evaluating the moments of ξ_T , and $a(t)$ is a deterministic function of time. Applying Itô's Lemma to (A.8) and using $d\xi_t = -\xi_t[r dt + (\mu - r)/\sigma dw_t]$ yields:

$$d\hat{W}_t = (a'(t) - b(t)\xi_t)dt + \frac{\mu - r}{\gamma\sigma} e^{-(2r - (\mu-r)^2/\sigma^2)(T-t)} e^{rT} \xi_t dw_t,$$

where $b(t)$ is a time-deterministic function. Matching the coefficients with the dynamic budget constraint (3) yields $\hat{\theta}_t$ in (37). *Q.E.D.*

Proof of Corollary 1. In the case of a complete market, measure P^* coincides with the risk-neutral one. To compute the optimal investment policy from Proposition 2, we need to evaluate the expected squared market price of risk, $E_t^*[(\mu_s - r)^2/\sigma_s^2]$, under the risk-neutral measure. Since the squared market price of risk in the CEV model is $(\mu - r)^2/(\bar{\sigma}^2 S_t^\alpha)$, we need to determine $g(t, s) \equiv E_t^*[S_s^{-\alpha}]$ for $s > t$. By Itô's Lemma, the process for $S_t^{-\alpha}$ under the risk-neutral measure satisfies:

$$dS_t^{-\alpha} = \left(-\alpha r S_t^{-\alpha} + \frac{\alpha(1 + \alpha)\bar{\sigma}^2}{2} \right) dt - \alpha \bar{\sigma} S_t^{-\alpha} dw_t^*. \quad (\text{A.9})$$

Integrating (A.9) from t to s and taking the time- t expectation under the risk-neutral measure on both sides, we obtain the equation for $g(t, s)$:

$$g(t, s) = S_t^{-\alpha} - \int_t^s \left(\alpha r g(t, y) - \frac{\alpha(1 + \alpha)\bar{\sigma}^2}{2} \right) dy. \quad (\text{A.10})$$

Differentiating (A.10) with respect to s yields the linear differential equation

$$\frac{\partial g(t, s)}{\partial s} = -\alpha r g(t, s) + \frac{\alpha(1 + \alpha)\bar{\sigma}^2}{2}, \quad (\text{A.11})$$

with initial condition $g(t, t) = S_t^{-\alpha}$. The unique solution to equation (A.11) is given by

$$g(t, s) = S_t^{-\alpha} e^{-\alpha r(s-t)} + (1 + \alpha) \bar{\sigma}^2 \frac{1 - e^{-\alpha r(s-t)}}{2r}. \quad (\text{A.12})$$

Substitution of (A.12) into the optimal investment policy (26) leads to the θ_t^* reported in Corollary 1. We also note that the process for the market price of risk is explosive for $\alpha \leq -1$ since the conditional expectation (A.12) is unbounded for large horizons. For $-1 < \alpha < 0$, the conditional expectation (A.12) is not well-defined since it may become negative for large investment horizons, implying that the process hits the zero-boundary with a positive probability.

Property (i) is immediate from the expression for the optimal investment policy (42). Property (ii) follows from (42) and the fact that the hedging demand is negative for $\alpha < 0$. Finally, property (iii) obtains since the product of exponents in (42) tends to zero (negative infinity) with increasing horizon for $\alpha > -1$ (< -1). *Q.E.D.*

Proof of Corollary 2. Since the squared market price of risk is equal to $\delta^2 X_t$, finding θ_t^* amounts to evaluating $E_t^*[X_s]$. It follows from (22) that the state variable under measure P^* follows the process

$$dX_t = (\lambda + \rho \bar{\nu} \delta) \left(\frac{\lambda \bar{X}}{\lambda + \rho \bar{\nu} \delta} - X_t \right) dt + \bar{\nu} \sqrt{X_t} dw_{Xt}^*,$$

for which the conditional moments are well-known (e.g., Cox, Ingersoll and Ross, 1985), yielding

$$E_t^*[X_s] = \frac{\lambda \bar{X}}{\lambda + \rho \bar{\nu} \delta} + \left(X_t - \frac{\lambda \bar{X}}{\lambda + \rho \bar{\nu} \delta} \right) e^{-(\lambda + \rho \bar{\nu} \delta)(s-t)}. \quad (\text{A.13})$$

Substituting (A.13) into (26) yields the desired result.

Assertion (i) follows from the fact that $(1 - e^{-(\lambda + \rho \bar{\nu} \delta)(T-t)})/(\lambda + \rho \bar{\nu} \delta)$ is always positive. As a result, the sign of the hedging demand (second term in (45)) depends only on the correlation. Assertion (ii) for the case of $\rho < 0$ is immediate from the fact that the hedging demand is positive in this case. For $\rho > 0$, it follows from the fact that $\rho \bar{\nu} \delta (1 - e^{-(\lambda + \rho \bar{\nu} \delta)(T-t)})/(\lambda + \rho \bar{\nu} \delta)$ is less than unity. Property (iii) follows directly from the properties of function $X^{(\beta-1)/\beta}$. Assertion (iv) obtains due to the fact that the product of exponents in the hedging term tends to zero (infinity) with increasing horizon for $\lambda + \rho \bar{\nu} \delta > -r$ ($< -r$).

Turning to property (v), we first prove that f_t decreases in correlation ρ . Since $f_t = \delta^2 \int_t^T E_t^*[X_s] ds$, it remains to show that $E_t^*[X_s]$ decreases in ρ . We observe that by virtue of (A.13), $E_t^*[X_s] = \int_t^s e^{-(\lambda + \rho \bar{\nu} \delta)(y-t)} dy + X_t e^{-(\lambda + \rho \bar{\nu} \delta)(s-t)}$, which is clearly decreasing in correlation ρ . Similarly, using Proposition 3, it can be shown that the variance of terminal wealth can be represented as $\text{var}_t[W_T^*] = (1 - \rho^2)G(\rho)$, where $G(\rho) \equiv E_t[\int_t^T \bar{\nu}^2 X_s (\partial f_s / \partial X_s)^2 ds]$ is a positive decreasing function of ρ . Clearly, the minimum is attained in a complete market with $\rho^2 = 1$. The first order condition for finding the ρ^* at which $\text{var}_t[W_T^*]$ is maximized is $2\rho G(\rho) = (1 - \rho^2)G'(\rho)$. Since the right-hand-side is negative and $G(\rho)$ is positive, the first order condition can only be satisfied for $\rho^* < 0$. *Q.E.D.*

Proof of Corollary 3. Since the squared market price of risk is X_t^2 , finding the optimal investment policy amounts to computing $E_t^*[X_s^2]$, which is well-known (e.g., Vasicek, 1977):

$$E_t^*[X_s^2] = \left(\frac{\lambda \bar{X}}{\lambda + \rho\nu} + \left(X_t - \frac{\lambda \bar{X}}{\lambda + \rho\nu} \right) e^{-(\lambda + \rho\nu)(s-t)} \right)^2 + \nu^2 \frac{1 - e^{-2(\lambda + \rho\nu)(s-t)}}{2(\lambda + \rho\nu)}. \quad (\text{A.14})$$

Substituting (A.14) into (26) yields the reported result.

Property (i) follows from the fact that since the unconditional expectation of the state variable, \bar{X} , is assumed positive, the sign of the mean hedging demand depends only on the sign of the correlation ρ . To show property (ii), we observe that the optimal investment policy can be rewritten as follows:

$$\theta_t^* = \left(1 - \rho\nu \frac{1 - e^{-2(\lambda + \rho\nu)(T-t)}}{\lambda + \rho\nu} \right) \frac{X_t}{\gamma\sigma} e^{-r(T-t)} - \frac{\rho\nu\lambda}{\gamma\sigma} \left(\frac{1 - e^{-(\lambda + \rho\nu)(T-t)}}{\lambda + \rho\nu} \right)^2 \bar{X} e^{-r(T-t)}.$$

Similarly to the proof of Corollary 2(ii), it can be shown that $1 - \rho\nu(1 - e^{-2(\lambda + \rho\nu)(T-t)})/(\lambda + \rho\nu)$ is positive, which then implies that θ_t^* is increasing in the market price of risk X_t . *Q.E.D.*

Proof of Proposition 5. The proof is similar to the proof of Proposition 1. The first step is to derive the Bellman equation adjusted for time-inconsistency in terms of anticipated gains, f . The second step is to derive the first order condition for the strategy θ_t^* . In discrete time, however, the explicit representation for the process for f_t is not available. As a result, the optimal strategy is in terms of $\text{cov}_t(\Delta S_t/S_t, \Delta f_t)$. *Q.E.D.*

Proof of Corollary 4. The first step is to obtain the anticipated gains process f . Substituting the conjecture $\theta_t^* = g(t)X_t^{(\beta-1)/2\beta}R^{-(T-\Delta t-t)}$ for the stochastic-volatility model (52)–(53) into the expression for f_t (51), we obtain:

$$f_t = E_t \left[\sum_{s=t}^{T-\Delta t} \delta X_s g(s) \Delta t \right]. \quad (\text{A.15})$$

To compute $E_t[X_s]$, we take expectations of both sides of the state variable process, (53), and obtain a difference equation for $E_t[X_s]$:

$$E_t[X_{s+\Delta t}] = \lambda \bar{X} \Delta t + (1 - \lambda \Delta t) E_t[X_s], \quad (\text{A.16})$$

with initial condition $E_t[X_t] = X_t$. The unique solution to equation (A.16) is

$$E_t[X_s] = (X_t - \bar{X})(1 - \lambda \Delta t)^{(s-t)/\Delta t} + \bar{X}. \quad (\text{A.17})$$

Substituting (A.17) into (A.15) yields:

$$f_t = d(t) + \delta X_t \sum_{s=t}^{T-\Delta t} g(s)(1 - \lambda \Delta t)^{(s-t)/\Delta t} \Delta t, \quad (\text{A.18})$$

where $d(t)$ denotes a time-deterministic function. Using (A.18), we compute Δf_t and substitute it into the recursive expression for the optimal strategy (50). Taking into account the conjecture for θ_t^* , after some algebra we obtain a recursive equation for $g(t)$:

$$g(t) = \delta/\gamma - \rho\bar{\nu}\delta \sum_{s=t+\Delta t}^{T-\Delta t} g(s)(1 - \lambda\Delta t)^{(s-\Delta t-t)/\Delta t} \Delta t. \quad (\text{A.19})$$

Evaluating (A.19) at time $t - \Delta t$ and then subtracting it from (A.19), we obtain the following forward difference equation for $g(t)$: $g(t - \Delta t) = \lambda\delta\Delta t/\gamma + (1 - (\lambda + \rho\bar{\nu}\delta)\Delta t)g(t)$, with condition $g(T - \Delta t) = \delta/\gamma$. The explicit solution to this equation is

$$g(t) = \frac{\delta}{\gamma} - \rho\bar{\nu}\delta \frac{1 - (1 - (\lambda + \rho\bar{\nu}\delta)\Delta t)^{T-\Delta t-t}}{\lambda + \rho\bar{\nu}\delta} \frac{\delta}{\gamma},$$

which then yields the reported result.

For the case of Gaussian mean-returns dynamics (54)–(55), we first obtain f_t by substituting our conjecture $\theta_t^* = X_t/\gamma\sigma - (g_1(t) + g_2(t)X_t)/\gamma\sigma$ into (51). Then, substituting f_t into (50) we obtain recursive equations for $g_1(t)$ and $g_2(t)$. Solving them as in the previous stochastic-volatility case we obtain $g_1(t)$ and $g_2(t)$, as reported in Corollary 4, with constants A and B explicitly given by

$$\begin{aligned} A &= \frac{\rho\nu(1 - \lambda\Delta t)(2\bar{X} - \varphi\nu^2\lambda\Delta t\sqrt{\Delta t})\Delta t}{1 - (1 - \lambda\Delta t)(1 - \rho\nu\Delta t)} - \varphi\nu^2\lambda\Delta t\sqrt{\Delta t} \\ &\quad - \left(1 - \frac{(1 - (1 - \lambda\Delta t)^2)(1 + 2\rho\nu\lambda\Delta t)}{1 - (1 - \lambda\Delta t)^2(1 - 2\rho\nu\Delta t)}\right) \left(\bar{X} - \frac{\varphi\nu\sqrt{\Delta t}}{2\rho} + \frac{\rho\nu(1 - \lambda\Delta t)(\bar{X} + \varphi\nu\sqrt{\Delta t}/2\rho)\Delta t}{1 - (1 - \lambda\Delta t)(1 - \rho\nu\Delta t)}\right), \\ B &= \left(\frac{\rho\nu(1 - \lambda\Delta t)(\bar{X} + \varphi\nu\sqrt{\Delta t}/2\rho)\Delta t}{(1 - \lambda\Delta t)^2(1 - 2\rho\nu\Delta t) - (1 - \lambda\Delta t)(1 - \rho\nu\Delta t)}\right. \\ &\quad \left. + \bar{X} - \frac{\varphi\nu\sqrt{\Delta t}}{2\rho}\right) \left(1 - \frac{(1 - (1 - \lambda\Delta t)^2)(1 + 2\rho\nu\lambda\Delta t)}{1 - (1 - \lambda\Delta t)^2(1 - 2\rho\nu\Delta t)}\right), \end{aligned}$$

where $\varphi = \text{cov}(\Delta w, \Delta w_x^2)$.¹⁸

Q.E.D.

Proof of Proposition 6. The proof is a multi-dimensional version of the proofs for Propositions 1–2.

Q.E.D.

Proof of Proposition 7. The proof is similar to those of Propositions 1–2, but now accounting for the mean-variance criterion being over W_T/B_T . As in the proof of Lemma 1, substituting the integral representation for W_T/B_T into the criterion we show that $\tilde{\theta}_t$ does not depend on W_t/B_t . Then, as in Section 2 we obtain an HJB equation in terms of df_t and $d(W_t/B_t)$, where $f_t \equiv E_t[W_T^*/B_T] - W_t/B_t$, whose solution yields (61). Employing measure P^* it can be shown that f_t is the same as in Proposition 2, but now with stochastic r_t .

Q.E.D.

¹⁸It can easily be demonstrated that $\varphi = 0$ if Δw and Δw_x are normally distributed.

References

- Acharya, V.V., and L.H. Pedersen, 2005, "Asset Pricing with Liquidity Risk," *Journal of Financial Economics*, 77, 375-410.
- Ait-Sahalia, Y., and M.W. Brandt, 2001, "Variable Selection for Portfolio Choice," *Journal of Finance*, 56, 1297-1351.
- Bajoux-Besnainou, I., and R. Portait, 1998, "Dynamic Asset Allocation in a Mean-Variance Framework," *Management Science*, 44, 79-95.
- Bansal, R., M. Dahlquist, and C.R. Harvey, 2004, "Dynamic Trading Strategies and Portfolio Choice," Working Paper, Duke University.
- Bansal, R., and D. Kiku, 2007, "Cointegration and Long-Run Asset Allocation," Working Paper, Duke University.
- Barberis, N., 2000, "Investing for the Long Run When Returns Are Predictable," *Journal of Finance*, 55, 225-264.
- Beckers, S., 1980, "The Constant Elasticity of Variance Model and Its Implications for Option Pricing," *Journal of Finance*, 35, 661-673.
- Bielecki, T., H. Jin, S.R. Pliska, and X.Y. Zhou, 2005, "Continuous-Time Mean-Variance Portfolio Selection with Bankruptcy Prohibition," *Mathematical Finance*, 15, 213-244.
- Black, F., 1976, "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3, 167-179.
- Bossaerts, P., K. Preuschoff, and S.R. Quartz, 2006, "Neural Differentiation of Expected Reward and Risk in Human Subcortical Structures," *Neuron*, 51, 381-390.
- Bossaerts, P., K. Preuschoff, and S.R. Quartz, 2008, "Markowitz in the Brain?" *Revue d'Economie Politique*, forthcoming.
- Brandt, M.W., 1999, "Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach," *Journal of Finance*, 54, 1609-1645.
- Brandt, M.W., 2004, "Portfolio Choice Problems," in Ait-Sahalia, Y., and L.P. Hansen (eds.), *Handbook of Financial Econometrics*, forthcoming.
- Brandt, M.W., A. Goyal, P. Santa-Clara, and J.R. Stroud, 2005, "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Predictability," *Review of Financial Studies*, 18, 831-873.
- Brandt, M.W., and P. Santa-Clara, 2006, "Dynamic Portfolio Selection by Augmenting the Asset Space," *Journal of Finance*, 61, 2187-2218.
- Brennan, M.J., and Y. Xia, 2002, "Dynamic Asset Allocation under Inflation," *Journal of Finance*, 57, 1201-1238.
- Campbell, J.Y., and L.M. Viceira, 1999, "Consumption and Portfolio Decisions when Expected Returns are Time Varying," *Quarterly Journal of Economics*, 114, 433-495.
- Campbell, J.Y., K. Serfaty-de Medeiros, and L.M. Viceira, 2007, "Global Currency Hedging," Working Paper, Harvard University.
- Campbell, J.Y., and L.M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.
- Chacko, G., and L.M. Viceira, 2005, "Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets," *Review of Financial Studies*, 18, 1369-1402.

- Cochrane, J.H., 2005, "A Mean Variance Benchmark for Intertemporal Portfolio Theory," Working Paper, University of Chicago.
- Cox, J.C., 1996, "The Constant Elasticity of Variance Option Pricing Model," *Journal of Portfolio Management*, 22, 15-17.
- Cox, J.C., and C.-F. Huang, 1989, "Optimal Consumption and Portfolio Policies when Asset Prices follow a Diffusion Process," *Journal of Economic Theory*, 39, 33-83.
- Cox, J.C., J. Ingersoll, and S. Ross, 1985, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385-408.
- Cox, J.C., and S.A. Ross, 1976, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3, 145-166.
- Cvitanic, J., L. Goukasian, and F. Zapatero, 2003, "Monte Carlo Computation of Optimal Portfolios in Complete Markets," *Journal of Economic Dynamics and Control*, 27, 971-986.
- Cvitanic, J., A. Lazrak, and T. Wang, 2007, "Implications of the Sharpe Ratio as a Performance Measure in Multi-Period Setting," *Journal of Economic Dynamics and Control*, forthcoming.
- Cvitanic, J., and F. Zapatero, 2004, *Introduction to the Economics and Mathematics of Financial Markets*, The MIT Press.
- Detemple, J.B., R. Garcia, and M. Rindisbacher, 2003, "A Monte Carlo Method for Optimal Portfolios," *Journal of Finance*, 58, 401-446.
- Duffie, D., and H. Richardson, 1991, "Mean-Variance Hedging in Continuous Time," *Annals of Probability*, 1, 1-15.
- Harris, C., and D. Laibson, 2001, "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 69, 935-957.
- Heston, S.L., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327-343.
- Hong, H., J. Scheinkman, and W. Xiong, 2006, "Asset Float and Speculative Bubbles," *Journal of Finance*, 61, 1073-1117.
- Huang, L., and H. Liu, 2007, "Rational Inattention and Portfolio Selection," *Journal of Finance*, 62, 1999-2040.
- Follmer, H., and D. Sondermann, 1986, "Hedging of Non-Redundant Contingent Claims," in Hildenbrand, W., and A. Mas-Colell (eds.), *Contributions to Mathematical Economics*, 205-233.
- Jagannathan, R., and T. Ma, 2003, "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraint Helps," *Journal of Finance*, 58, 1651-1683.
- Karatzas, I., J.P. Lehoczky, and S.E. Shreve, 1987, "Optimal Portfolio and Consumption Decisions for a Small Investor on a Finite Horizon," *SIAM Journal of Control and Optimization*, 25, 1557-1586.
- Karatzas, I., and S.E. Shreve, 1991, *Brownian Motion and Stochastic Calculus*, Springer-Verlag, New York.
- Karatzas, I., and S.E. Shreve, 1998, *Methods of Mathematical Finance*, Springer-Verlag, New York.

- Kim, T.S., and E. Omberg, 1996, "Dynamic Nonmyopic Portfolio Behavior," *Review of Financial Studies*, 9, 141-161.
- Korn, R., and H. Kraft, 2004, "On the Stability of Continuous-Time Portfolio Problems with Stochastic Opportunity Sets," *Mathematical Finance*, 14, 403-414.
- Leippold, M., F. Trojani, and P. Vanini, 2004, "Geometric Approach to Multiperiod Mean-Variance Optimization of Assets and Liabilities," *Journal of Economic Dynamics and Control*, 28, 1079-1113.
- Li, D., and W.L. Ng, 2000, "Optimal Dynamic Portfolio Selection: Multiperiod Mean-Variance Formulation," *Mathematical Finance*, 10, 387-406.
- Lim, A.E.B., and X.Y. Zhou, 2002, "Mean-Variance Portfolio Selection with Random Parameters in a Complete Market," *Mathematics of Operations Research*, 27, 101-120.
- Liu, J., 2001, "Dynamic Portfolio Choice and Risk Aversion," Working Paper, UCLA.
- Liu, J., 2007, "Portfolio Selection in Stochastic Environments," *Review of Financial Studies*, 20, 1-39.
- Maenhout, P., 2006, "Robust Portfolio Rules and Detection-Error Probabilities for a Mean-Reverting Risk Premium," *Journal of Economic Theory*, 128, 136-163.
- Markowitz, H.M., 1952, "Portfolio Selection," *Journal of Finance*, 7, 77-91.
- Merton, R.C., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373-413.
- Miyahara, Y., 1996, "Canonical Martingale Measures of Incomplete Assets Markets," in Watanabe, S., M. Fukushima, Yu. V. Prohorov, and A.M. Shiryaev (eds.), *Probability Theory and Mathematical Statistics: Proceedings of the Seventh Japan-Russian Symposium*, 343-352.
- Sangvinatsos, A., and J.A. Wachter, 2005, "Does the Failure of the Expectations Hypothesis Matter for Long-Term Investors?" *Journal of Finance*, 60, 179-230.
- Schroder, M., 1989, "Computing the Constant Elasticity of Variance Option Pricing Formula," *Journal of Finance*, 44, 211-219.
- Schweizer, M., 1992, "Mean-Variance Hedging for General Claims," *Annals of Probability*, 2, 171-179.
- Schweizer, M., 1999, "A Minimality Property of the Minimal Martingale Measure," *Statistics and Probability Letters*, 42, 27-31.
- Vasicek, O., 1977, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177-188.
- Wachter, J.A., 2002, "Portfolio and Consumption Decisions under Mean-Reverting Returns: An Explicit Solution for Complete Market," *Journal of Financial and Quantitative Analysis*, 37, 63-91.
- Weiss, N., 2005, *A Course in Probability*, Addison-Wesley.
- Zhao, Y., and W.T. Ziemba, 2002, "Mean-Variance versus Expected Utility in Dynamic Investment Analysis," Working Paper, University of British Columbia.
- Zhou, X.Y., and D. Li, 2000, "Continuous-Time Mean-Variance Portfolio Selection: A Stochastic LQ Framework," *Applied Mathematics and Optimization* 42, 19-53.