

Portfolio Choice and Asset Pricing in Incomplete Markets

by

Georgy Chabakauri

A dissertation submitted to the University of London
for the degree of Doctor of Philosophy

Department of Finance
London Business School
University of London
©June 2009

Declaration

Chapter 1 of this Thesis is entirely my own. Chapters 2 and 3 are coauthored with Professor Suleyman Basak.

Georgy Chabakauri

Acknowledgements

I would especially like to thank my advisor Suleyman Basak for extensive discussions, invaluable comments and support throughout my course of study at London Business School. I am also grateful to Mike Chernov, Victor DeMiguel, Francisco Gomes, Christopher Hennessy, Igor Makarov, Anna Pavlova and Raman Uppal for insightful discussions and comments on my work.

Abstract

This Thesis studies the portfolio choice and asset pricing in economic settings with incomplete markets and is comprised of three chapters. Chapter 1 explores the formation of stock prices in a general equilibrium economy with heterogeneous investors facing portfolio constraints. We compute the equilibrium when both investors have (identical for simplicity) CRRA preferences, one of them is unconstrained while the other faces an upper bound constraint on the proportion of wealth invested in stocks. We demonstrate that under certain parameters the model can generate countercyclical stock return volatilities, procyclical price-dividend ratios, excess volatility and other patterns observed in the data. Our baseline analysis is also extended to models with heterogeneous beliefs.

Chapter 2 studies the portfolio choice in economies with incomplete markets with investors guided by mean-variance criteria. Mean-variance criteria remain prevalent in multi-period problems, and yet not much is known about their dynamically optimal policies. In this chapter we provide a fully analytical characterization of the optimal dynamic mean-variance portfolios within a general incomplete-market economy. 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. A calibration exercise shows that the mean-variance hedging demands may comprise a significant fraction of the total risky asset demand.

Chapter 3 provides a simple solution to the hedging problem in a general incomplete-market economy in which a hedger, guided by the minimum-variance criterion, aims at reducing the risk of a non-tradable asset. We derive fully analytical optimal time-consistent hedges and demonstrate that they can easily be computed in various stochastic environments. Our dynamic hedges preserve the simple structure of complete-market perfect hedges and are in terms of generalized “Greeks,” familiar in risk management applications. We demonstrate that our optimal hedges typically outperform their static and myopic counterparts under plausible economic environments.

Table of Contents

Introduction	10
1. Asset Pricing in General Equilibrium with Constraints	14
1.1. Introduction	14
1.2. General Equilibrium with Constraints	21
1.2.1. Economic Setup	21
1.2.2. Characterization of Equilibrium	23
1.2.3. Further Properties of Equilibrium	28
1.3. Equilibrium with Proportional Constraints	31
1.3.1. Characterization and Computation of Equilibrium	31
1.3.2. Analysis of Equilibrium	39
1.4. Extensions and Ramifications	46
1.4.1. Heterogeneous Beliefs Formulation	46
1.4.2. Multiple Stock Formulation	54
1.5. Conclusion	57
1.6. Appendix A: Proofs	59
1.7. Appendix B: Numerical Method	65
2. Dynamic Mean-Variance Asset Allocation	69
2.1. Introduction	69
2.2. Asset Allocation with Mean-Variance Criteria	75
2.2.1. Economic Setup	75
2.2.2. Determination of Optimal Dynamic Investment Policy	77
2.2.3. Further Properties of Optimal Policy	89

2.2.4. Optimal Pre-commitment Policy	91
2.3. Applications	95
2.3.1. Constant Elasticity of Variance	95
2.3.2. Stochastic Volatility	97
2.3.3. Time-Varying Gaussian Mean Returns	102
2.4. Extensions and Ramifications	105
2.4.1. Discrete-Time Formulation	105
2.4.2. Multiple Stock Formulation	109
2.4.3. Stochastic Interest Rates	111
2.5. Conclusion	112
2.6. Appendix: Proofs	114
3. Dynamic Hedging in Incomplete Markets: A Simple Solution	123
3.1. Introduction	123
3.2. Dynamic Minimum-Variance Hedging	128
3.2.1. Economic Setup	129
3.2.2. Dynamically Optimal Hedging Policy	131
3.2.3. Applications	135
3.2.4. Time-Consistency Conditions	138
3.3. Comparison with Pre-commitment, Static and Myopic Hedges	140
3.3.1. Comparison with Pre-commitment Policy	141
3.3.2. Comparison with Static Policy	144
3.3.3. Comparison with Myopic Policy	147
3.4. Extensions	149

3.4.1. Mean-Variance Hedging and Benchmark Tracking	149
3.4.2. Multi-dimensional case	151
3.5. Conclusion	153
3.6. Appendix: Proofs	155
References	165

List of Tables

1.1	Effective Domains and Support Functions	29
2.1	Percentage Hedging Demand over Total Demand for the Stochastic-Volatility Model with elasticity $\beta = -1$	101
2.2	Percentage Mean Hedging Demand over Mean Total Demand for the Mean-Reverting Gaussian Returns Model	104
3.1	Optimal Dynamic and Static Hedges and Their Relative Performances	146
3.2	Dynamically Optimal and Myopic Policies and Their Relative Performances	148

List of Figures

1.1	Parameters of Equilibrium with Constraints, $\gamma < 1$	38
1.2	Parameters of Equilibrium with Constraints, $\gamma > 1$	40
1.3	Probability Density Functions for Constrained Investor's Share in Aggregate Consumption, $\gamma = 3$	45
1.4	Price-Dividend Ratios and Ratios of Stock Return and Dividend Growth Volatilities with Heterogeneous Beliefs, $\gamma > 1$	52
3.1	Relative Performance of Dynamically Optimal and Pre-commitment Hedges.	143

Introduction

The main objective of this Thesis is to study the optimal portfolio choice and asset pricing in financial markets with imperfections, such as portfolio constraints or unhedgeable risks that render the markets incomplete. It is generally admitted that these imperfections significantly contribute towards understanding the investor behavior and the dynamics of financial markets. Despite their importance in Finance, portfolio choice and asset pricing models are notoriously difficult to solve in incomplete markets. This Thesis develops new methods for solving these models in various incomplete-market settings and provides new economic insights into the formation of optimal portfolios and asset prices. The Thesis consists of three Chapters that explore different aspects of portfolio choice and asset pricing in incomplete financial markets.

Chapter 1 is entitled “Asset Pricing in General Equilibrium with Constraints” and explores the formation of asset prices in a pure exchange general equilibrium economy with two heterogeneous investors facing portfolio constraints, and one consumption good generated by a Lucas tree. Even though the presence of portfolio constraints is well documented in the contemporary literature, it is admittedly difficult to incorporate them into general equilibrium analysis unless constrained investors are logarithmic, as commonly assumed in the literature. However, the assumption of logarithmic investors significantly impedes the evaluation of the impact of portfolio constraints on financial markets. In particular, there is a large literature (e.g., Detemple and Murthy, 1997; Basak and Cuoco, 1998; Basak and Croitoru, 2000, 2006) demonstrating that in economic settings with logarithmic investors various portfolio constraints do not affect stock prices since classical income and substitution effects perfectly offset each other. In Chapter 1 we develop a methodology that does not rely on the restrictive assumption of logarithmic investors and therefore allows us to evaluate the impact of constraints on equilibrium parameters such as interest rates, market prices of risk, price-dividend ratios and stock return volatilities.

In particular, we numerically compute the equilibrium in a setting with two investors

with identical constant relative risk aversion (CRRA) utility functions where one investor is unconstrained while the other faces an upper bound constraint on the proportion of wealth that can be invested in stocks. We compute the equilibrium parameters and study their behavior in good (when dividend growth shocks are positive) and bad (when dividend growth shocks are negative) states of the economy. We demonstrate that the impact of portfolio constraints crucially depends on the intertemporal elasticity of substitution (IES). When IES is greater than unity our model generates countercyclical market prices of risk and stock return volatilities and procyclical stock price-dividend ratios (i.e., market prices of risk and stock return volatilities are higher in bad states and lower in good states, and vice versa for price-dividend ratios) as well as excess stock return volatility, consistently with the literature (e.g., Ferson and Harvey, 1991; Schwert, 1989; Campbell and Cochrane, 1999). Moreover, we demonstrate that our methodology can easily be extended to incorporate heterogeneous beliefs and multiple assets.

Chapter 2 is entitled “Dynamic Mean-Variance Asset Allocation” and studies optimal dynamic mean-variance asset allocation in a general incomplete-market economy with stochastic investment opportunity sets. Mean-variance portfolio choice has always been recognized as the cornerstone of modern portfolio theory, and is still widely employed both in industry and academia due to its tractability and intuitive appeal (e.g., among others, Ait-Sahalia and Brandt, 2001; Campbell and Viceira, 2002; Jagannathan and Ma, 2003; Bansal, Dahlquist and Harvey, 2004; Acharya and Pedersen, 2005; Hong, Scheinkman and Xiong, 2006; Brandt, 2009; Campbell, Serfaty-de Medeiros and Viceira, 2009). The extant literature primarily solves for optimal mean-variance policies assuming that investors are myopic and maximize their criterion over the next period, ignoring the hedging demands arising due to the fluctuations of investment opportunity sets. In contrast with the previous literature, we study dynamic optimal mean-variance asset allocation via dynamic programming by explicitly recognizing the time-inconsistency of mean-variance criteria. While previous literature has reported difficulties in applying dynamic programming to mean-variance criteria due to the failure of iterated expectations property (see e.g., Zhou and Li, 2000), we provide a

recursive representation of the investor's mean-variance criterion which makes possible the application of dynamic programming.

To our best knowledge, our work is the first to obtain closed form expressions for optimal investment policies in a general incomplete-market setting in terms of exogenous model parameters and a new hedge-neutral probability measure that facilitates the tractability of the analysis. We also demonstrate that our dynamically optimal time-consistent policies are generically different from the pre-commitment policies obtained in the literature, from which an investor may deviate in the future unless being able to pre-commit. Moreover, methodology presented in Chapter 2 can efficiently be applied to study the asset allocation for various investment opportunity sets, including those endogenously derived in equilibrium in Chapter 1. Specifically, in this work we obtain explicit expressions for various stochastic investment opportunity sets widely used in the literature, such as models with Gaussian mean-reverting stock returns (e.g., Kim and Omberg, 1996) and models with stochastic volatility (e.g., Liu, 2005; Chacko and Viceira, 2005). Moreover, the calibration to the real data reveals the significance of the hedging demands, which shows that popular myopic policies ignore a substantial fraction of the optimal total demand for stocks in a dynamic setting.

Finally, Chapter 3 entitled “Dynamic Hedging in Incomplete Markets: A Simple Solution” employs the methodology developed in Chapter 2 to study mean-variance hedging problem in a dynamic incomplete-market economy with stochastic investment opportunity sets. The investor in this setting hedges the fluctuations of a non-tradable asset by minimizing the variance of the hedging error given by the difference of the non-tradable asset payoff and tradable wealth at a terminal date. Hedging in incomplete markets is a classical problem covered in major text books as well as academic works but primary in static or myopic settings. When the market is incomplete and perfect hedging by means of standard no-arbitrage methods is no longer feasible the literature derives optimal hedges that maximize investors' utility functions or minimize various hedging criteria such as squared hedging error or hedging error variance. Despite much work in this direction, we still lack tractable

dynamically optimal policies.

By employing the methodology of Chapter 2 we provide tractable dynamically optimal hedges in a general incomplete-market economy by employing the minimum-variance criterion. We demonstrate that our optimal hedges retain the intuitive structure of classical complete-market hedges and are in terms of generalized “Greeks”, widely used in risk management applications. We obtain our dynamic time-consistent hedges via dynamic programming by explicitly recognizing the time-inconsistency of mean-variance criteria while previous works (e.g., Duffie and Richardson, 1991; Schweizer, 1994; Musiela and Rutkowski, 1998) obtain policies that minimize the variance criterion globally and from which the investor may choose to optimally deviate in the future. Moreover, we demonstrate that our hedges typically outperform static and myopic counterparts in plausible economic settings.

1. Asset Pricing in General Equilibrium with Constraints

1.1. Introduction

Portfolio constraints and market frictions have long been considered among key contributors towards understanding investor behavior and equilibrium asset prices. In particular, dynamic equilibrium models with heterogeneous investors facing portfolio constraints have extensively been employed by financial economists to confront a wide range of phenomena such as the equity premium puzzle, mispricing of redundant assets, role of arbitrageurs, impact of heterogeneous beliefs on asset prices, and stock comovements (e.g., among others, Detemple and Murthy, 1997; Basak and Cuoco, 1998; Basak and Croitoru, 2000, 2006; Kogan, Makarov and Uppal, 2007; Gallmeyer and Hollifield 2008; Pavlova and Rigobon, 2008). However, tractable characterizations of equilibria are only obtained assuming that a constrained investor has logarithmic preferences which simplifies the analysis at the cost of assuming investor's myopia.¹ Despite recent developments in portfolio optimization, such as duality method of Cvitanic and Karatzas (1992), portfolio constraints are notoriously difficult to incorporate into general equilibrium analysis as well as portfolio choice when constrained investors have more general preferences inducing hedging demands.

The assumption of logarithmic preferences is not innocuous and impedes the evaluation of the impact of constraints on stock prices and stock return volatilities. Thus, in economic settings with two logarithmic investors and single consumption good (e.g., Detemple and Murthy, 1997; Basak and Cuoco, 1998; Basak and Croitoru, 2000, 2006) stock prices and hence stock return volatilities are unaffected by constraints since the income and substitution effects perfectly offset each other. When the constrained investor is logarithmic, the volatility effects of constraints have been studied in specific settings where the other (un-

¹The assumption that one investor has logarithmic preferences is also commonly made for tractability in models with unconstrained investors who differ in risk aversions. Thus, Dumas (1989) studies dynamic equilibrium in a production economy, where one investor has logarithmic while the other general CRRA preferences. Wang (1996) studies an exchange economy where one investor has logarithmic while the other square-root preferences. One notable exception is Bhamra and Uppal (2009), who study the effect of introducing non-redundant securities on the volatilities of asset returns in an exchange economy with CRRA investors not restricted to being logarithmic.

constrained) investor has different preferences (e.g., Gallmeyer and Hollifield, 2008), which requires further justification. To our best knowledge, this paper is the first to study the effect of different constraints on stock return volatility in a continuous-time economy without relying on the assumption of logarithmic investors. As a result, our solution method yields new insights on the impact of portfolio constraints on stock prices and, in particular, highlights the role of constraints in explaining empirically observed procyclical variation of price-dividend ratios and countercyclical variation of stock return volatilities (i.e., positive shocks to dividend growth rates lead to higher price-dividend ratios and lower stock return volatilities).

We solve for the equilibrium in a continuous-time pure exchange economy with one consumption good and two heterogeneous investors facing portfolio constraints. First, for general preferences and constraints we provide a characterization of interest rates and market prices of risk which highlight the role of constraints and risk sharing, and in specific economic settings can explicitly be characterized in terms of empirically observable quantities such as stock returns and consumption volatilities. Based on these results, we specialize to settings with two CRRA investors one of whom is unconstrained while the other faces portfolio constraints. Specifically, we first derive the equilibrium when the constrained investor faces an upper bound on the proportion of wealth invested in stocks.² Then, we study the impact of short-sale constraints on equilibrium when investors have different beliefs about mean dividend growth. The methodological contribution of the paper is a solution method for the efficient computation of equilibria in economies with constraints. Specifically, we derive stock price-dividend ratios, stock return volatilities and other parameters in terms of wealth-

²Srinivas, Whitehouse and Yermo (2000) in a survey of pension fund regulations show that limits on both domestic and foreign equity holdings of pension funds are in place in a number of OECD countries such as Germany (30% on EU and 6% on non-EU equities), Switzerland (30% on domestic and 25% on foreign equities) and Japan (30% on domestic and 30% on foreign equities), among others. Moreover, our approach allows to study the impact of passive investors that hold a fixed fraction of their wealth in stocks, as in Chien, Cole and Lustig (2008). Samuelson and Zeckhouser (1988) document the popularity of this strategy using as an example the participants of popular TIAA/CREF retirement plan who choose a fraction of wealth to be invested in stocks and rarely change it due to “status quo bias”, while Campbell (2006) points out that households may limit their participation in stock market and invest cautiously due to the lack of necessary skills. Important special case of our framework is stock market non-participation which in year 2002 accounted for 50% of U.S. households (e.g., Guvenen, 2006).

consumption ratios that can be computed numerically via a simple iterative procedure with fast convergence.

At the first step of our analysis when we allow for general preferences, we demonstrate that the riskless rates and market prices of risk include new terms that capture the effects of constraints and risk sharing. In specific settings we obtain the expressions for interest rates and market prices of risk in terms of intuitive and empirically observable parameters such as stock return and consumptions volatilities. The tractability of our results allows to compare interest rates in constrained and unconstrained economies for a given allocation of consumption among investors and demonstrate that for various constraints interest rates will be lower in constrained economies whenever both investors have the same prudence-risk aversion ratios.

Using the insights from the case with general preferences we show that when investors have (identical for simplicity) CRRA preferences, one of them faces an upper bound on the proportion of wealth invested in stocks, and dividends follow a geometric Brownian motion, the interest rates and market prices of risk can explicitly be expressed in terms of marginal utility ratios, their volatilities and the volatilities of stock returns. We completely characterize the equilibrium by computing these volatilities numerically. While in models with two logarithmic investors price-dividend ratios and stock return volatilities are deterministic functions of time, in our setting these parameters depend on constrained investor's consumption share which evolves stochastically. The effect of constraints on price-dividend ratios and stock-return volatilities depends on the relative strength of classical income and substitution effects. When the intertemporal elasticity of substitution (IES) is less than one and hence the income effect dominates, price-dividend ratios increase while stock return volatilities decrease with tighter constraints, and vice versa when IES is greater than one and the substitution effect is stronger.³ Moreover, the effects of constraints are more

³When the investment opportunities worsen, the income effect induces investors to decrease consumption and save more while the substitution effect induces them to consume more and save less due to cheaper current consumption. For CRRA preferences with risk aversion γ , $IES=1/\gamma$, the income effect dominates for $IES < 1$ and the substitution effect dominates for $IES > 1$ while for $IES = 1$ both effects perfectly offset each other.

pronounced in bad times, when dividends are hit by adverse shocks, than in good times.

To understand the intuition, we first evaluate the impact of portfolio constraints on investment opportunity sets and demonstrate that interest rates decrease while market prices of risk increase with tighter constraints, and that the effects of constraints are stronger in bad times. When the portfolio constraint binds, negative shocks to dividends shift the distribution of the aggregate wealth and consumption to the constrained investor since she is less exposed to stock market fluctuations. Thus, in bad times, when the constrained investor holds a significant fraction of aggregate wealth and consumption, the price-dividend ratio is approximately equal to her wealth-consumption ratio. With tighter constraints the investment opportunities of the constrained investor deteriorate since interest rates fall and she is unable to benefit from the increase in market prices of risk. As a result, her wealth-consumption ratio, and hence the price-dividend ratio, increases when the income effect dominates and decreases when the substitution effect dominates. The effect of constraints is weaker in good times since as the share of the unconstrained investor in aggregate wealth and consumption increases, all the economic parameters, including price-dividend ratios, converge to the parameters in the unconstrained economy.

Thus, when the substitution effect dominates, price-dividend ratios turn out to be procyclical (lower in bad times than in good times) while stock return volatilities exceed the volatility of dividends and are countercyclical (higher in bad times than in good times), consistently with the empirical evidence (e.g., Schwert, 1989; Campbell and Cochrane, 1999). Moreover, irrespective of investors' intertemporal elasticities of substitution, market prices of risk turn out to be countercyclical (e.g., Ferson and Harvey, 1991) since in bad times unconstrained investors lose wealth and require higher compensation for risk taking, causing market prices of risk to go up. We also study the survival of constrained investors in equilibrium and demonstrate that their impact on financial markets is gradually eliminated in the course of time but is significant even after one hundred years.

Finally, we extend our baseline analysis to economic settings with heterogeneous beliefs and multiple stocks. In both cases, for general preferences we derive expressions for interest

rates and market prices of risk similar to those in the baseline model. In the case of heterogeneous beliefs we solve for equilibrium in a model where two investors have the same CRRA utilities and disagree on the dividend growth rate. The optimist is unconstrained while the pessimist faces a constraint on the proportion of wealth that can be held in short positions in stocks. We demonstrate that tighter short-sale constraints imply higher price-dividend ratios since they increase the constrained investor's demand for stocks. We also find that stock return volatility in the constrained economy can be both higher or lower than the volatility in an unconstrained economy, depending on whether the latter is higher or lower than the volatility of dividend growth. This is because the short-sale constraints do not allow the investor to trade on her pessimism making her stockholding closer to what it would be in the case of homogeneous beliefs, and hence, the stock return volatility shifts towards volatility in an unconstrained homogeneous economy, given by the volatility of dividends.

Our solution method is based on a combination of the duality approach and dynamic programming. First, following Cvitanic and Karatzas (1992) we derive optimal consumptions in terms of the state price densities in equivalent unconstrained fictitious economies in which the interest rates and market prices of risk are adjusted to account for the difference in investors' behavior in constrained economies. Then, market clearing for consumption yields expressions for equilibrium parameters in terms of the adjustment parameters that solve a certain fixed point problem. Moreover, in our specific examples these adjustments to interest rates and market prices of risk can be derived in terms of instantaneous volatilities of stock returns and the ratios of marginal utilities of the two investors. Next, these volatilities and all the equilibrium parameters are explicitly characterized in terms of investors' wealth-consumption ratios that satisfy a system of quasilinear Hamilton-Jacobi-Bellman equations. We solve this system of equations numerically via a simple iterative procedure that requires solving a simple system of linear equations at each step.

There is a growing literature studying dynamic equilibria in continuous-time economies with heterogeneous investors and portfolio constraints assuming that constrained investors

have logarithmic preferences. Basak and Cuoco (1998) consider a model in which one investor is unconstrained and guided by a general time-additive utility function while the other investor cannot invest in the stock market and has logarithmic preferences. They derive the riskless rates and market prices of risk in this economy and characterize all the equilibrium parameters explicitly when both investors are logarithmic. Detemple and Murthy (1997), Basak and Croitoru (2000, 2006) present equilibrium models with two logarithmic investors, heterogeneous beliefs and portfolio constraints. Hugonnier (2008) considers a similar model and shows that under restricted participation the stock prices implied by market clearing may contain a bubble and in the setting with multiple stocks the equilibrium might not be unique. In contrast to our work all the above papers do not find the impact of constraints on stock prices and their moments.

Jarrow (1980) studies the equilibrium effect of short-sale constraints in a one-period economy with mean-variance investors that have heterogeneous beliefs. Dumas and Maenhout (2002) develop an approach with two central planners for solving incomplete-market equilibrium with two CRRA investors. However, in their analysis the variance-covariance matrix of returns is taken as given and hence they do not study the impact of constraints on volatility. Kogan, Makarov and Uppal (2007) derive equilibrium parameters in an economy with borrowing constraints when one investor is logarithmic while the other has general CRRA utility and find that all the moments of asset returns are deterministic and stock return volatilities are unaffected by constraints. When little borrowing is permitted they numerically find interest rates and market prices of risk as functions of wealth distributions but do not consider the volatilities of stock returns. He and Krishnamurthy (2008) consider a model of intermediated asset pricing in which individual households are logarithmic and invest into stock only via an intermediary guided by CRRA utility. Wu (2008) studies the equilibrium in a setting with one unconstrained and one buy-and-hold CRRA investors. Gallmeyer and Hollifield (2008) study the asset pricing with short-sale constraints in the presence of heterogeneous beliefs when the pessimist and optimist have logarithmic and CRRA utilities respectively. They study equilibrium parameters by employing Monte-Carlo

simulations and derive conditions for stock return volatilities to be larger or lower than in the unconstrained case assuming that investors have the same share of aggregate wealth at the initial date.

Bhamra (2007) analyzes the effect of liberalization on emerging markets' cost of capital in a model with two logarithmic investors, two stocks and one consumption good. Pavlova and Rigobon (2008) and Schornick (2009) consider models with constrained logarithmic investors and two consumption goods in international finance framework and derive various asset-pricing implications assuming that investors face preference shocks. Longstaff (2009) develops a two-asset economy where one of the assets is non-tradable for a certain period and logarithmic investors are heterogeneous in time discount parameter.

There are a number of papers that solve models with heterogeneous investors and portfolio constraints numerically in discrete time. Cuoco and He (2001) consider a model with general utilities and derive equilibrium asset prices in terms of stochastic weights of a representative investor's utility which are obtained numerically from a nonlinear system of equations. Guvenen (2006) solves numerically a model with restricted market participation when investors are guided by recursive utilities. Chien, Cole and Lustig (2008) also in a discrete-time framework consider a model with non-participants, passive and active investors guided by CRRA preferences, where passive investors hold fixed portfolios while active ones adjust them each period. Gomes and Michaelides (2008) study numerically the equilibrium with incomplete markets and investors subject to fixed cost of stock market participation and by calibration generate high equity premium and match observed market participation rate. Dumas and Lyasoff (2008) solve for equilibrium in various incomplete market settings in discrete time by employing binomial trees. These works do not study the impact of constraints on conditional stock return volatilities and do not provide expressions for equilibrium parameters in terms of observable quantities as we do in this paper by employing considerable flexibility of continuous-time methods.

The remainder of the paper is organized as follows. In Section 1.2, we derive interest rates and market prices of risk for general utility functions under the assumption that the

dual optimization problem has a solution and discuss their properties. In Section 1.3 we illustrate our solution method by computing the equilibrium in a model with two CRRA investors where one investor is unconstrained while the other faces an upper bound on the fraction of wealth invested in stocks. Section 1.4 extends our baseline analysis to the settings with heterogeneous beliefs and multiple stocks. We also solve for equilibrium in a model with heterogeneous beliefs in which one of the investors faces short-sale constraints. Section 1.5 concludes, Appendix A provides the proofs and Appendix B provides further details for our numerical method.

1.2. General Equilibrium with Constraints

1.2.1. Economic Setup

We consider a continuous-time economy with one consumption good and an infinite horizon. The uncertainty is represented by a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$, on which is defined a Brownian motion w . All the stochastic processes that appear in the paper are adapted to $\{\mathcal{F}_t, t \in [0, \infty)\}$, the augmented filtration generated by w .

The investors trade continuously in two securities, a riskless bond in zero net supply with instantaneous interest rate r and a stock in a positive net supply, normalized to one unit. The stock is a claim to an exogenous strictly positive stream of dividends δ following the dynamics

$$d\delta_t = \delta_t[\mu_{\delta_t}dt + \sigma_{\delta_t}dw_t], \quad (1)$$

where the dividend mean-return, μ_{δ} , and volatility, σ_{δ} , are stochastic processes. The dividend process (1) and its moments are assumed to be well-defined, without explicitly stating the regularity conditions. We consider equilibria in which bond prices, B , and stock prices, S , follow processes

$$dB_t = B_t r_t dt, \quad (2)$$

$$dS_t + \delta_t dt = S_t[\mu_t dt + \sigma_t dw_t], \quad (3)$$

where the interest rate r , the stock mean return μ and volatility σ are stochastic processes determined in equilibrium, and bond price at time 0 is normalized so that $B_0 = 1$.

There are two investors in the economy. Investor 1 is endowed with s units of stock and $-b$ units of bond, while investor 2 is endowed with $1 - s$ units of stock and b units of bond. The investors choose consumption, c_i , and an investment policy, $\{\alpha_i, \theta_i\}$, where α_i and θ_i denote the fractions of wealth invested in bonds and stocks, respectively, and hence, $\alpha_i + \theta_i = 1$. Investor i 's wealth process W evolves as

$$dW_{it} = \left[W_{it} \left(r_t + \theta_{it} (\mu_t - r_t) \right) - c_{it} \right] dt + W_{it} \theta_{it} \sigma_t dw_t, \quad (4)$$

and her investment policies are subject to portfolio constraints

$$\theta_i \in \Theta_i, \quad i = 1, 2, \quad (5)$$

where $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i]$. We also assume that initial endowments of stocks are such that θ_i at time 0 belong to sets Θ_i . Thus, the financial market in our economy is incomplete due to the presence of portfolio constraints (5).

Each investor i ($i = 1, 2$) is guided by an expected utility over a stream of consumption c . In particular, her dynamic optimization is given by

$$\max_{c_i, \theta_i} E \left[\int_0^\infty e^{-\rho t} u_i(c_{it}) dt \right], \quad (6)$$

subject to the budget constraint (4), no-bankruptcy constraint $W_t \geq 0$ and portfolio constraints (5), for some discount parameter $\rho > 0$. The utility functions $u_i(c)$ are assumed to be increasing, concave, three times continuously differentiable, satisfying Inada's conditions

$$\lim_{c \downarrow 0} u'_i(c) = \infty, \quad \lim_{c \uparrow \infty} u'_i(c) = 0, \quad i = 1, 2. \quad (7)$$

By A_{it} and P_{it} we denote absolute risk aversion and prudence parameters of investor i , given by

$$A_{it} = -\frac{u''_i(c)}{u'_i(c)}, \quad P_{it} = -\frac{u'''_i(c)}{u''_i(c)}, \quad (8)$$

and assume that both are strictly positive for each investor.

Next, we define an *equilibrium* in this economy as a set of parameters $\{r_t, \mu_t, \sigma_t\}$ and of consumption and investment policies $\{c_{it}^*, \alpha_{it}^*, \theta_{it}^*\}_{i=1}^2$ such that consumption and investment policies solve dynamic optimization problem (6) for each investor, given price parameters $\{r_t, \mu_t, \sigma_t\}$, and consumption and financial markets clear, i.e.,

$$\begin{aligned} c_{1t}^* + c_{2t}^* &= \delta_t, \\ \alpha_{1t}^* W_{1t}^* + \alpha_{2t}^* W_{2t}^* &= 0, \\ \theta_{1t}^* W_{1t}^* + \theta_{2t}^* W_{2t}^* &= S_t, \end{aligned} \tag{9}$$

where W_{1t}^* and W_{2t}^* denote optimal wealths of investors 1 and 2 under optimal consumption and investment policies.

1.2.2. Characterization of Equilibrium

This Section characterizes the parameters of equilibria and studies their properties in economies with constrained investors. In particular, by employing the duality method of Karatzas and Cvitanic (1992), we recover expressions for interest rates and market prices of risk in equilibrium in terms of the parameters of equivalent fictitious unconstrained economies. These expressions are intuitive and highlight the impact of risk-sharing and attitude towards risk on equilibrium parameters. Moreover, they form a basis for an efficient methodology for computing equilibria, which we develop in Section 1.3.

We start by noting that since the market is incomplete due to the presence of portfolio constraints, a Pareto optimal allocation may not be feasible and hence, the ratio of the marginal utilities of consumption of the investors follows a stochastic process. This ratio can be interpreted as a stochastic weight in the construction of a representative-investor preferences in an equivalent economy, and serves as a state variable in terms of which the equilibrium can be characterized (e.g., Basak and Cuoco, 1998; Cuoco and He, 2001). By employing the methodology of Cvitanic and Karatzas (1992) we obtain optimal consump-

tions and then derive the equilibrium parameters from the market clearing conditions.⁴ This approach is similar to the approach in Basak (2000), who characterizes the equilibrium in an economy where investors have heterogeneous beliefs, but in contrast to our work are unconstrained.

We start by characterizing optimal consumptions of constrained investors in a partial equilibrium in which the investment opportunities are taken as given, and then obtain the interest rate r , and the market price of risk κ , from the consumption clearing condition. For each investor i , following the approach of Cvitanic and Karatzas (1992), we characterize the optimality conditions for consumption by embedding investor i 's partial equilibrium economy into an equivalent fictitious complete-market economy with bond and stock prices following dynamics with adjusted parameters:

$$dB_t = B_t[r_t + f(\nu_{it}^*)]dt, \quad (10)$$

$$dS_t + \delta_t dt = S_t[(\mu_t + \nu_{it}^* + f(\nu_{it}^*))dt + \sigma_t dw_t], \quad (11)$$

where $f_i(\nu)$ are *support functions* for the sets of portfolio constraints Θ_i , defined as

$$f_i(\nu) = \sup_{\theta \in \Theta_i} (-\nu\theta), \quad (12)$$

and ν_{1t}^* and ν_{2t}^* solve so called dual optimization problem, defined in Cvitanic and Karatzas (1992), and lie in the *effective domains* for support functions, given by

$$\Upsilon_i = \{\nu \in \mathbb{R} : f_i(\nu) < \infty\}. \quad (13)$$

It follows from the dynamics of bond and stock prices in fictitious economy (10)–(11) that the corresponding state prices ξ_{it} evolve as

$$d\xi_{it} = -\xi_{it}[r_{it}dt + \kappa_{it}dw_t], \quad (14)$$

where r_{it} and κ_{it} denote the adjusted riskless rate and market price of risk in fictitious economy i , given by

$$r_{it} = r_t + f_i(\nu_{it}^*), \quad \kappa_{it} = \kappa_t + \frac{\nu_{it}^*}{\sigma_t}, \quad (15)$$

⁴Cuoco (1997) studies consumption-portfolio choice of constrained investors, mainly at a partial equilibrium level, and extends the results of Cvitanic and Karatzas to the case of more general utility functions and forms of market incompleteness. He derives a CAPM in an economy with portfolio constraints but does not study interest rates and other parameters of equilibrium.

where $\kappa = (\mu - r)/\sigma$ is the *market price of risk* in the original constrained economy.

Throughout this Section we assume that the solutions to dual optimization problems exist and since the fictitious economies are complete, the marginal utilities of optimal consumption are given by

$$e^{-\rho t} u'_i(c_{it}^*) = \psi_i \xi_{it}, \quad i = 1, 2, \quad (16)$$

for some constants $\psi_i > 0$. The first order conditions (16) and state prices (14) demonstrate that consumption and investment decisions of the constrained investor are equivalent to those of an unconstrained one, who faces interest rates and market prices of risk adjusted to account for the constraints. Moreover, optimality conditions in (16) allow to express consumptions c_{it}^* in terms of state prices in fictitious economies as follows:

$$c_{it}^* = I_i(\psi_i e^{\rho t} \xi_{it}), \quad i = 1, 2, \quad (17)$$

where $I_i(\cdot)$ denote inverse functions for marginal utilities $u'_i(\cdot)$.

The expressions for marginal utilities in (16) also imply that the *ratio of investors' marginal utilities*, defined as

$$\lambda_t = \frac{u'_1(c_{1t}^*)}{u'_2(c_{2t}^*)}, \quad (18)$$

is stochastic in equilibrium, and not a constant as in complete markets (e.g., Karatzas and Shreve, 1998) where consumption allocations are Pareto efficient. Basak and Cuoco (1998) and Cuoco and He (2001) demonstrate that the process λ serves as a convenient state variable in terms of which the equilibrium parameters can be expressed. Moreover, in an equivalent complete-market economy with a representative investor, parameter λ can be interpreted as a *stochastic weight* in the utility $u(c; \lambda)$ of a *representative investor*, given by

$$u(c; \lambda) = \max_{c_1 + c_2 = c} u_1(c_1) + \lambda u_2(c_2), \quad (19)$$

and follows a stochastic process

$$d\lambda_t = -\lambda_t[\mu_{\lambda_t} dt + \sigma_{\lambda_t} dw_t]. \quad (20)$$

The parameters μ_{λ} and σ_{λ} are determined in equilibrium and quantify the violation of Pareto-optimality in the economy.

Next we characterize the parameters of our economy in equilibrium in terms of adjustments ν_{it}^* from the market clearing in consumption. To determine the interest rate r and market price of risk κ we substitute optimal consumptions (17) into consumption clearing condition in (9), apply Itô's Lemma to both sides and recover equilibrium parameters by matching the drift and volatility terms. Similarly, from optimality conditions (16), by applying Itô's Lemma to equation (18) for λ_t and comparing the result with the process for λ_t in (20) we recover parameters μ_λ and σ_λ . The following Proposition summarizes our results.

Proposition 1.1. *If there exists an equilibrium, the riskless interest rate r , market price of risk κ , drift μ_λ and volatility σ_λ of weighting process λ that follows (20) are given by*

$$r_t = \bar{r}_t - \frac{A_t}{A_{1t}} f_1(\nu_{1t}^*) - \frac{A_t}{A_{2t}} f_2(\nu_{2t}^*) - \frac{A_t^3 (P_{1t} + P_{2t})}{2A_{1t}^2 A_{2t}^2} \sigma_{\lambda t}^2 - \frac{A_t^3}{A_{1t} A_{2t}} \left(\frac{P_{1t}}{A_{1t}} - \frac{P_{2t}}{A_{2t}} \right) \delta_t \sigma_{\delta t} \sigma_{\lambda t} \quad (21)$$

$$\kappa_t = \bar{\kappa}_t - \frac{A_t}{A_{1t}} \frac{\nu_{1t}^*}{\sigma_t} - \frac{A_t}{A_{2t}} \frac{\nu_{2t}^*}{\sigma_t}, \quad (22)$$

$$\mu_{\lambda t} = A_t \delta_t \sigma_{\delta t} \sigma_{\lambda t} + f_1(\nu_{1t}^*) - f_2(\nu_{2t}^*) - \frac{A_t}{A_{1t}} \sigma_{\lambda t}^2, \quad \sigma_{\lambda t} = \frac{\nu_{1t}^* - \nu_{2t}^*}{\sigma_t}, \quad (23)$$

where \bar{r} is the riskless rate and $\bar{\kappa}$ is the market price of risk in an unconstrained economy, given by

$$\bar{r}_t = \rho + A_t \delta_t \mu_{\delta t} - \frac{A_t P_t}{2} \delta_t^2 \sigma_{\delta t}^2, \quad \bar{\kappa}_t = A_t \delta_t \sigma_{\delta t} \quad (24)$$

A_{it} , P_{it} , and A_t and P_t are absolute risk aversions and prudence parameters of investor i and a representative investor with utility (19), respectively.⁵

Optimal consumptions c_i^* , wealths W_i , stock S and optimal investment policies θ_i^* are given by

$$c_{it}^* = g_i(\delta_t, \lambda_t), \quad (25)$$

$$W_{it}^* = \frac{1}{\xi_{it}} E_t \left[\int_0^\infty \xi_{is} c_{is}^* ds \right], \quad (26)$$

⁵As demonstrated in Basak (2000), the risk aversion, A , and prudence, P , of the representative investor can be obtained from the following expressions:

$$\frac{1}{A_t} = \frac{1}{A_{1t}} + \frac{1}{A_{2t}}, \quad \frac{P_t}{A_t^2} = \frac{P_{1t}}{A_{1t}^2} + \frac{P_{2t}}{A_{2t}^2}.$$

$$S_t = W_{1t}^* + W_{2t}^*, \quad (27)$$

$$\theta_{it}^* = \frac{1}{\sigma_t} \left(W_{it}^* \left(\kappa_t + \frac{\nu_{it}^*}{\sigma_t} \right) + \frac{\phi_{it}}{\xi_{it}} \right), \quad (28)$$

where functions $g_i(\delta_t, \lambda_t)$ are such that c_{1t}^* and c_{2t}^* satisfy consumption clearing in (9) and equation (18) for process λ , state prices ξ_{it} follow processes (14) and ϕ_i are such that

$$M_{it} \equiv E_t \left[\int_0^\infty \xi_{is} c_{is}^* ds \right] = M_{i0} + \int_0^t \phi_{is} dw_s.$$

Initial value λ_0 is such that the budget constraints at time 0 are satisfied:

$$s_i S_0 + b_i = W_{i0}^*, \quad (29)$$

where $s_1 = s$, $s_2 = 1 - s$, $b_1 = -b$ and $b_2 = b$. Moreover, adjustments ν_{it}^* satisfy complementary slackness condition

$$f_i(\nu_{it}^*) + \theta_{it}^* \nu_{it}^* = 0. \quad (30)$$

Proposition 1.1 provides the characterization of equilibrium parameters in terms of adjustments ν_i^* in fictitious economy. Expression (21) decomposes interest rates r into groups of terms that separate the effects of constraints and the inefficiency of risk sharing. The first term in (21) is the riskless rate in the unconstrained economy with the representative investor. The next two terms capture the effect of binding constraints on interest rates and tend to increase or decrease them depending on the signs of support functions $f_i(\nu)$. In particular, these terms are positive in economic settings with binding portfolio constraints when investors buy more bonds. This is due to the fact that the investors behave as if their subjective interest rates r_{it} in their fictitious economy were higher than in the real one, and hence positive adjustments $f_i(\nu_i^*)$. Finally, the last two terms in expression (21) capture the effect of risk sharing, quantified by volatility σ_λ . The weight λ acts as a state variable that gives rise to specific hedging demands that can push interest rates in either direction.

Similarly, the expression (22) for the market price of risk is comprised of the market price of risk in an unconstrained economy (first term in (22)) and the effects of constraints

(second and third terms in (22)). Expressions for the drift μ_λ and volatility σ_λ parameters of the stochastic weighting process λ in (23) demonstrate that this process, in general, is no longer a local martingale as in works assuming logarithmic constrained investor (e.g., Basak and Cuoco, 1998; Gallmeyer and Hollifield, 2008; Pavlova and Rigobon, 2008). Finally we observe that optimal consumptions, wealths, stock prices and investments can be obtained from expressions (25) – (28) when the parameters of equilibrium, and hence all state prices, are known.

The results in Proposition 1.1 can also be used to compute the equilibrium parameters numerically. On one hand, Proposition 1.1 expresses equilibrium parameters and investment policies in terms of adjustments ν_i^* , and on the other, the adjustments can be obtained from the complementary slackness condition (30). Thus, finding the adjustments becomes essentially a fixed point problem, which can potentially be solved by the method of successive iterations. Moreover, as demonstrated in Huang and Pages (1992), under certain conditions optimal wealths (26) satisfy linear PDEs with coefficients determined by equilibrium parameters while optimal policies (28) can be expressed in terms of derivatives of wealths W_i^* . Hence, the adjustments can be expressed in terms of derivatives of W_{it} from conditions (30) and substituted back into the PDEs for optimal wealths. Thus, the characterization of equilibrium reduces to solving a system of quasilinear PDEs which, as we demonstrate in Section 1.3, can efficiently be solved numerically for specific constraints.

1.2.3. Further Properties of Equilibrium

We here explore the implications of Proposition 1.1 by noting that in various economic settings the signs of adjustments ν_i^* and support functions $f_i(\nu)$ can easily be determined explicitly from the definitions of support functions and effective domains in (12) and (13). Moreover, the interest rates r and market prices of risk κ can be expressed in terms of empirically observed quantities, such as stock and consumption volatilities, thus providing empirical implications of the model.

Table 1.1
Effective Domains and Support Functions

Case	Constraint	Υ	$f(\nu)$
(a)	$\theta \in \mathbb{R}$	0	0
(b)	$\theta = 0$	\mathbb{R}	0
(c)	$\underline{\theta} \leq \theta \leq \bar{\theta}, \underline{\theta} \leq 0$	\mathbb{R}	+
(d)	$\theta \leq \bar{\theta}, \bar{\theta} > 0$	$\nu \leq 0$	+
(e)	$\theta \geq \underline{\theta}, \underline{\theta} < 0$	$\nu \geq 0$	+
(f)	$\theta \geq \underline{\theta}, \underline{\theta} > 0$	$\nu \geq 0$	-

Table 1.1 presents the effective domains and the signs of the support functions for plausible constraints and allows to analyze their effect on equilibrium parameters. For example, when investors face constraints on the proportion of wealth invested in stocks (case (d) in Table 1.1) the results in Proposition 1.1 and Table 1.1 imply that these constraints tend to decrease the interest rates and increase the market prices of risk relative to an unconstrained model if stock volatility σ is strictly positive. Hence, these constraints work in the right direction for explaining the equity premium puzzle (e.g., Mehra and Prescott, 1985). The overall effect of constraints on interest rates is convoluted by the risk sharing captured by the last two terms in the expression for interest rates (21). The following Corollary to Proposition 1.1 establishes simple sufficient conditions under which the interest rate r will be lower than the interest rate \bar{r} in a representative-investor unconstrained economy.

Corollary 1.1. *If the utility functions and the allocation of consumption are such that $P_1/A_1 = P_2/A_2$ and the sets of portfolio constraints have positive support functions $f_i(\nu)$ then the interest rate in a constrained economy, r , is lower than in an unconstrained one, \bar{r} , and the following upper bound for rate r holds:*

$$r_t \leq \bar{r}_t - \frac{A_t^3(P_{1t} + P_{2t})}{2A_{1t}^2 A_{2t}^2} \sigma_\lambda^2. \quad (31)$$

The Corollary demonstrates that the inability to share risks contributes to the decrease of

interest rates by creating hedging needs against fluctuating ratios of marginal utilities λ . The condition that investors have the same prudence-risk aversion ratio is in particular satisfied when both investors have identical HARA preferences.⁶ In the case of two logarithmic investors when one of them is unconstrained the result in Corollary 1.1 has also been pointed out in the literature (e.g., Basak and Cuoco, 1998).

Conveniently, in various economic settings interest rates and market price of risk can be expressed only in terms of the parameters of utility functions and empirically observed parameters. For example, when investor 1 is unconstrained and investor 2 faces a constraint allowing her to invest in stock no more than a certain fraction of wealth (case (d) of Table 1.1), it can be observed that parameters r and κ are given by:

$$r_t = \bar{r}_t - \frac{A_t}{A_{2t}} \bar{\theta} \sigma_t \sigma_{\lambda t} - \frac{A_t^3 (P_{1t} + P_{2t})}{2A_{1t}^2 A_{2t}^2} \sigma_{\lambda t}^2 - \frac{A_t^3}{A_{1t} A_{2t}} \left(\frac{P_{1t}}{A_{1t}} - \frac{P_{2t}}{A_{2t}} \right) \delta_t \sigma_{\delta t} \sigma_{\lambda t}, \quad \kappa_t = \bar{\kappa}_t + \frac{A_t}{A_{2t}} \sigma_{\lambda t}, \quad (32)$$

where stock return volatility σ can easily be obtained from the data, while the weighting process volatility σ_{λ} can be obtained in terms of utility parameters and the parameters of the consumption processes for each investor. In particular, assuming that consumption processes c_i for each investor follow Itô's processes

$$dc_{it} = c_{it} [\mu_{c_{it}} dt + \sigma_{c_{it}} dw_t], \quad (33)$$

applying Itô's Lemma to the definition of weighting process λ in (18) we find that

$$\sigma_{\lambda t} = A_{1t} c_{1t} \sigma_{c_{1t}} - A_{2t} c_{2t} \sigma_{c_{2t}}. \quad (34)$$

In specific frameworks the volatilities of consumption growth can potentially be estimated from the data. In particular, for the model with restricted participation ($\bar{\theta} = 0$) Malloy, Moskowitz and Vissing-Jorgensen (2009) estimate consumption volatilities of stock market participants and non-participants to be 3.6% and 1.4% respectively, while Mankiw and Zeldes (1991) and Guvenen (2006) show that the share of consumption of non-participants

⁶For HARA utility function absolute risk aversion is given by $-u''(c)/u'(c) = \gamma/(\gamma_0 + c)$. Differentiating both sides of this expression and then dividing by $-u''(c)/u'(c)$ we obtain that $P_i/A_i = 1 + \gamma$, and hence, the prudence-risk aversion ratio is the same for both investors.

in aggregate consumption is 0.68. As a result, the expressions for r and κ in (32) can potentially be used for identifying the parameters of the utility functions of investors as well as for quantifying the impact of risk sharing inefficiencies on the interest rates and market prices of risk.

1.3. Equilibrium with Proportional Constraints

This Section applies the results of Section 1.2 to compute and analyze the equilibrium in a specific economic setting in which investor 1 is unconstrained while investor 2 faces a constraint allowing her to invest in stock no more than a certain fraction of wealth. For simplicity we assume that dividends follow a geometric Brownian motion and both investors have identical CRRA preferences. Using the results of Section 1.2, in Section 1.3.1 we present a simple solution method for finding an equilibrium in this economy, and in Section 1.3.2 we study the impact of constraints on the equilibrium. In our setting with fully rational investors we also study the survival of constrained investors in the long run and demonstrate that it takes a long time to eliminate their impact on financial markets.

1.3.1. Characterization and Computation of Equilibrium

In this Section we present a solution method which allows to compute the equilibrium in an efficient way. This method does not rely on a widely used assumption of a logarithmic constrained investor (e.g., Detemple and Murthy, 1997; Basak and Cuoco, 1998; Basak and Croitoru, 2000, 2006; Kogan, Makarov and Uppal, 2003; Bhamra, 2007; Gallmeyer and Hollifield, 2008; Hugonnier, 2008; Pavlova and Rigobon, 2008; Schornick, 2009) which allows to derive the adjustments ν_i^* in fictitious economy explicitly at the cost of investor's myopia inherent in logarithmic preferences. In discrete time, Cuoco and He (2001), Guvenen (2006), Chien, Cole and Lustig (2008) and Gomes and Michaelides (2008) study the models with constrained heterogeneous investors numerically without assuming that constrained investor is logarithmic. In contrast to these works, in settings with two CRRA investors the flexibility of continuous-time analysis allows us to recover tractable expressions for interest rates and market prices of risk and to find new insights on the impact of constraints on

price-dividend ratios and stock return volatilities.

Finding an equivalent unconstrained economy is a challenging problem which so far has only been solved for logarithmic investors (e.g., Cvitanic and Karatzas, 1992; Karatzas and Shreve, 1998) or CRRA investors but assuming constant investment opportunity sets (e.g., Tepla, 2000a). We tackle this problem by first expressing the parameters of the fictitious economy in terms of the stochastic weighting process λ , and the volatilities of λ and stock returns, which then are obtained in terms of the investors' wealth-consumption ratios satisfying Hamilton-Jacobi-Bellman equations. Even though in equilibrium the coefficients of HJB equations themselves depend on the sensitivities of wealth-consumption ratios with respect to parameter λ , we demonstrate that the time-independent solutions can easily be obtained via an iterative procedure that at each step requires solving a simple system of linear algebraic equations.

Throughout Section 1.3 we assume for simplicity that dividends follow a geometric Brownian motion

$$d\delta_t = \delta_t[\mu_\delta dt + \sigma_\delta dw_t], \quad (35)$$

both investors have CRRA utilities with relative risk aversion parameter γ , given by⁷

$$u_i(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}, \quad i = 1, 2, \quad (36)$$

and solve optimization problem in (6) subject to budget constraint (4), no-bankruptcy constraint $W_t \geq 0$, and portfolio constraint $\theta \leq \bar{\theta}$ for investor 2, while investor 1 is unconstrained. By $J_i(W_t, \lambda_t, t)$ we denote the indirect utility function of investor i .

For convenience, we solve the optimization problem of constrained investor 2 in an equivalent fictitious unconstrained economy in which she maximizes objective function (6) subject to budget constraint

$$dW_{2t} = \left[W_{2t} \left(r_t + f_2(\nu_{2t}^*) + \theta_{2t}(\mu_t - r_t + \nu_{2t}^*) \right) - c_{2t} \right] dt + W_{2t} \theta_{2t} \sigma_t dw_t, \quad (37)$$

⁷The assumption that investors have identical risk aversions is made for simplicity. More general case can be considered along the same lines.

where ν_{2t}^* and $f_2(\nu_{2t}^*)$ are adjustments to stock mean returns and riskless rates respectively. By applying dynamic programming we find that the indirect utility functions should satisfy the following HJB equations:

$$0 = \max_{c_i, \theta_i} \left\{ e^{-\rho t} \frac{c_{it}^{1-\gamma}}{1-\gamma} + \frac{\partial J_{it}}{\partial t} + \left[W_{it} \left(r_t + f_i(\nu_{it}^*) + \theta_{it}(\mu_t - r_t + \nu_{it}^*) \right) - c_{it} \right] \frac{\partial J_{it}}{\partial W_{it}} - \lambda_t \mu_{\lambda t} \frac{\partial J_{it}}{\partial \lambda_t} + \frac{1}{2} \left[W_{it}^2 \theta_{it}^2 \sigma_t^2 \frac{\partial^2 J_{it}}{\partial W_{it}^2} - 2W_{it} \theta_{it} \lambda_t \sigma_t \sigma_{\lambda t} \frac{\partial^2 J_{it}}{\partial W_{it} \partial \lambda_t} + \lambda_t^2 \sigma_{\lambda t}^2 \frac{\partial^2 J_{it}}{\partial \lambda_t^2} \right] \right\}, \quad (38)$$

with transversality condition $E_t[J_{iT}] \rightarrow 0$ as $T \rightarrow \infty$, which guarantees the convergence of the integral in investors' optimization (6). We next obtain expressions for ν_i^* and $f_i(\nu_i^*)$ without solving the dual problem by noting that since investor 1 is unconstrained $\nu_1^* = 0$ (case (a) in Table 1.1) while ν_2^* can be obtained from equilibrium expression for $\sigma_{\lambda t}$ in (23), and hence,

$$\nu_{1t}^* = 0, \quad f_1(\nu_{1t}^*) = 0, \quad \nu_{2t}^* = -\sigma_t \sigma_{\lambda t}, \quad f_2(\nu_{2t}^*) = \bar{\theta} \sigma_t \sigma_{\lambda t}. \quad (39)$$

The HJB equations in (38) are standard except for the fact that the equation for investor 2 is in terms of parameters of fictitious economy, which allows to formulate her problem as an unconstrained one. We conjecture that the indirect utility functions are given by

$$J_i(W_i, \lambda, t) = e^{-\rho t} \frac{W_i^{1-\gamma}}{1-\gamma} H_i(\lambda, t)^\gamma, \quad i = 1, 2. \quad (40)$$

Then, from the first order conditions with respect to consumption we obtain

$$c_{it}^* = \frac{W_{it}}{H_{it}}, \quad i = 1, 2, \quad (41)$$

where H_{it} is a shorthand notation for $H_i(\lambda, t)$, and hence, functions H_{it} can be interpreted as the *wealth-consumption ratio* of investor i . By substituting indirect utility functions (40) into HJB equations it can be verified that wealth-consumption ratios satisfy the following PDEs:

$$\frac{\partial H_{it}}{\partial t} + \frac{\lambda_t^2 \sigma_{\lambda t}^2}{2} \frac{\partial^2 H_{it}}{\partial \lambda_t^2} - \lambda_t \left(\mu_{\lambda t} + \frac{1-\gamma}{\gamma} \kappa_{it} \sigma_{\lambda t} \right) \frac{\partial H_{it}}{\partial \lambda_t} + \left(\frac{1-\gamma}{2\gamma} \kappa_{it}^2 + (1-\gamma)r_{it} - \rho \right) \frac{H_{it}}{\gamma} + 1 = 0, \quad i = 1, 2, \quad (42)$$

where r_{it} and κ_{it} denote riskless rate and price of risk in a fictitious economy and are defined in (15) in terms of adjustments given in (39). Moreover, optimal investment policies for investors 1 and 2 are given by

$$\theta_{it} = \frac{1}{\gamma\sigma_t} \left(\kappa_{it} - \gamma\sigma_{\lambda t} \frac{\partial H_{it}}{\partial \lambda_t} \frac{\lambda_t}{H_{it}} \right), \quad i = 1, 2. \quad (43)$$

Since the horizon is infinite we will look for *time-independent* and bounded solutions of equations (42). Moreover, throughout this Section we assume that $\bar{\theta} \leq 1$. We note that if investor 2 faces borrowing constraint, i.e. $\bar{\theta} \geq 1$, the equilibrium coincides with the equilibrium in an unconstrained economy in which the investors, being identical, optimally choose $\theta_{it}^* = 1$.

Conveniently, since the fictitious economy is complete, the equations for wealth-consumption ratios in (42) are linear if volatilities σ and σ_λ are known. However, in equilibrium these volatilities themselves depend on wealth-consumption ratios H_i . The stock return volatility σ can be obtained by applying Itô's Lemma to stock price $S_t = R_t \delta_t$, where R_t is a shorthand notation for the *stock price-dividend ratio* which can be expressed in terms of wealth-consumption ratios from the market clearing conditions in (9). Furthermore, the volatility σ_λ can be obtained from the complementary slackness condition (30). The following Proposition 1.2 summarizes our results and provides a characterization of equilibrium in terms of wealth-consumption ratios.

Proposition 1.2. *If there exists an equilibrium, the riskless interest rate r , market price of risk κ and drift μ_λ of weighting process λ that follows (20) are given by*

$$r_t = \bar{r} - \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \bar{\theta} \sigma_t \sigma_{\lambda t} - \frac{1 + \gamma}{2\gamma} \frac{\lambda_t^{1/\gamma}}{(1 + \lambda_t^{1/\gamma})^2} \sigma_{\lambda t}^2, \quad (44)$$

$$\kappa_t = \bar{\kappa} + \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \sigma_{\lambda t}, \quad (45)$$

$$\mu_{\lambda t} = \gamma \sigma_\delta \sigma_{\lambda t} - \bar{\theta} \sigma_t \sigma_{\lambda t} - \frac{1}{1 + \lambda_t^{1/\gamma}} \sigma_{\lambda t}^2, \quad (46)$$

where \bar{r} is the riskless rate and $\bar{\kappa}$ is the market price of risk in an unconstrained economy,

given by

$$\bar{r} = \rho + \gamma\mu_\delta - \frac{\gamma(1+\gamma)}{2}\sigma_\delta^2, \quad \bar{\kappa} = \gamma\sigma_\delta. \quad (47)$$

Optimal consumptions c_i^* , wealths W_i^* , stock price-dividend ratio R and optimal investment policies θ_i^* are given by

$$c_{1t}^* = \frac{1}{1 + \lambda_t^{1/\gamma}}\delta_t, \quad c_{2t}^* = \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}}\delta_t, \quad (48)$$

$$W_{1t}^* = H_{1t} \frac{1}{1 + \lambda_t^{1/\gamma}}\delta_t, \quad W_{2t}^* = H_{2t} \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}}\delta_t, \quad (49)$$

$$R_t = H_{1t} \frac{1}{1 + \lambda_t^{1/\gamma}} + H_{2t} \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}}, \quad (50)$$

$$\theta_{1t}^* = \frac{1}{\gamma\sigma_t} \left(\kappa_t - \gamma\sigma_{\lambda t} \frac{\partial H_{1t}}{\partial \lambda_t} \frac{\lambda_t}{H_{1t}} \right), \quad \theta_{2t}^* = \bar{\theta}, \quad (51)$$

while the volatilities of the stock returns, σ , and weighting process, σ_λ , are given by

$$\sigma_t = \sigma_\delta - \sigma_{\lambda t} \frac{\partial R_t}{\partial \lambda_t} \frac{\lambda_t}{R_t}, \quad \sigma_{\lambda t} = \frac{(1 - \bar{\theta})\gamma\sigma_\delta}{\frac{1}{1 + \lambda_t^{1/\gamma}} + \gamma \frac{\partial H_{2t}}{\partial \lambda_t} \frac{\lambda_t}{H_{2t}} - \bar{\theta}\gamma \frac{\partial R_t}{\partial \lambda_t} \frac{\lambda_t}{R_t}}, \quad (52)$$

where wealth-consumption ratios H_{1t} and H_{2t} satisfy equations (42). Moreover, the initial value λ_0 for the weighting process (20) solves equation

$$sH_2(\lambda_0, 0) \frac{\lambda_0^{1/\gamma}}{1 + \lambda_0^{1/\gamma}}\delta_0 - (1 - s)H_1(\lambda_0, 0) \frac{1}{1 + \lambda_0^{1/\gamma}}\delta_0 = b. \quad (53)$$

The expressions for riskless rate r and price of risk κ in Proposition 1.2 are in terms of the of volatilities σ and σ_λ , as well as parameter $\lambda^{1/\gamma}$ which in our economic setting can be interpreted as the ratio of consumptions of investors 2 and 1, as it follows from the expressions in (48). As in the general case in Proposition 1.1, interest rates are comprised of three terms, where the first term is a riskless rate in an unconstrained economy, while the second and third terms highlight the impact of constraints and risk sharing. Moreover, the effect of risk sharing, as captured by volatility σ_λ , can be expressed in terms of consumption volatilities. In particular, from expression (34) it follows that

$$\sigma_{\lambda t} = \gamma(\sigma_{c_{1t}} - \sigma_{c_{2t}}). \quad (54)$$

It will be demonstrated later that volatility σ_λ is positive in equilibrium since investor 1 is more exposed to risk and hence her consumption growth is more volatile.

Proposition 1.2 also demonstrates that when $\bar{\theta} < 1$ the portfolio constraint of investor 2 is always binding since otherwise, having identical preferences, both investors should find optimal to invest $\theta_i < 1$ which contradicts market clearing conditions (9). Moreover, Proposition 1.2 provides expressions for equilibrium volatilities σ and σ_λ in terms of the elasticities of wealth-consumption and price-dividend ratios with respect to weighting process λ , given by

$$\epsilon_{H_2t} = \frac{\partial H_{2t}}{\partial \lambda_t} \frac{\lambda_t}{H_{2t}}, \quad \epsilon_{Pt} = \frac{\partial R_t}{\partial \lambda_t} \frac{\lambda_t}{R_t}. \quad (55)$$

From the expression for the volatility σ_λ in (52) it follows that σ_λ is decreasing in elasticity ϵ_{H_2} and increasing in ϵ_P . The effect of elasticities in (55) on volatility σ_λ then determines their impact on all the other parameters in equilibrium.

To understand the effect of these elasticities on volatility σ_λ we observe that elasticity ϵ_{H_2} is proportional to the stock hedging demand of investor 2 given by the second term in the expression for optimal policy (43). Moreover, since σ_λ is positive, it follows from this expression that higher elasticity ϵ_{H_2} tends to decrease optimal investment in stock. Thus, higher ϵ_{H_2} makes the stock less attractive, and hence reduces the cost of being constrained. Therefore, σ_λ also decreases to reflect decreased risk sharing distortions of the constraint. Moreover, as follows from the expressions for volatilities (52) the increase in elasticity ϵ_P tends to decrease stock volatility σ since the dividends and weighting process are negatively correlated. Hence, if volatility σ decreases, the stock becomes more attractive for both investors. However, since investor 2 is constrained, her ideal unconstrained holding moves further away from her constrained holding $\bar{\theta}$ and hence the risks are shared in a less optimal way and σ_λ increases.

Proposition 1.2 also allows to explicitly identify the coefficients of PDEs (42) for wealth-consumption ratios H_i , which depend on equilibrium parameters identified in expressions (44)–(52). Moreover, it appears that the coefficients themselves depend on ratios H_i and

hence, we obtain a system of quasilinear PDEs the solutions to which completely characterize the equilibrium. We next solve for time-independent solutions of PDEs (42) which correspond to the infinite horizon case. To solve the equations (42), we first fix a large horizon parameter T , choose a starting value for $H_i(\lambda, T)$ and then solve the equation backwards using a modification of Euler's finite-difference method until the solution converges to a stationary one. This approach is similar to the subsequent iterations method for solving Bellman equations in discrete time (e.g., Ljungqvist and Sargent, 2004) when at a distant time in the future the value function is set equal to some function (usually zero) and then the value functions at earlier dates are obtained by solving equations backwards.

Since weight λ varies from zero to infinity, we first perform a change of variable and rewrite the PDEs (42) as well as the equilibrium parameters in Proposition 1.2 in terms of *constrained investor's share in aggregate consumption*, given by

$$y_t = \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}}. \quad (56)$$

Variable y takes values in the interval $[0, 1]$ and provides one-to-one mapping to variable λ . The solution of PDEs in terms of new variable we label as $\tilde{H}_i(y, t)$. Assuming that the solutions to new PDEs are continuous and twice continuously differentiable, setting in those equations $y = 0$ and $y = 1$ we recover boundary conditions for $\tilde{H}_i(y, t)$. Next, we replace the derivatives by their finite-difference analogues letting the time and state variable increments denote $\Delta t \equiv T/M$ and $\Delta y \equiv 1/N$, where M and N are integer numbers. Solving the equation backwards, sitting at time t we compute the coefficients of finite-difference analogues of PDEs (42) using the solutions $\tilde{H}_i(y, t + \Delta t)$ obtained from the previous step $t + \Delta t$. As a result, the coefficients of equations for $\tilde{H}_i(y, t)$ are known at time t and hence $\tilde{H}_i(y, t)$ can be found by solving a system of linear finite-difference equations with three-diagonal matrix. Appendix B provides further details of the numerical algorithm. The wealth-consumption ratios then allow us to derive all the parameters of equilibrium.

Remark 1 (Bond prices). Proposition 1.2 allows to determine the instantaneous interest rate r_t . Therefore, the bond price B_t can be obtained by solving numerically the equation

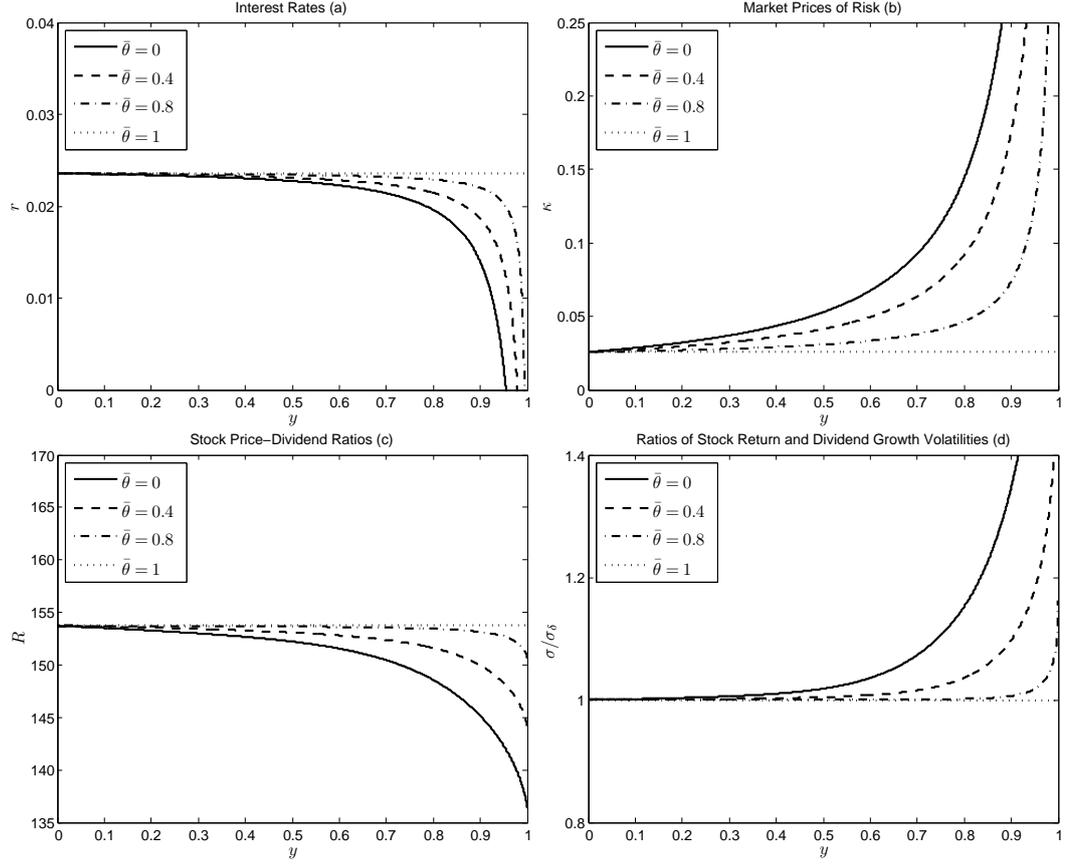


Figure 1.1: Parameters of Equilibrium with Constraints, $\gamma < 1$.

The figure plots interest rates r , market prices of risk κ , price-dividend ratios R and ratios of stock return and dividend growth volatilities σ/σ_δ as functions of constrained investor's consumption share y . Dividend mean growth rate $\mu_\delta = 1.8\%$ and volatility $\sigma_\delta = 3.2\%$ are from the estimates in Campbell (2003), based on consumption data in 1891–1998, while risk aversion and time discount are set to $\gamma = 0.8$ and $\rho = 0.01$.

for the bond price dynamics (2).

Remark 2 (Existence of Equilibrium). The numerical analysis shows that the function on the left-hand side of the equation for λ_0 in (53) is a monotone function of λ_0 and maps interval $[0, \infty)$ into $[C_0, C_1)$, where C_0 and C_1 are some constants, and hence, if $b \in [C_0, C_1)$ there always exists the unique solution λ_0 that satisfies the equation. Given the existence of λ_0 and the solutions to HJB equations (42), expressions (44)–(52) fully characterize the equilibrium in the economy.

1.3.2. Analysis of Equilibrium

We now study the impact of constraints on various equilibrium parameters. Important implication of our model is that in contrast to models with logarithmic investors the constraints do affect the price-dividend ratios and stock return volatilities. Figures 1.1 and 1.2 present equilibrium interest rates, market prices of risk, price-dividend ratios and the ratios of stock return and dividend growth volatilities as functions of constrained investor's consumption share y for different levels of the tightness of constraints $\bar{\theta}$ when risk aversions are less than unity ($\gamma = 0.8$) and greater than unity ($\gamma = 3$), respectively. The equilibrium is derived under plausible parameters for the dividend process.⁸ We note that in our model the instantaneous changes in the dividend growth $d\delta/\delta$ and constrained investor's consumption share dy are negatively correlated since negative shocks to dividends shift relative consumption to constrained investors, due to the fact that the latter are less affected by adverse stock market fluctuations. Hence, higher consumption share y is associated with bad times while lower y is associated with good times. Following the literature (e.g., Chan and Kogan, 2002) we label economic variables as *procyclical* if they increase in good times (when dividend growth rate shocks are positive) and decrease in bad times (when dividend growth rate shocks are negative), and as *countercyclical* if they decrease in good times and increase in bad times.

For risk aversion less than unity Figure 1.1 demonstrates that tighter constraints decrease interest rates and price-dividend ratios, and increase market prices of risk and stock return volatilities. For risk aversion greater than unity, Figure 1.2 shows that tighter constraints decrease interest rates and stock return volatilities, and increase market prices of risk and price-dividend ratios. In both cases the impact of constraints is asymmetric and is more pronounced in bad times, when consumption share y is larger. We first analyze the equilibrium parameters for the case $\gamma < 1$, presented on Figure 1.1, and then for the case $\gamma > 1$, presented on Figure 1.2.

⁸In particular, the parameters for the dividend process ($\mu_\delta = 1.8\%$, $\sigma_\delta = 3.2\%$) are taken from the estimates in Campbell (2003), based on consumption data in 1891–1998 years, and the discounting parameter is set to $\rho = 0.01$.

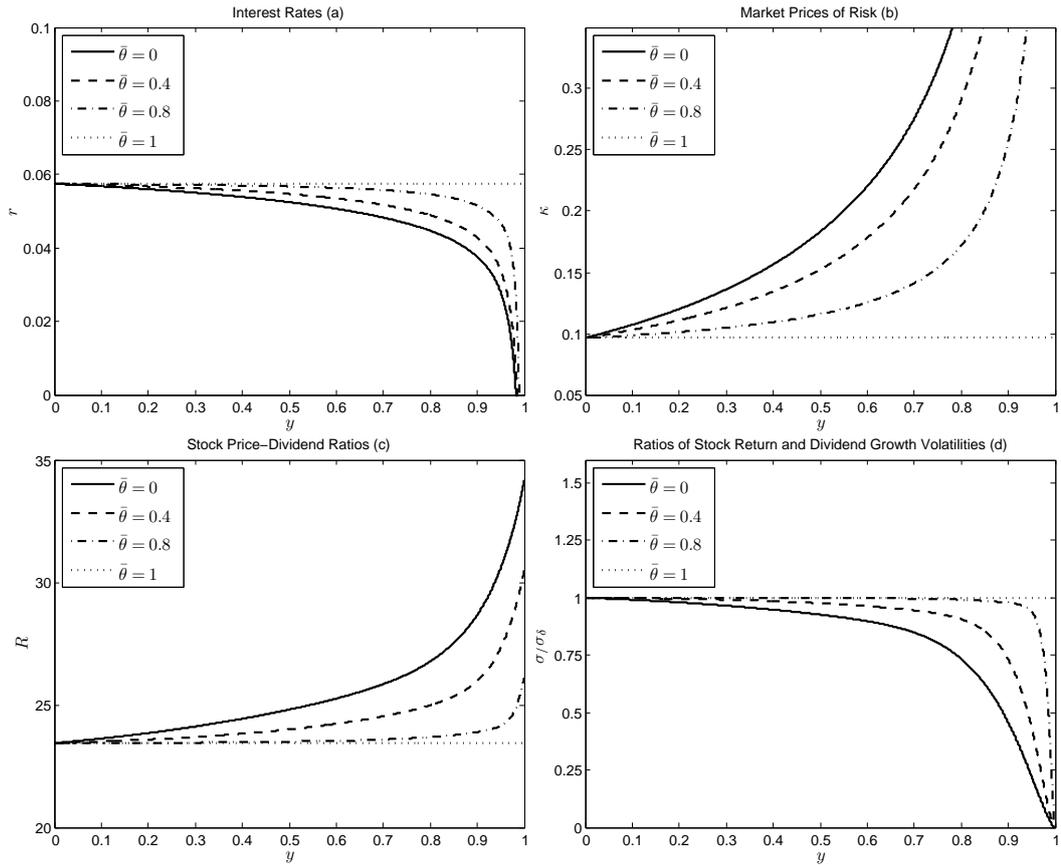


Figure 1.2: Parameters of Equilibrium with Constraints, $\gamma > 1$.

The figure plots interest rates r , market prices of risk κ , price-dividend ratios R and ratios of stock return and dividend growth volatilities σ/σ_δ as functions of constrained investor's consumption share y . Dividend mean growth rate $\mu_\delta = 1.8\%$ and volatility $\sigma_\delta = 3.2\%$ are from the estimates in Campbell (2003), based on consumption data in 1891–1998, while risk aversion and time discount are set to $\gamma = 3$ and $\rho = 0.01$.

Panel (a) of Figure 1.1 presents interest rates when $\gamma < 1$ and demonstrates that in line with the results of Section 1.2 interest rates in constrained economy are lower than in an unconstrained one for a given consumption share y . Moreover, they become lower with tighter constraints and are decreasing functions of constrained investor's share of consumption y . Intuitively, constrained investor invests more in bonds driving interest rates down. Moreover, constraints prevent the investor from sharing risks efficiently and smoothing consumption over time. As a result, when her current consumption is high the price of future consumption increases making her more willing to lend at a lower interest causing interest rates to fall.

Panel (b) of Figure 1.1 shows that the prices of risk are higher in the constrained than in the unconstrained economies and increase as constraint becomes tighter. When the constrained investor invests only a fraction $\bar{\theta} < 1$ of her wealth in the stock, for the markets to clear investor 1 should be leveraged so that $\theta_1^* > 1$. This, however, implies that the unconstrained investor should be more exposed to risk as the constraint tightens, and hence, the market price of risk should be higher. Moreover, market price of risk also increases with constrained investor's consumption share y since in those states in which the unconstrained investor consumes less and possesses less wealth she is more risk averse and requires market prices of risk to increase for the stock market to clear. Thus, the market price of risk is countercyclical, consistently with the empirical literature (e.g., Ferson and Harvey, 1991).

Panel (c) of Figure 1.1 demonstrates that the price-dividend ratios become lower with tighter constraints and the effect of constraints is more pronounced in states with higher constrained investor's consumption share y . To understand the patterns of price-dividend ratios we first observe that in equilibrium the price-dividend ratio can be interpreted as the ratio of aggregate wealth over aggregate consumption since the market clearing conditions (9) imply that the stock price equals aggregate wealth while the aggregate consumption equals the dividend. As a result, the price-dividend ratio will be close to wealth-consumption ratio of unconstrained or constrained investor depending on which of them dominates in the market by holding larger fraction of consumption and wealth. When the unconstrained investor dominates (y is low), the equilibrium will be close to that in the benchmark unconstrained economy in which case all equilibrium parameters, including price-dividend ratios, are constant (dotted lines in Figures 1.1 and 1.2). However, in states with dominating constrained investor (y is high) the price-dividend ratio is close to constrained investor's wealth-consumption ratio, which increases or decreases with tighter constraints depending on the relative strength of classical income and substitution effects. When the investment opportunities worsen, the income effect induces investors to decrease consumption and save more while the substitution effect induces them to consume more and save less due to cheaper current consumption. For CRRA preferences the intertemporal elasticity of substi-

tution (IES) equals $1/\gamma$, the income effect dominates for $IES < 1$ and the substitution effect dominates for $IES > 1$ while in the case of $IES = 1$ both effects perfectly offset each other. With tighter constraints the investment opportunities for constrained investor worsen due to the decline in interest rates and inability to fully benefit from the increase in market prices of risk, and hence her wealth-consumption ratios decrease for $\gamma < 1$ via the substitution effect.⁹ As a result, the price-dividend ratios decrease with tighter constraints and the effect is stronger in bad times, when constrained investor dominates the market and the decline in interest rates is sharper.

The stock return volatilities on panel (d) of Figure 1.1 increase with tighter constraints and are higher in bad times (when y is high) than in good times (when y is low). This is due to the fact that the instantaneous changes in price-dividend ratio R and dividend δ are positively correlated due to the fact that ratio R is a decreasing function of consumption share y , which is negatively correlated with changes in dividend δ . Consequently, since the stock price is the product of price-dividend ratio and the dividend, stock return volatility should be higher in constrained economy. Moreover, this effect is stronger in bad times (when y is high) due to the concavity of ratio R , and when $\bar{\theta}$ is low, due to the higher sensitivity of ratio R to changes in y . Thus, for $\gamma < 1$ consistently with the empirical literature (e.g., Schwert, 1989; Campbell and Cochrane, 1999) our model generates procyclical price-dividend ratios, countercyclical stock return volatilities exceeding the volatility of dividends, as well as negative correlation between changes in stock returns and their volatilities. Moreover, the results on Figure 1.1 demonstrate that lower price-dividend ratios R predict higher market prices of risk κ as well as higher risk premia (given by $\mu - r = \kappa\sigma$).

⁹The relation between wealth-consumption ratios and the attractiveness of investment opportunities can conveniently be illustrated in an unconstrained partial equilibrium economy with constant interest rate r and market price of risk $\kappa = (\mu - r)/\sigma$, and an investor maximizing her objective (6) subject to budget constraint (4) and no-bankruptcy constraint. It can easily be verified that when condition $\rho - (1 - \gamma)(r + 0.5\kappa^2/\gamma) > 0$ is satisfied, the investor's wealth-consumption ratio is given by:

$$\frac{W}{c} = \frac{\gamma}{\rho - (1 - \gamma)(r + 0.5\kappa^2/\gamma)},$$

Hence, if investment opportunities deteriorate due to decrease of r or κ , the wealth-consumption ratio increases if the income effect dominates ($\gamma > 1$) and decreases if the substitution effect dominates ($\gamma < 1$).

Turning to the case $\gamma > 1$ we observe from the results shown on Figure 1.2 that the constraints affect the interest rates and market prices of risk in the same directions as in the case $\gamma < 1$. However, by contrast with the case of $\gamma < 1$, due to the dominance of income effect, price-dividend ratios increase while stock return volatilities decrease with tighter constraints, and the effects are stronger in bad times.¹⁰ One might think that the results in the case $\gamma > 1$ are more plausible than in the case $\gamma < 1$ given the evidence (e.g., Mehra and Prescott, 1985) that risk aversion is greater than unity. However, we note, that the intuition for the dynamics of price-dividend ratios and stock return volatilities in our model is driven by the relative strength of income and substitution effect and not by the risk aversion per se. It is well known that CARA utility does not allow to separate IES from the risk aversion and hence, in our setting $\text{IES} > 1$ is necessarily associated with $\gamma < 1$.

We also note that since lower $\bar{\theta}$ decreases interest rates and increases market prices of risk, irrespective of risk aversion γ , the case of restricted participation which corresponds to $\bar{\theta} = 0$ better explains the levels of observed interest rates and market prices of risk. In particular, in our model with plausible parameters described above and $\gamma = 3$, when we set $y = 0.7$ (e.g., Mankiw and Zeldes, 1991; Guvenen, 2006) we obtain $r = 4.8\%$ and $\kappa = 28\%$, while the volatilities of individual consumptions are $\sigma_{c_1} = 9\%$ and $\sigma_{c_2} = 0.7\%$. The estimates in Campbell (2003) show that $r = 2\%$ and $\kappa = 36\%$, while Malloy, Moskowitz, and Vissing-Jorgensen (2009) show that $\sigma_{c_1} = 3.6\%$ and $\sigma_{c_2} = 1.4\%$. Thus, our model implies riskless rates and market prices of risk sufficiently close to those in the data for such a simple model.

Remark 3 (Duffie-Epstein preferences). The discussion above demonstrates that for risk aversion $\gamma < 1$ the model generates empirically plausible patterns for price-dividend ratios and stock return volatilities while for $\gamma > 1$ it generates high market prices of risk and low interest rates close to those observed in the data. We note that the intuition for price-dividend ratios and stock return volatilities only relies on the relative strength of income and

¹⁰In our model when $\gamma > 1$ the instantaneous volatility of stock returns is lower than that of dividend growth and hence there is no excess volatility. Bhamra and Uppal (2009) demonstrate a significant excess volatility in a complete-market exchange economy with CRRA investors that differ in risk aversions. Thus, excess volatility is likely to be present in the extension of our model to the case where investors have different risk aversions.

substitution effects. As pointed out above, for CRRA preferences the intertemporal elasticity of substitution (IES) equals $1/\gamma$ and hence high IES leading to the dominance of substitution effect is only possible for $\gamma < 1$. However, more general Duffie-Epstein recursive preferences allow for IES independent of risk aversion parameter γ (Duffie and Epstein, 1992). Our results lead to a conjecture that in a model with Duffie-Epstein preferences with both IES and risk aversion exceeding unity (as in Bansal and Yaron, 2004) it might be possible to match interest rates and market prices of risk, as well as generate procyclical price-dividend ratios and countercyclical stock return volatilities which exceed the volatility of dividends, consistently with the empirical literature.¹¹

Our results also allow to obtain the expressions for consumption growth volatilities of investors, which also capture the effect of risk sharing between them. The expressions for the volatilities can be obtained by applying Itô's Lemma to optimal consumptions (48) and are reported in the following Corollary 1.2.

Corollary 1.2. *The optimal consumption growth volatilities of unconstrained and constrained investors are given by*

$$\sigma_{c_1t} = \sigma_\delta + \frac{1}{\gamma} \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \sigma_{\lambda t}, \quad \sigma_{c_2t} = \sigma_\delta - \frac{1}{\gamma} \frac{1}{1 + \lambda_t^{1/\gamma}} \sigma_{\lambda t}. \quad (57)$$

It can be shown in our example that the volatility σ_λ is positive, and hence, consumption volatilities in (57) imply that unconstrained investor, being exposed to more risk, has larger volatility of consumption than the constrained one. Basak and Cuoco (1998) show in the case of restricted participation and $\gamma = 1$ that the volatility σ_{c_2} of constrained investor is zero and all the risk is borne by the unconstrained investor. However, in our case with $\gamma > 1$ volatility σ_{c_2} is greater than zero, as also in the data for non-stockholders (e.g. Malloy, Moskowitz, Vissing-Jorgensen, 2009).

¹¹Campbell and Cochrane (1999) and Chan and Kogan (2002) present the models with habit formation and “catching up with the Joneses” preferences respectively, that explain the patterns for price-dividend ratios and stock return volatilities.

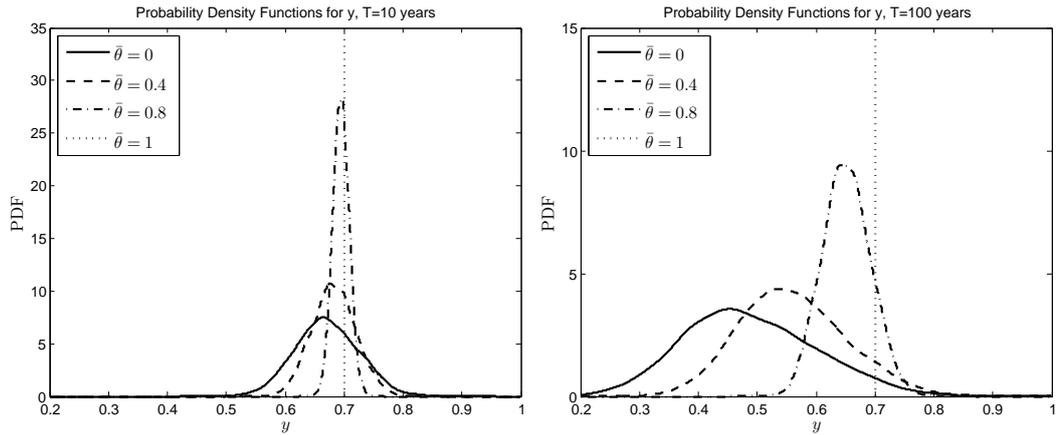


Figure 1.3: Probability Density Functions for Constrained Investor's Share in Aggregate Consumption, $\gamma = 3$.

The figure plots interest rates r , market prices of risk κ , price-dividend ratios R and ratios of stock return and dividend growth volatilities σ/σ_δ as functions of constrained investor's consumption share y . Dividend mean growth rate $\mu_\delta = 1.8\%$ and volatility $\sigma_\delta = 3.2\%$ are from the estimates in Campbell (2003), based on consumption data in 1891–1998, while risk aversion and time discount are set to $\gamma = 3$ and $\rho = 0.01$.

Finally, we address the question of how the constraints affect the distribution of consumption between the investors. So far we have compared the parameters of equilibria with different constraint $\bar{\theta}$ for a given level of consumption share y . This comparison does not account for the fact that share y itself depends on $\bar{\theta}$. Figure 1.3 shows probability density functions of y for $\gamma = 3$, different constraints $\bar{\theta}$ and time horizons equal to ten and one hundred years respectively. The probability densities imply that consumption share y tends to decline, and hence, the impact of constrained investor becomes smaller in the course of time even though it is still significant even after hundred years. As discussed in Hong, Kubik and Stein (2004) stock market participation depends on person-specific characteristics such as social integrations and education. Thus, specializing to the case of restricted participation ($\bar{\theta} = 0$) our model demonstrates that these characteristics lead to gradual, although slow,

elimination of non-stockholders' impact on financial markets via natural selection.¹²

1.4. Extensions and Ramifications

In this Section we demonstrate that our model is extendable to different alternative economic settings. Section 1.4.1 extends the results of Section 1.2 to the case of heterogeneous beliefs and provides a numerical solution to the model with CRRA investors with heterogeneous beliefs when one of them faces short-sale constraints. Section 1.4.2 demonstrates that the results of Section 1.2 generalize to the environments with multiple assets.

1.4.1. Heterogeneous Beliefs Formulation

We now consider an economy in which investors are constrained and have different beliefs about mean dividend growth rate in the economy. We first generalize the results of Section 1.2 and derive expressions for the parameters of equilibrium in terms of adjustments in fictitious economy and the differences in beliefs. Then, we specialize to a framework in which both investors have identical CRRA preferences and the pessimist faces short-sale constraints. We solve this model numerically by employing the approach of Section 1.3 and discuss some properties of the equilibrium parameters.

Basak (2000, 2005) derives expressions for equilibrium parameters for general utility functions in the economy in which investors face heterogeneous belief but does not study the impact of constraints as we do in this work. Our model is also related to the model of Gallmeyer and Hollifield (2008) in which the pessimist has logarithmic preferences and faces short-sale constraints while the investor with general CRRA is optimistic and unconstrained. By contrast, our model does not rely on the assumption of a logarithmic constrained investor.

The economic setting is similar to that of Section 1.2. In particular, investors trade in

¹²In unconstrained economic settings the survival of irrational investors has been studied in Kogan, Ross, Wang and Westerfield (2004), Berrada (2009), Dumas, Kurshev and Uppal (2009) and Yan (2008), among others. The results in the latter three works suggest that it takes a long time to eliminate the impact of irrational investors that have wrong beliefs about mean dividend growth rates. Hugonnier (2008) considers survival of constrained logarithmic investor and demonstrates that their impact can quickly be eliminated. However, in his calibration the volatility of dividends is 20% while we set this parameter to the volatility of aggregate consumption 3.2% taken from Campbell(2003). When in the calibration we choose $\gamma = 1$ and $\sigma_\delta = 20\%$ consistently with Hugonnier our results also imply fast elimination of constrained investor's impact.

two securities, a riskless bond and stock, and dividends follow process (1). They agree on dividends, bond and stock prices and the dividend growth rate volatility σ_δ but disagree on the growth rate μ_δ . Throughout this Section we will be using superscript i to denote quantities on which investors disagree, while by subscript i investor-specific quantities on which there is no disagreement. Investors update their beliefs $\mu_{\delta t}^i$ in a Bayesian fashion:

$$\mu_{\delta t}^i = E^i[\mu_{\delta t} | \mathcal{F}_t^\delta], \quad i \in \{o, p\}, \quad (58)$$

where $E^i[\cdot]$ denotes the expectation under the subjective probability measure of investor i and \mathcal{F}_t^δ is the augmented filtration generated by δ_t . Both investors have different priors $\mu_{\delta 0}^i$ and investor 1 is optimistic ($i = o$) while investor 2 is pessimistic ($i = p$) about the dividend growth. From the point of view of investor i the dividends and stock prices follow the processes

$$d\delta_t = \delta_t[\mu_{\delta t}^i + \sigma_{\delta t}dw_t^i], \quad (59)$$

$$dS_t + \delta_t dt = S_t[\mu_t^i dt + \sigma_t dw_t^i], \quad (60)$$

where w_t^i denotes Brownian motions under the *subjective probability measure* of investor i .

From the filtering theory in Lipster and Shiriyayev (1977) it follows that Brownian motions w_t^i are given by

$$dw_t^i = \frac{\mu_\delta - \mu_{\delta t}^i}{\sigma_\delta} dt + dw_t, \quad i \in o, p. \quad (61)$$

By $\Delta\mu_{\delta t}$ we denote the *disagreement process* defined as

$$\Delta\mu_{\delta t} = \frac{\mu_{\delta t}^o - \mu_{\delta t}^p}{\sigma_{\delta t}}. \quad (62)$$

Moreover, if dividends follow geometric Brownian motion (35) and investors' initial priors are normally distributed with parameters

$$\mu_\delta^i \sim N(\hat{\mu}_{\delta 0}^i, \hat{\sigma}_{\delta 0}^i),$$

then $\mu_{\delta t}^i$ is also normally distributed and the processes for $\mu_{\delta t}^i$ and $\Delta\mu_{\delta t}$ are given by

$$d\mu_{\delta t}^i = \frac{\hat{\sigma}_{\delta t}^i}{\sigma_\delta} dw_t^i, \quad (63)$$

$$d\Delta\mu_{\delta t} = -\frac{\hat{\sigma}_{\delta t}^p}{\sigma_\delta} \Delta\mu_{\delta t} dt + \frac{\hat{\sigma}_{\delta t}^o - \hat{\sigma}_{\delta t}^p}{\sigma_\delta} dw_t^i, \quad (64)$$

where

$$\hat{\sigma}_{\delta t}^i = \frac{\hat{\sigma}_{\delta 0}^i \sigma_{\delta}^2}{\hat{\sigma}_{\delta 0}^i t + \sigma_{\delta}^2}. \quad (65)$$

The budget constraint for each investor is given by (4) in which Brownian motion w and stock mean-return μ are replaced by investor's subjective Brownian motion w^i and mean-return μ^i . Each investor solves optimization problem (6) in which now expectation operator $E[\cdot]$ is replaced by operator $E^i[\cdot]$ under investor's subjective beliefs, subject to the budget constraint, no-bankruptcy constraint $W_t \geq 0$ and portfolio constraints (5).

The equilibrium in this economy is a set of parameters $\{r_t, \mu_t^o, \mu_t^p, \sigma_t\}$ and of consumption and investment policies $\{c_{it}^*, \alpha_{it}^*, \theta_{it}^*\}_{i \in \{o,p\}}$ which solve investor i 's dynamic optimization problem and satisfy market clearing conditions in (9).

As in Section 1.2, the parameters of equilibrium are characterized in terms of adjustments ν_i^* and support functions $f_i(\nu_i^*)$. We first characterize investor's marginal utilities in terms of state prices that follow processes as in (14) but with Brownian motions under subjective probability measures. Then, we introduce the ratio of their marginal utilities λ , which follows the process

$$d\lambda_t = -\lambda_t[\mu_{\lambda t}^i dt + \sigma_{\lambda t} dw_t^i]. \quad (66)$$

By employing market clearing conditions we obtain the parameters of equilibrium. Proposition 1.3 summarizes our results.

Proposition 1.3. *If there exists an equilibrium, the riskless interest rate r , perceived market prices of risk κ^i , drifts μ_{λ}^i and volatility σ_{λ} of weighting process (66) are given by*

$$\begin{aligned} r_t = & \bar{r}_t - \frac{A_t}{A_{ot}} f_o(\nu_{ot}^*) - \frac{A_t}{A_{pt}} f_p(\nu_{pt}^*) - \frac{A_t^3 (P_{ot} + P_{pt})}{2A_{ot}^2 A_{pt}^2} \sigma_{\lambda t}^2 - \frac{A_t^3}{A_{ot} A_{pt}} \left(\frac{P_{ot}}{A_{ot}} - \frac{P_{pt}}{A_{pt}} \right) \delta_t \sigma_{\delta t} \sigma_{\lambda t} \\ & - \frac{A_t^2}{A_{pt}} \delta_t \sigma_{\delta t} \Delta \mu_{\delta t} + \frac{A_t^2}{A_{ot} A_{pt}} \sigma_{\lambda t} \Delta \mu_{\delta t}, \end{aligned} \quad (67)$$

$$\kappa_t^o = \bar{\kappa}_t + \frac{A_t}{A_{pt}} \sigma_{\lambda t}, \quad \kappa_t^p = \bar{\kappa}_t - \frac{A_t}{A_{ot}} \sigma_{\lambda t}, \quad (68)$$

$$\mu_{\lambda t}^o = A_t \delta_t \sigma_{\delta t} \sigma_{\lambda t} + f_o(\nu_{ot}^*) - f_p(\nu_{pt}^*) - \frac{A_t}{A_{ot}} \sigma_{\lambda t}^2, \quad \mu_{\lambda t}^p = \mu_{\lambda t}^o - \sigma_{\lambda t} \Delta \mu_{\delta t}, \quad (69)$$

$$\sigma_{\lambda t} = \Delta\mu_{\delta t} + \frac{\nu_{ot}^* - \nu_{pt}^*}{\sigma_t}, \quad (70)$$

where \bar{r} is the riskless rate and $\bar{\kappa}$ is the market price of risk in an unconstrained economy populated by optimists, given by

$$\bar{r}_t = \rho + A_t\delta_t\mu_{\delta t}^o - \frac{A_tP_t}{2}\delta_t^2\sigma_{\delta t}^2, \quad \bar{\kappa}_t = A_t\delta_t\sigma_{\delta t}, \quad (71)$$

A_{it} , P_{it} , and A_t and P_t are absolute risk aversions and prudence parameters of investor i and a representative investor with utility (19), respectively.

Expressions for optimal consumption c_i^* and stock price S are as in Proposition 1.1. Optimal wealths W_i^* and optimal investment policies θ_i^* are given by expressions (26) and (28) in which expectation operator $E[\cdot]$ and market prices of risk κ are replaced by subjective operator $E^i[\cdot]$ and price of risk κ^i . Initial value λ_0 for weighting process (66) is such that budget constraint at time zero (29) is satisfied. Moreover, adjustments ν_i^* satisfy complementary slackness conditions (30), as in Proposition 1.1.

The expressions for interest rates in Proposition 1.3 demonstrate the impact of heterogeneous beliefs on interest rates and subjective market prices of risk. In particular, the expression for interest rates have additional terms (last two terms in (67)) which demonstrate the direct effect of disagreement process $\Delta\mu_{\delta}$. Since the disagreement process is positive, its impact depends on the sign of volatility σ_{λ} . Moreover, the expression for volatility σ_{λ} in (70) demonstrates that this parameter itself depends on $\Delta\mu_{\delta}$ since the disagreement affects the efficiency of the risk sharing, quantified by σ_{λ} . Unlike the setup of Section 1.2, investors now disagree also on the market prices of risk, which are given in (68).

We now consider a modification of the model in Section 1.3 in which now investors have heterogeneous beliefs about the dividend growth rate. In particular, investor 1 is optimistic and unconstrained while investor 2 is pessimistic and faces constraints that impose a limit on the short-sales $\theta \geq \underline{\theta}$, where $\underline{\theta} < 0$. For simplicity, as in Yan (2008) we assume that investors do not update their beliefs and believe that dividends follow a GBM

$$d\delta_t = \delta_t[\mu_{\delta}^i dt + \sigma_{\delta} dw_t^i], \quad (72)$$

and their difference in beliefs we denote by $\Delta\mu_\delta$. This assumption can further be justified by noting that under plausible parameters it takes very long time for the beliefs to converge.¹³

As in Section 1.3 we characterize the equilibrium in terms of the wealth-consumption ratios of investors which satisfy HJB equations (42) in which the drift parameter μ_λ is now investor-specific and should be replaced by μ_λ^i . Our results are summarized in Proposition 1.4.

Proposition 1.4. *If there exists an equilibrium, the riskless interest rate r , perceived market price of risk κ^i and drifts μ_λ^i of weighting process λ that follows (20) are given by*

$$\begin{aligned} r_t = & \bar{r} + \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \theta \sigma_t (\Delta\mu_\delta - \sigma_{\lambda t}) - \frac{1 + \gamma}{2\gamma} \frac{\lambda_t^{1/\gamma}}{(1 + \lambda_t^{1/\gamma})^2} \sigma_{\lambda t}^2 \\ & - \gamma \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \sigma_\delta \Delta\mu_\delta + \frac{\lambda_t^{1/\gamma}}{(1 + \lambda_t^{1/\gamma})^2} \sigma_{\lambda t} \Delta\mu_\delta, \end{aligned} \quad (73)$$

$$\begin{aligned} \kappa_t^o = & \bar{\kappa} + \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \sigma_{\lambda t}, \quad \kappa_t^p = \bar{\kappa} - \frac{1}{1 + \lambda_t^{1/\gamma}} \sigma_{\lambda t}, \\ \mu_{\lambda t}^o = & \gamma \sigma_\delta \sigma_{\lambda t} - \frac{1}{1 + \lambda_t^{1/\gamma}} \sigma_{\lambda t}^2 - \Delta\mu_\delta \kappa_t^p + \theta \sigma_t (\Delta\mu_\delta - \sigma_{\lambda t}), \quad \mu_{\lambda t}^p = \mu_{\lambda t}^o - \Delta\mu_\delta \sigma_{\lambda t}, \end{aligned} \quad (74)$$

where \bar{r} is the riskless rate and $\bar{\kappa}$ is the market price of risk in an unconstrained economy populated by optimists, given by

$$\bar{r} = \rho + \gamma \mu_\delta^o - \frac{\gamma(1 + \gamma)}{2} \sigma_\delta^2, \quad \bar{\kappa} = \gamma \sigma_\delta. \quad (76)$$

Optimal consumptions c_i^* , wealths W_i^* , stock price-dividend ratio R and optimal investment policies θ_i^* are given by

$$c_{ot}^* = \frac{1}{1 + \lambda_t^{1/\gamma}} \delta_t, \quad c_{pt}^* = \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \delta_t, \quad (77)$$

¹³In particular, assuming that investors have the same variances for the prior belief, $\hat{\sigma}_{\delta 0}^i = \hat{\sigma}_{\delta 0}$, equations for the disagreement and estimation error processes in (64) and (65) imply that

$$\Delta\mu_{\delta t} = \Delta\mu_{\delta 0} \left(\frac{\sigma_\delta^2}{\hat{\sigma}_{\delta 0} t + \sigma_\delta^2} \right)^{\sigma_\delta}.$$

Assuming further that $\hat{\sigma}_{\delta 0} = \sigma_\delta$ and taking $\sigma_\delta = 3.2\%$, as in Campbell (2003), we obtain that it takes 100 years for the disagreement $\Delta\mu_\delta$ to decrease by 20%.

$$W_{ot}^* = H_{ot} \frac{1}{1 + \lambda_t^{1/\gamma}} \delta_t, \quad W_{pt}^* = H_{pt} \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}} \delta_t, \quad (78)$$

$$R_t = H_{ot} \frac{1}{1 + \lambda_t^{1/\gamma}} + H_{pt} \frac{\lambda_t^{1/\gamma}}{1 + \lambda_t^{1/\gamma}}, \quad (79)$$

$$\theta_{ot}^* = \frac{1}{\gamma \sigma_t} \left(\kappa_t^o - \gamma \sigma_{\lambda t} \frac{\partial H_{ot}}{\partial \lambda_t} \frac{\lambda_t}{H_{ot}} \right), \quad \theta_{pt}^* = \frac{1}{\gamma \sigma_t} \left(\kappa_t^p - \gamma \sigma_{\lambda t} \frac{\partial H_{pt}}{\partial \lambda_t} \frac{\lambda_t}{H_{pt}} \right), \quad (80)$$

while the volatilities of the stock returns, σ , and weighting process, σ_λ , are given by

$$\sigma_t = \sigma_\delta - \sigma_{\lambda t} \frac{\partial R_t}{\partial \lambda_t} \frac{\lambda_t}{R_t}, \quad \sigma_{\lambda t} = \min \left\{ \frac{(1 - \underline{\theta}) \gamma \sigma_\delta}{\frac{1}{1 + \lambda_t^{1/\gamma}} + \gamma \frac{\partial H_{pt}}{\partial \lambda_t} \frac{\lambda_t}{H_{pt}} - \underline{\theta} \gamma \frac{\partial R_t}{\partial \lambda_t} \frac{\lambda_t}{R_t}}, \Delta \mu_\delta \right\}, \quad (81)$$

where wealth-consumption ratios H_{ot} and H_{pt} satisfy equations (42). Moreover, the initial value λ_0 for the weighting process (20) solves equation

$$s H_p(\lambda_0, 0) \frac{\lambda_0^{1/\gamma}}{1 + \lambda_0^{1/\gamma}} \delta_0 - (1 - s) H_o(\lambda_0, 0) \frac{1}{1 + \lambda_0^{1/\gamma}} \delta_0 = b. \quad (82)$$

Proposition 1.4 characterizes equilibrium parameters in terms of wealth-consumption ratios and highlights the effects of heterogeneous beliefs and short-sale constraints. Crucial difference from the results of Proposition 1.2 is that now market prices of risk (74) and the drifts of weighting process (75) are investor-specific due to investors' disagreement on the dividend growth. Moreover, the short-sale constraint will not always be binding in equilibrium since when constrained investor's share of aggregate consumption is large she becomes more willing to smooth consumption over time and invests more in stock.

By calibrating our economy to plausible parameters we find that constraints have little effect on riskless rates, while market prices of risk are investor-specific. Therefore, we here focus on price-dividend ratios and stock return volatilities which are presented on Figure 1.4 for different levels of $\underline{\theta}$. We here consider only the case $\gamma > 1$ and note that the case $\gamma < 1$ can be analyzed in a similar way. The dotted lines correspond to quantities in an unconstrained economy ($\underline{\theta} = -\infty$) which are computed using explicit formula for stock prices in terms of weighting process λ , available in Yan (2008). We assume that the optimist has correct beliefs about mean dividend growth while the pessimist underestimates it by 40%. The first picture

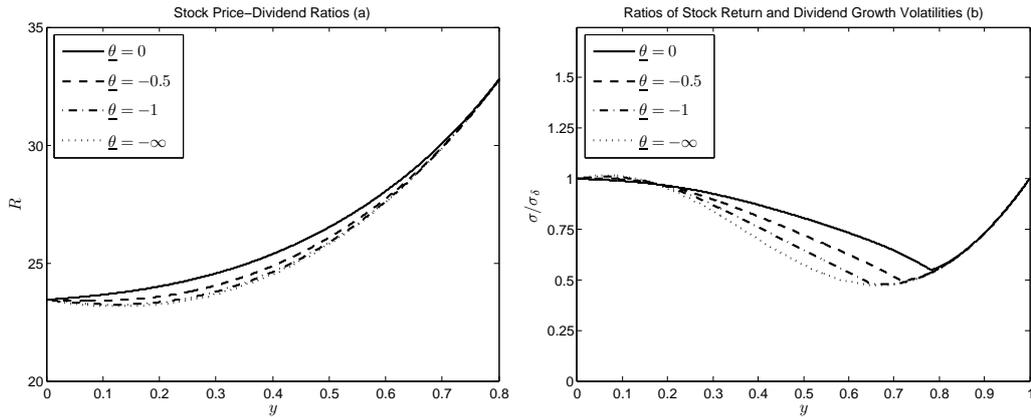


Figure 1.4: Price-Dividend Ratios and Ratios of Stock Return and Dividend Growth Volatilities with Heterogeneous Beliefs, $\gamma > 1$.

The figure plots interest rates r , market prices of risk κ , price-dividend ratios R and ratios of stock return and dividend growth volatilities σ/σ_δ as functions of constrained investor's consumption share y . Dividend mean growth rate $\mu_\delta = 1.8\%$ and volatility $\sigma_\delta = 3.2\%$ are from the estimates in Campbell (2003), based on consumption data in 1891–1998, while risk aversion and time discount are set to $\gamma = 3$ and $\rho = 0.01$.

on Figure 1.4 demonstrates that tighter short-selling constraints (higher $\underline{\theta}$) increase price-dividend ratios. In the presence of short-sale constraints the optimist should hold less stocks in equilibrium which decreases her perceived market price of risk. As a result, investment opportunities deteriorate and her wealth-consumption ratio increases due to the dominance of substitution effect. Thus, when the optimist dominates in the market, the price-dividend ratio should go up for the similar reasons as in Section 1.3. When the pessimist dominates, the constraint does not bind and the price-dividend ratio becomes closer to that in the unconstrained case.

It can also be observed that the price-dividend ratios on panel (a) of Figure 1.4 are U-shaped when $\underline{\theta}$ is low, even though this effect is not economically significant. To understand the intuition, we observe that when the optimist dominates in the market, when pessimist's consumption and wealth share gradually increases, she shorts more in proportion of her wealth. As a result, the optimist should hold more stocks in equilibrium which requires higher market prices of risk, and hence, better investment opportunities. Therefore, the income effect decreases the optimist's wealth-consumption ratio. However, as the pessimist's

consumption share increases further, at some point the price-dividend ratio should start increasing again since when the pessimist dominates, the optimist's wealth is low and she becomes unable to hold large amount of stock. As a result, shorting becomes less attractive for the pessimist in equilibrium and her subjective market price of risk increases pushing up the wealth-consumption ratio and hence the price-divided ratio.

Panel (b) of Figure 1.4 demonstrates that stock return volatility can both be higher and lower than the volatility of dividends, which is due to the U-shaped form of the price-dividend ratio. Moreover, as short-sale constraints become tighter the volatility of stock returns decreases for small consumption shares y , increases for medium y , and is almost unchanged for values of y close to unity when the constraint does not bind. Intuitively, short-sale constraints limit the ability of the pessimist to trade on her pessimism and hence her stockholding look as if she had smaller disagreement with the unconstrained investor. As a result, the economic parameters should become closer to the values in the unconstrained economy without disagreement. In particular, stock return volatilities should move closer to the volatility of dividends σ_δ , which we observe on Figure 1.4. This effect can also be formally demonstrated by observing that adjustment parameters for unconstrained and constrained investors are such that $\nu_o^* = 0$ (case (a) in Table 1.1) and $\nu_p^* \geq 0$ (case (e) in Table 1.1), and hence the volatility σ_λ given by (70) decreases towards zero since the volatility of stock returns σ is positive. Then, from the expression for volatility σ in (81) it follows that the difference between σ and dividend growth volatility σ_δ becomes smaller.

In a similar model with a logarithmic constrained pessimist Gallmeyer and Hollifield (2008) find that the stock return volatility increases when the unconstrained optimist has risk aversion $\gamma > 1$ and each investor is initially endowed with 50% of the market portfolio. By contrast with their work we present the analysis of price-dividend ratios and stock return volatilities as functions of both the pessimist's consumption share y and the tightness of the short-sale constraint. Moreover, we show that the volatility σ can decrease with tighter constraints even though the economic magnitude of this effect is small. Finally, our numerical method relies only on solving linear algebraic equations at each step rather than

employing Monte-Carlo simulations as in their work.

1.4.2. Multiple Stock Formulation

We now demonstrate that the baseline analysis of Section 1.2 with single stock can easily be generalized to the case of multiple stocks. The uncertainty is now generated by a multi-dimensional Brownian motion $w = (w_1, \dots, w_N)$. The investors trade in a riskless bond and N stocks in a positive net supply, normalized to unity, each of which is a claim to an exogenous strictly positive stream of dividends δ_n following the dynamics

$$d\delta_{nt} = \delta_{nt}[\mu_{\delta_{nt}}dt + \sigma_{\delta_{nt}}^\top dw_t], \quad n = 1, \dots, N, \quad (83)$$

where μ_{δ_n} and σ_{δ_n} are stochastic processes. We consider equilibria in which bond prices, B , and stock prices, S , follow processes

$$dB_t = B_t r_t dt \quad (84)$$

$$dS_{nt} + \delta_{nt} dt = S_{nt}[\mu_{nt} dt + \sigma_{nt}^\top dw_t], \quad n = 1, \dots, N. \quad (85)$$

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 measuring the covariance between the stock return and Brownian motion w_n . By δ we denote the process for aggregate dividend, $\delta = \delta_1 + \delta_2 + \dots + \delta_N$, which follows the process

$$d\delta_t = \delta_t[\mu_{\delta t} + \sigma_{\delta t}^\top dw_t], \quad (86)$$

where

$$\mu_{\delta t} = \frac{\delta_{1t}}{\delta_t} \mu_{\delta_{1t}} + \dots + \frac{\delta_{Nt}}{\delta_t} \mu_{\delta_{Nt}}, \quad \sigma_{\delta t} = \frac{\delta_{1t}}{\delta_t} \sigma_{\delta_{1t}} + \dots + \frac{\delta_{Nt}}{\delta_t} \sigma_{\delta_{Nt}}.$$

Investor 1 is endowed with s_n units of stock n and $-b$ units of bond, while investor 2 is endowed with $1 - s_n$ units of stock n and b units of bond. Investor i 's wealth process W follows

$$dW_{it} = \left[W_{it} \left(r_t + \theta_{it}^\top (\mu_t - r_t) \right) - c_{it} \right] dt + W_{it} \theta_{it}^\top \sigma_t dw_t, \quad (87)$$

and her investment policies are subject to portfolio constraints

$$\theta_i \in \Theta_i, \quad i = 1, 2, \quad (88)$$

where Θ_i is a closed convex set in \mathbb{R}^N and $\theta = (\theta_1, \dots, \theta_N)^\top$ is the vector of wealth proportions invested in the N stocks. Each investor i solves her dynamic optimization (6) subject to budget constraint (87), no-bankruptcy constraint $W_t \geq 0$ and portfolio constraints (88).

Following the approach of Section 1.2 we first embed the optimization problem for each investor into an equivalent fictitious complete-market economy in which stock prices evolve as

$$d\xi_{it} = -\xi_{it}[r_{it}dt + \kappa_{it}^\top dw_t]. \quad (89)$$

Assuming that dual problems in Cvitanic and Karatzas (1992) have solutions we obtain that riskless rates r_{it} and market prices of risk κ_{it} in fictitious economy are given by

$$r_{it} = r_t + f_i(\nu_{it}^*), \quad \kappa_{it} = \kappa_t + \sigma_t^{-1}\nu_{it}^*, \quad (90)$$

where κ is the market price of risk in the original economy, $f_i(\nu)$ are support functions for the sets Θ_i , defined as

$$f_i(\nu) = \sup_{\theta \in \Theta_i} (-\nu^\top \theta), \quad (91)$$

ν_{1t}^* and ν_{2t}^* solve duality optimization problem in Cvitanic and Karatzas (1992) and belong to the effective domains for support functions, given by

$$\Upsilon_i = \{\nu \in \mathbb{R}^N : f_i(\nu) < \infty\}. \quad (92)$$

Proposition 1.5 characterizes the equilibrium in terms of the adjustments ν_{it}^* and $f(\nu_{it}^*)$ in fictitious economies and the parameters of the process for the ratio of marginal utilities of consumption, λ_t , which evolves as

$$d\lambda_t = -\lambda_t[\mu_{\lambda t}dt + \sigma_{\lambda t}^\top dw_t]. \quad (93)$$

Proposition 1.5. *If there exists an equilibrium, the riskless interest rate r , market price of risk κ , drift μ_λ and volatility σ_λ of weighting process λ that follows (93) are given by*

$$r_t = \bar{r}_t - \frac{A_t}{A_{1t}} f_1(\nu_{1t}^*) - \frac{A_t}{A_{2t}} f_2(\nu_{2t}^*) - \frac{A_t^3 (P_{1t} + P_{2t})}{2A_{1t}^2 A_{2t}^2} \sigma_{\lambda t}^\top \sigma_{\lambda t} - \frac{A_t^3}{A_{1t} A_{2t}} \left(\frac{P_{1t}}{A_{1t}} - \frac{P_{2t}}{A_{2t}} \right) \delta_t \sigma_{\delta t}^\top \sigma_{\lambda t}, \quad (94)$$

$$\kappa_t = \bar{\kappa}_t - \frac{A_t}{A_{1t}} \sigma_t^{-1} \nu_{1t}^* - \frac{A_t}{A_{2t}} \sigma_t^{-1} \nu_{2t}^*, \quad (95)$$

$$\mu_{\lambda t} = A_t \delta_t \sigma_{\delta t}^\top \sigma_{\lambda t} + f_1(\nu_{1t}^*) - f_2(\nu_{2t}^*) - \frac{A_t}{A_{1t}} \sigma_{\lambda t}^\top \sigma_{\lambda t}, \quad \sigma_{\lambda t} = \sigma_t^{-1} (\nu_{1t}^* - \nu_{2t}^*), \quad (96)$$

where \bar{r} is the riskless rate and $\bar{\kappa}$ is the market price of risk in an unconstrained economy, given by

$$\bar{r}_t = \rho + A_t \delta_t \mu_{\delta t} - \frac{A_t P_t}{2} \delta_t^2 \sigma_{\delta t}^\top \sigma_{\delta t}, \quad \bar{\kappa}_t = A_t \delta_t \sigma_{\delta t}, \quad (97)$$

A_{it} , P_{it} , and A_t and P_t are absolute risk aversions and prudence parameters of investor i and a representative investor with utility (19), respectively. Optimal consumptions c_i^* , wealths W_i and optimal investment policies θ_i^* are given by

$$c_{it}^* = g_i(\delta_t, \lambda_t), \quad (98)$$

$$W_{it}^* = \frac{1}{\xi_{it}} E_t \left[\int_0^\infty \xi_{is} c_{is}^* ds \right], \quad (99)$$

$$\theta_{it}^* = \sigma_t^{-1} \left(W_{it}^* (\kappa_t + \sigma_t^{-1} \nu_{it}^*) + \frac{\phi_{it}}{\xi_{it}} \right), \quad (100)$$

where functions $g_i(\delta_t, \lambda_t)$ are such that c_{1t}^* and c_{2t}^* satisfy consumption clearing in (9) and equation (18) for process λ , state prices ξ_{it} follow processes (14) and ϕ_i are such that

$$M_{it} \equiv E_t \left[\int_0^\infty \xi_{is} c_{is}^* ds \right] = M_{i0} + \int_0^t \phi_{is}^\top dw_s.$$

Initial value λ_0 is such that budget constraints at time 0 are satisfied:

$$s_{i1} S_{10} + \dots + s_{iN} S_{N0} + b_i = W_{i0}^*, \quad (101)$$

where $s_{1n} = s_n$, $s_{2n} = 1 - s_n$, $b_1 = -b$ and $b_2 = b$. Moreover, adjustments ν_{it}^* satisfy complementary slackness condition

$$f_i(\nu_{it}^*) + \theta_{it}^{*\top} \nu_{it}^* = 0. \quad (102)$$

The expression for interest rates (94) can again be decomposed into three groups of terms that represent riskless rate in an unconstrained economy, the impact of constraints and the effect of risk sharing. The last term in (94) also shows that in the case of heterogeneous

utility functions the interest rates depend on the covariance between aggregate dividend and weighting process λ , captured by $\sigma_\delta^\top \sigma_\lambda$. The expression for equilibrium interest rates also allows to formulate a simple sufficient condition under which the equilibrium interest rates in the constrained economy are lower than in the unconstrained one.

Corollary 1.3. *If the utility functions and the allocations of consumption are such that $P_1/A_1 = P_2/A_2$ and the sets of portfolio constraints Θ_i contain the origin, i.e. $0 \in \Theta_i$, then the interest rate in a constrained economy, r , is lower than in an unconstrained one, \bar{r} , and the following upper bound for rate r holds:*

$$r_t \leq \bar{r}_t - \frac{A_t^3(P_{1t} + P_{2t})}{2A_{1t}^2 A_{2t}^2} \sigma_\lambda^\top \sigma_\lambda. \quad (103)$$

The expressions for the market price of risk now reflect the impact of multiple constraints. By contrast with the single stock case, market clearing conditions can only determine the aggregate value of all stocks and not the values of individual ones. Moreover, as demonstrated in Hugonnier (2008) if the weighting process is not a martingale then there might be multiple equilibria with different stock prices but unique riskless rates and market prices of risk. The application of the methodology developed in Section 1.3 for finding equilibria in a multi-stock economy we leave for the future research.

1.5. Conclusion

Despite numerous applications of dynamic equilibrium models with heterogeneous investors facing portfolio constraints, little is known about the equilibrium when we depart from the assumption of logarithmic preferences. In various frameworks we provide explicit expressions for interest rates and market prices of risk in terms of instantaneous volatilities of stock returns and consumptions as well as risk aversions and prudence parameters. We then consider an economic setting where one investor is unconstrained while the other faces upper bound on the proportion that can be invest in stocks, and both investors have identical CRRA utilities. We completely characterize the equilibrium in terms of investors' wealth-

consumption ratios satisfying a pair of differential equations that we solve numerically by employing a simple iterative algorithm. We further demonstrate that the direction in which portfolio constraints change price-dividend ratios and stock returns volatilities crucially depends on the intertemporal elasticity of substitution (IES). In particular, when the IES is greater than unity the model generates countercyclical market prices of risk and stock return volatilities, procyclical price-dividend ratios, excess volatility and other patterns consistent with empirical findings. We also find that the impact of constrained investor diminishes in the course of time but is still significant even after one hundred years. Our approach is then extended to the case of heterogeneous beliefs and multiple assets. Given the tractability of our analysis we believe that our approach for finding equilibria in economies with constraints may find applications in various models with heterogeneous investors and incomplete financial markets as well as in solving portfolio choice problems with constraints at a partial equilibrium level.

1.6. Appendix A: Proofs

Proof of Proposition 1.1. First, we obtain a system of equations for parameters of the fictitious economy by substituting the expressions for optimal consumption (17) into consumption clearing condition in (9), applying Itô's Lemma to both sides and matching the coefficients. Noting from the properties of inverse functions that

$$I'_i(\psi_i e^{\rho t} \xi_{it}) = \frac{1}{u'_i(c_{it}^*)}, \quad I''_i(\psi_i e^{\rho t} \xi_{it}) = -\frac{u''_i(c_{it}^*)}{u'_i(c_{it}^*)} \frac{1}{(u'_i(c_{it}^*))^2},$$

we obtain the following equations

$$\frac{r_t - \rho}{A_t} + \frac{f_1(\nu_{1t}^*)}{A_{1t}} + \frac{f_2(\nu_{2t}^*)}{A_{2t}} + \frac{1}{2} \left(P_{1t} \left(\frac{\kappa_{1t}}{A_{1t}} \right)^2 + P_{2t} \left(\frac{\kappa_{2t}}{A_{2t}} \right)^2 \right) = \delta_t \mu_{\delta t}, \quad (104)$$

$$\frac{\kappa_{1t}}{A_{1t}} + \frac{\kappa_{2t}}{A_{2t}} = \delta_t \sigma_{\delta t}. \quad (105)$$

By applying Itô's Lemma to both sides of the definition of λ in (18) and noting that marginal utilities $u'_i(c_i^*)$ are given by (16), matching the terms we obtain the drift μ_λ and volatility σ_λ of the weighting process (20):

$$\mu_{\lambda t} = \sigma_{\lambda t} \kappa_{2t} + f_1(\nu_{1t}^*) - f_2(\nu_{2t}^*), \quad \sigma_{\lambda t} = \kappa_{1t} - \kappa_{2t}. \quad (106)$$

Taking into account the definition of κ_{it} in terms of adjustments in (15) from equations (104)–(106) we obtain expressions (21)–(23) in Proposition 1.1. Analogously, it can be shown that in the unconstrained economy the interest rate is given by (24).

Optimal consumptions c_{it}^* are obtained from consumption clearing and the equation for weight λ in (18). Expressions for optimal wealths and optimal policy (26) and (28) follow from the results in Cox and Huang (1989), Huang and Pages (1992) and Karatzas and Shreve (1998), while stock prices (3) are derived from the market clearing conditions in (9). The complementary slackness condition in (30) is established in Chapter 6.3 of Karatzas and Shreve (1998).

Q.E.D.

Proof of Corollary 1.1. The proof directly follows from Proposition 1.1 by noting that the last term in the expression for r in (21) disappears.

Q.E.D.

Proof of Proposition 1.2. We obtain expressions (44)–(48) for equilibrium parameters from expressions (21)–(25) in Proposition 1.1 by substituting adjustment parameters (39) and risk-aversion and prudence parameters for CRRA preferences, given by

$$\begin{aligned} A_{1t} &= \frac{\gamma}{c_{1t}}, & A_{2t} &= \frac{\gamma}{c_{2t}}, & A_t &= \frac{\gamma}{\delta_t}, \\ P_{1t} &= \frac{1+\gamma}{c_{1t}}, & P_{2t} &= \frac{1+\gamma}{c_{2t}}, & P_t &= \frac{1+\gamma}{\delta_t}. \end{aligned} \quad (107)$$

We first demonstrate that the constraint for investor 2 should always be binding in equilibrium. The complementary slackness condition (30), given expressions for adjustments (39), takes the form $(\bar{\theta} - \theta_{2t}^*)\nu_{2t}^* = 0$. Therefore, if constraint does not bind it follows that $\nu_{2t}^* = 0$. Hence, from (23) we obtain that $\sigma_{\lambda t} = 0$ and $\mu_{\lambda t} = 0$ and the economy will permanently remain in a Pareto-efficient unconstrained equilibrium. As a result, since the investors have identical preferences and the equilibrium investment opportunity sets are constant when $\sigma_{\lambda t} = 0$ and $\mu_{\lambda t} = 0$, it can easily be verified that the investors will choose $\theta_{it}^* = 1$, which violates constraint $\theta_{2t} \leq \bar{\theta} < 1$. Therefore, the constraint should always be binding in equilibrium.

Expressions for wealths W_{it}^* follow from the first order condition for consumption in (41), while the expression for price-dividend ratio R follows from the expression for stock price (3), derived from consumption clearing, and the expressions for wealths in (49). Optimal policy for investor 1, θ_{1t}^* , in (51) is obtained by solving an HJB equation, while policy for investor 2 equals $\bar{\theta}$ since the investor always binds on her constraint, as demonstrated below. Stock return volatility σ in (52) is derived by applying Itô's Lemma to stock price, given by $S_t = R_t \delta_t$.

From the definition of κ_{it} in (15), expression for κ_t in (45) and expressions for adjustments in (39) we find that

$$\kappa_{2t} = \gamma \sigma_\delta - \frac{1}{1 + \lambda_t^{1/\gamma}} \sigma_\lambda. \quad (108)$$

Substituting κ_{2t} from (108) into expression for optimal investment policy (43) and noting

that constraint $\theta_{2t} \leq \bar{\theta}$ is always binding we obtain the following equation for σ_λ :

$$\frac{1}{\gamma\sigma_t} \left(\gamma\sigma_\delta - \sigma_{\lambda t} \left(\frac{1}{1 + \lambda_t^{1/\gamma}} + \gamma \frac{\partial H_{2t}}{\partial \lambda_t} \frac{\lambda_t}{H_{2t}} \right) \right) = \bar{\theta}. \quad (109)$$

Substituting volatility σ given by first expression in (52) into equation (109) and solving it yields σ_λ given by second expression in (52). Finally, the equation for λ_0 is obtained so as to satisfy time-0 budget constraints (29). By substituting W_{10}^* , W_{20}^* and $S_0 = R_0\delta_0$ from Proposition 1.2 into the budget constraints (29) it can easily be observed that both constraints are satisfied whenever equation (53) for λ_0 holds.¹⁴

Q.E.D.

Proof of Corollary 1.2. Applying Itô's Lemma to both sides of the first order conditions for consumption (16) and matching the terms we find that

$$c_{it}\sigma_{c_{it}} = \frac{\kappa_{it}}{A_{it}}. \quad (110)$$

Since investor 1 is unconstrained, $\kappa_1 = \kappa$ and is given by (45) while κ_2 is given by (108). Substituting κ_1 and κ_2 into (110) and noting that for CRRA utility $A_i = \gamma/c_i$ we obtain expressions (57) for volatilities σ_{c_i} .

Q.E.D.

¹⁴We also note that the results of Proposition 1.2 can be derived without relying on the methodology in Cvitanic and Karatzas (1992) by solving the HJB for investor 2 directly in constrained economy. Since the constraint is always binding the problem is equivalent to the one with constraint $\theta_{2t} = \bar{\theta}$. The HJB equation is then given by (38) in which $\theta_{2t} = \bar{\theta}$ and $\nu_{it}^* = 0$, since we solve in constrained economy. Then, conjecturing that J_{2t} has form (40) yields the equation for H_{2t} . From the first order condition (41) we obtain $e^{-\rho t} W_{2t}^{-\gamma} H_{2t}^\gamma = \xi_{2t}$, where ξ_{2t} is the marginal utility of investor 2 which follows the process (14). Applying Itô's Lemma to both sides shows that

$$\bar{\theta}\sigma_t = \frac{\kappa_{2t}}{\gamma} - \sigma_{\lambda t} \frac{\partial H_{2t}}{\partial \lambda_t} \frac{\lambda_t}{H_{2t}}.$$

Substituting this expression into HJB after some algebra we obtain equation (42) for investor 2. Price of risk κ_2 can be found from (105)–(106) while r_2 can be found by applying Itô's Lemma to $\xi_{2t}W_{2t} = e^{-\rho t} W_{2t}^{1-\gamma} H_{2t}^\gamma$, noting that the right-hand side satisfies HJB equation, $\theta_{2t} = \bar{\theta}$, and matching the terms.

Moreover, since investor 1 faces complete market, in the derivation of r_t and κ_t to obtain equations (104)–(105) we assume that $u'(c_{1t}^*) = \psi_1 e^{\rho t} \xi_{1t}$ where

$$\xi_{1t} = -\xi_{1t}[r_t dt + \kappa_t dw_t].$$

Huang and Pages (1992) derive this result assuming that $\int_0^t |r_\tau| d\tau < \infty$ a.s., and $\kappa_t < \bar{K}$ a.s., where \bar{K} is a constant. It is difficult to check these conditions analytically. However, the graphs on Figure 1.3 demonstrate that the states with y close to 1, where r_t and κ_t are unbounded, have zero probability, and hence, the conditions are likely to be satisfied. We also check numerically that the integrals in investor's optimization (6) converge to J_{it} derived in Section 1.3.

Proof of Proposition 1.3. From expression (61) we first express Brownian motion w^p in terms of Brownian motion w^o as follows:

$$dw_t^p = \Delta\mu_{\delta t}dt + dw_t^o, \quad (111)$$

and then rewrite all subsequent stochastic processes in terms of Brownian motion w^o under the optimist's probability measure. Then, state prices ξ_{it} in fictitious economies follow processes:

$$d\xi_{ot} = -\xi_{ot}[r_{ot}dt + \kappa_t^o dw_t^o], \quad d\xi_{pt} = -\xi_{pt}[(r_{pt} + \Delta\mu_{\delta t}\kappa_t^p)dt + \kappa_t^p dw_t^o]. \quad (112)$$

Optimal consumptions in fictitious economies are given by (17). Substituting them into consumption clearing condition in (9), applying Itô's Lemma to both sides and matching terms as in the proof of Proposition 1.1 after some algebra we obtain:

$$\frac{r_t - \rho}{A_t} + \frac{f_o(\nu_{ot}^*)}{A_{ot}} + \frac{f_p(\nu_{pt}^*)}{A_{pt}} + \frac{1}{2} \left(P_{ot} \left(\frac{\kappa_t^o}{A_{ot}} \right)^2 + P_{pt} \left(\frac{\kappa_t^p}{A_{pt}} \right)^2 \right) = \frac{\kappa_t^o}{A_{ot}} \frac{\mu_{\delta t}^o}{\sigma_{\delta t}} + \frac{\kappa_t^p}{A_{pt}} \frac{\mu_{\delta t}^p}{\sigma_{\delta t}}, \quad (113)$$

$$\frac{\kappa_t^o}{A_{ot}} + \frac{\kappa_t^p}{A_{pt}} = \delta_t \sigma_{\delta t}. \quad (114)$$

By applying Itô's Lemma to both sides of the definition of λ in (18) and noting that marginal utilities $u'_i(c_i^*)$ are given by (16) and state prices follow (112), matching the terms we obtain the drift μ_λ and volatility σ_λ of the weighting process (66) for the optimist:

$$\mu_{\lambda t}^o = \sigma_{\lambda t} \kappa_t^p - \Delta\mu_{\delta t} \kappa_t^p + f_o(\nu_{ot}^*) - f_p(\nu_{pt}^*), \quad \sigma_{\lambda t} = \kappa_t^o - \kappa_t^p. \quad (115)$$

Using equations (113), (114) and the second equation in (115) we obtain expressions for r and κ in Proposition 1.3.

To obtain drift $\mu_{\lambda t}^p$ we rewrite the process for λ_t given by (66) under the Brownian motion of the optimist as follows

$$d\lambda_t = -\lambda_t [(\mu_{\lambda t}^p + \sigma_{\lambda t} \Delta\mu_{\delta t})dt + \sigma_{\lambda t} dw_t^o].$$

Matching the drift parameters for the processes for λ_t from both optimist's and pessimist's points of view yields expression for $\mu_{\lambda t}^p$ in Proposition 1.3. To obtain expression for σ_λ we

note first that by the definition of prices of risk in fictitious economies

$$\kappa_t^o = \frac{\mu_t^o - r_t}{\sigma_t} + \frac{\nu_{it}^o}{\sigma_t}, \quad \kappa_t^p = \frac{\mu_t^p - r_t}{\sigma_t} + \frac{\nu_{it}^p}{\sigma_t}. \quad (116)$$

Moreover, rewriting the process for stock prices S_t for both the optimist and pessimist in terms of Brownian motion w^o

$$\begin{aligned} dS_t &= S_t[\mu_t^o dt + \sigma_t dw_t^o] \\ &= S_t[(\mu_t^p + \Delta\mu_{\delta t}\sigma_t)dt + \sigma_t dw_t^o], \end{aligned}$$

and matching the terms we obtain

$$\frac{\mu_t^o - \mu_t^p}{\sigma_t} = \Delta\mu_{\delta t}. \quad (117)$$

The expression for σ_λ in (115) along with equations (116) and (117) gives σ_λ reported in Proposition 1.3. The rest of the proof is as in Proposition 1.1. *Q.E.D.*

Proof of Proposition 1.4. From the definition of the support function in (12) applied to $\theta \geq \underline{\theta}$ and the expression (70) for volatility σ_λ we obtain the adjustment parameters:

$$\nu_{1t}^* = 0, \quad f(\nu_{1t}^*) = 0, \quad \nu_{2t}^* = \sigma_t(\Delta\mu_{\delta t} - \sigma_{\lambda t}), \quad f(\nu_{2t}^*) = -\underline{\theta}\sigma_t(\Delta\mu_{\delta t} - \sigma_{\lambda t}). \quad (118)$$

Substituting adjustments (118) and risk-aversion and prudence parameters in (107), into the expressions (67)–(71) we obtain equilibrium parameters (73)–(76) reported in Proposition 1.4.

Consumptions (77) are obtained from the consumption clearing condition in (9) and definition of λ_t in (18). Wealth-consumption ratios H^o and H^p satisfy HJB equations (42) in which μ_λ is replaced by μ_λ^o and μ_λ^p respectively. Hence, from the first order condition for consumption in (41) and market clearing condition we obtain expressions for W_{it}^* and R_t . Expressions for optimal policies are obtained by solving HJB equations in fictitious economies, as in Section 1.3, while stock return volatility σ is obtained by applying Itô's Lemma to stock price $S_t = R_t \delta_t$.

The complementary slackness condition in (30) in our setting takes the form $(\underline{\theta} - \theta_{it}^*)\nu_{it}^* = 0$. As a result, if constraint is not binding $\nu_{it}^* = 0$, and hence, from the expression in (70) it

follows that $\sigma_\lambda = \Delta\mu_{\delta t}$. To solve for σ_λ when the constraint is binding we first substitute κ^p from (74) into the investment policy (80) and obtain

$$\theta_{pt}^* = \frac{1}{\gamma\sigma_t} \left(\gamma\sigma_\delta - \sigma_{\lambda t} \left(\frac{1}{1 + \lambda_t^{1/\gamma}} + \gamma \frac{\partial H_t^p}{\partial \lambda_t} \frac{\lambda_t}{H_t^p} \right) \right). \quad (119)$$

Then, substituting σ from (81) into (119) and solving equation $\theta_{pt}^* = \underline{\theta}$ we obtain

$$\sigma_{\lambda t} = \frac{(1 - \underline{\theta})\gamma\sigma_\delta}{\frac{1}{1 + \lambda_t^{1/\gamma}} + \gamma \frac{\partial H_{pt}}{\partial \lambda_t} \frac{\lambda_t}{H_{pt}} - \underline{\theta} \gamma \frac{\partial R_t}{\partial \lambda_t} \frac{\lambda_t}{R_t}}. \quad (120)$$

Moreover, since $\nu_{2t}^* \geq 0$ (Table 1.1 case (e)) if the constraint binds σ_λ is given by (120) and should be lower than $\Delta\mu_{\delta t}$ which leads to expression for σ_λ in Proposition 1.4.¹⁵ *Q.E.D.*

Proof of Proposition 1.5. The proof is a multi-dimensional version of the proof of Proposition 1.1. *Q.E.D.*

Proof of Corollary 1.3. From the definition of support functions in (12) it follows easily that $f_i(\nu) \geq 0$ if $0 \in \Theta_i$. Then, the proof follows from the fact that in the expression for interest rates r in Proposition 1.5 the second and third terms are positive while the last term vanishes. *Q.E.D.*

¹⁵Similarly to the discussion in the footnote in the proof of Proposition 1.2 it can be argued that the results in Proposition 1.4 can be obtained without relying on the methodology in Cvitanic and Karatzas (1992).

1.7. Appendix B: Numerical Method

We here present the details of our numerical solution method in Section 1.3 first for $\gamma > 1$ and then for $\gamma < 1$. Since variable λ takes values in the interval $(0, +\infty)$ we first rewrite the HJB equations (42) in terms of variable $y = \lambda^{1/\gamma}/(1 + \lambda^{1/\gamma})$. By $\tilde{H}_i(y, t)$ we denote the wealth-consumption ratios as functions of y so that

$$H_i(\lambda, t) = \tilde{H}_i(y(\lambda), t). \quad (121)$$

The derivatives of $H_i(\lambda, t)$ then can be expressed in terms of derivatives of $\tilde{H}_i(y, t)$ by differentiating both sides in (121) as follows:

$$\frac{\partial H_{it}}{\partial t} = \frac{\partial \tilde{H}_{it}}{\partial t}, \quad \lambda_t \frac{\partial H_{it}}{\partial \lambda_t} = \frac{y(1-y)}{\gamma} \frac{\partial \tilde{H}_{it}}{\partial y}, \quad (122)$$

$$\lambda_t^2 \frac{\partial^2 H_{it}}{\partial \lambda_t^2} = \frac{y^2(1-y)^2}{\gamma^2} \frac{\partial^2 \tilde{H}_{it}}{\partial y^2} + \frac{2y(1-y)((1-\gamma)/2 - y)}{\gamma^2} \frac{\partial \tilde{H}_{it}}{\partial y}. \quad (123)$$

Taking into account our change of variable and the expressions for derivatives in (122)–(123) from the expressions in Proposition 1.2, definitions of parameters r_{it} and κ_{it} in (15), and expressions for adjustment parameters in (39) we obtain the following expressions for equilibrium parameters in fictitious economies:

$$\begin{aligned} r_{1t} &= \bar{r} - \frac{y}{1-y} \bar{\theta} \sigma_t \sigma_{yt} - \frac{1+\gamma}{2\gamma} \frac{y}{1-y} \sigma_{yt}^2, & \kappa_{1t} &= \gamma \sigma_\delta + \frac{y}{1-y} \sigma_{yt}, \\ r_{2t} &= \bar{r} + \bar{\theta} \sigma_t \sigma_{yt} - \frac{1+\gamma}{2\gamma} \frac{y}{1-y} \sigma_{yt}^2, & \kappa_{2t} &= \gamma \sigma_\delta - \sigma_{yt}, \\ \mu_{\lambda t} &= \frac{\mu_{yt}}{1-y}, & \sigma_{\lambda t} &= \frac{\sigma_{yt}}{1-y} \end{aligned} \quad (124)$$

where \bar{r} is given by (47), μ_{yt} , σ_t and σ_{yt} are given by

$$\mu_{yt} = \gamma \sigma_\delta \sigma_{yt} - \bar{\theta} \sigma_t \sigma_{yt} - \sigma_{yt}^2, \quad \sigma_t = \sigma_\delta - \frac{\sigma_{yt}}{\gamma} \frac{\partial \tilde{R}_t}{\partial y_t} \frac{y_t}{\tilde{R}_t}, \quad \sigma_{yt} = \frac{(1-\bar{\theta})\gamma \sigma_\delta}{1 + \frac{\partial \tilde{H}_{2t}}{\partial y_t} \frac{y_t}{\tilde{H}_{2t}} - \bar{\theta} \frac{\partial \tilde{R}_t}{\partial y_t} \frac{y_t}{\tilde{R}_t}}, \quad (125)$$

and \tilde{R}_t is a price-dividend ratio as a function of y . Substituting expressions for derivatives (122) and (123) into the HJB equations (42) we obtain the following PDEs for \tilde{H}_{it} :

$$\begin{aligned} &\frac{\partial \tilde{H}_{it}}{\partial t} + \frac{y_t^2 \sigma_{yt}^2}{2\gamma^2} \frac{\partial^2 \tilde{H}_{it}}{\partial y_t^2} + \frac{y_t}{\gamma^2} \left(\sigma_{yt}^2 \frac{(1-\gamma)/2 - y_t}{1-y_t} - \gamma \mu_{yt} - (1-\gamma) \kappa_{it} \sigma_{yt} \right) \frac{\partial \tilde{H}_{it}}{\partial y_t} \\ &+ \frac{1}{\gamma} \left(\frac{1-\gamma}{2\gamma} \kappa_{it}^2 + (1-\gamma) r_{it} - \rho \right) \tilde{H}_{it} + 1 = 0, \quad i = 1, 2. \end{aligned} \quad (126)$$

To find stationary, time-independent solutions of equations (126) we fix a large horizon T , pick two functions $\tilde{h}_1(y)$ and $\tilde{h}_2(y)$, specify terminal condition

$$\tilde{H}_i(y, T) = \tilde{h}_i(y), \quad i = 1, 2, \quad (127)$$

and solve HJB equations (126) backwards until the convergence to stationary solutions. We assume that functions \tilde{h}_i are continuous and differentiable on the interval $[0, 1]$ and satisfy conditions $\tilde{h}_1(1) = 0$ and $\tilde{h}'_2(1) = (\gamma - 1)\tilde{h}_2(1)$.

We assume that $\tilde{H}_i(y, t)$ are twice continuously differentiable in the interval $(0, 1)$, have bounded first and second right derivatives at $y = 0$, $\sigma_y^2 > 0$, and there exist limits $(1 - y)^2 \partial^2 \tilde{H}_1(y, t) / \partial y^2 \rightarrow 0$, $(1 - y) \partial^2 \tilde{H}_2(y, t) / \partial y^2 \rightarrow 0$ and $(1 - y) \partial \tilde{H}_1(y, t) / \partial y \rightarrow 0$, as $y \rightarrow 1$. After we compute the solutions we also verify numerically that these assumptions are satisfied for $\gamma > 1$.

Passing to the limit $y \rightarrow 0$ in equations (126) we obtain simple ordinary differential equations for $H_i(0, t)$ solving which yields boundary conditions at $y = 0$:

$$\tilde{H}_i(0, t) = \tilde{h}_i(0) e^{p_i(T-t)} + \frac{e^{p_i(T-t)} - 1}{p_i}, \quad i = 1, 2, \quad (128)$$

where

$$p_1 = \frac{1 - \gamma}{2} \bar{\theta}^2 \sigma_\delta^2 + \frac{(1 - \gamma)\bar{r} - \rho}{\gamma}, \quad p_2 = \frac{1 - \gamma}{2} \sigma_\delta^2 + \frac{(1 - \gamma)\bar{r} - \rho}{\gamma}. \quad (129)$$

Expressions in (128) and (129) demonstrate that conditions $p_i \leq 0$ are necessary for the existence of stationary solutions of equations (126). To obtain boundary conditions at $y = 1$ we multiply the equations for $H_1(y, t)$ and $H_2(y, t)$ by $(1 - y)^2$ and $(1 - y)$, respectively, and passing to the limit $y \rightarrow 1$ we obtain:

$$(1 - \bar{\theta})(\gamma - 1)\tilde{H}_1(1, t) = 0, \quad \frac{\partial \tilde{H}_2(1, t)}{\partial y} = (\gamma - 1)\tilde{H}_2(1, t). \quad (130)$$

The problem then becomes to solve HJB equations (126) subject to terminal condition (127) and boundary conditions (128) and (130).

For simplicity, in the description of the numerical method we omit subscript i . We let the time and state variable increments denote $\Delta t \equiv T/M$ and $\Delta y \equiv 1/N$, where M

and N are integer numbers, and index time and state variables by $t = 0, \Delta t, 2\Delta t, \dots, T$ and $y = 0, \Delta y, 2\Delta y, \dots, 1$, respectively. Next, we derive discrete-time analogues of HJB equations and boundary conditions replacing derivatives by their finite-difference analogues as follows:

$$\frac{\tilde{H}_{n,k+1} - \tilde{H}_{n,k}}{\Delta t} + a_{n,k+1} \frac{\tilde{H}_{n+1,k} - 2\tilde{H}_{n,k} + \tilde{H}_{n-1,k}}{\Delta y^2} + b_{n,k+1} \frac{\tilde{H}_{n,k} - \tilde{H}_{n-1,k}}{\Delta y} + c_{n,k+1} \tilde{H}_{n,k} + 1 = 0, \quad (131)$$

$$\tilde{H}_{n,M} = \tilde{h}_n, \quad \tilde{H}_{0,k} = d_{0,k}, \quad \tilde{H}_{N,k} = e_{N,k} \tilde{H}_{N-1,k}, \quad (132)$$

where $n = 1, 2, \dots, N - 1$, $k = 1, 2, \dots, M - 1$, $\tilde{H}_{n,k} = \tilde{H}(n\Delta y, k\Delta t)$. The coefficients in (131) correspond to coefficients in equation (126) and are computed using the solution $\tilde{H}_{n,k+1}$, while coefficients in (131) are obtained by replacing terminal condition (127) and boundary conditions (128) and (130) by their finite-difference analogues. The system of equations in (131)–(132) is then solved backwards in time, starting at $k = M - 1$. Given solution $H_{n,k+1}$ we compute all the coefficients in (131) at step $k + 1$, and hence at step k function $H_{n,k}$ for fixed k solves a system of linear algebraic equations. We then iterate backwards until the process converges to a stationary time-independent solution.

Figure 1.1 shows the numerical solutions for wealth-consumption ratios plotted against constrained investors share of consumption, y , for plausible exogenous parameters. These numerical solutions have the appearance of bounded and twice continuously differentiable on interval $[0, 1]$ functions irrespective of the grid parameter Δy . Assuming that they are indeed twice continuously differentiable, and given that they satisfy finite-difference equations (131)–(132), passing to a limit $\Delta y \rightarrow 0$ indeed gives solutions to the HJB equations for wealth-consumption ratios.¹⁶

When risk aversion γ is less than unity wealth-consumption ratio H_{1t} and its derivatives become unbounded while σ_{yt} approaches zero, as y approaches unity. As a result, the assumptions under which the boundary conditions (128) and (130) are derived are violated. However, it turns out that function $(1 - y)H_{1t}$ is bounded and equals zero when $y = 1$.

¹⁶As an additional check we also verify by Monte-Carlo simulations that for both investors integrals in their optimization problem (6) do not explode under optimal consumption policies in (48) and converge to the values obtained by our numerical method. The convergence of those integrals also implies that the transversality conditions for HJB equations (38) are satisfied.

Hence, we derive the differential equation for $(1 - y)H_{1t}$ and solve it using the methodology described above.

The model with heterogeneous beliefs in Section 1.4.1 is solved in a similar way. First, we derive an HJB equation in terms of consumption share y , which is given by (126) in which μ_y is replaced by μ_y^i . Then, we obtain boundary conditions and solve the finite-difference equations numerically.

2. Dynamic Mean-Variance Asset Allocation

2.1. 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; Acharya and Pedersen, 2005; Hong, Scheinkman and Xiong, 2006; Brandt, 2009; Campbell, Serfaty-de Medeiros and Viceira, 2009).¹⁷ 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

¹⁷We 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.

to it as the pre-commitment policy. Nevertheless, time-consistency is a basic requirement of rational decision making, and as Strotz (1956) in his original work on time-consistent plans puts it: an investor should choose “... the best plan among those that he will actually follow.” 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 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. Our recursive approach is the one originally suggested by Strotz (1956), although 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 with time-inconsistent preferences (e.g., Peleg and Yaari, 1973; Harris and Laibson, 2001).¹⁸

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. Moreover, given our dynamically optimal mean-variance policy, it is possible to recover time-consistent objective functions that would lead to the same policy (Remark 1). One such function is an increasing, concave, state-dependent criterion of CARA form, a result also generalizing the well-known equivalence of mean-variance and CARA optimization in a one-period Gaussian setting.

We also find the dynamic mean-variance policies to inherit a number of conventional

¹⁸The dynamic inconsistency of preferences is common in economic settings with non-exponential discounting and changing preferences. A growing literature emphasizes the importance of time-inconsistency by demonstrating that it significantly affects consumption and investment decisions of individuals (Harris and Laibson, 2001; Grenadier and Wang, 2007; Ekeland and Lazrak, 2008; Kihlstrom, 2008).

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 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 (2008), 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 (2008) 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 (2009) 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 intertem-

poral 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 2.2, 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 2.3, we provide applications of our analysis to various stochastic investment opportunities, while in Section 2.4, we discuss the extensions to discrete-time, multiple-stock and stochastic interest rate settings. Section 2.5 concludes and the Appendix provides all proofs.

2.2. Asset Allocation with Mean-Variance Criteria

2.2.1. 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

$$\frac{dS_t}{S_t} = \mu(S_t, X_t, t)dt + \sigma(S_t, X_t, t)dw_t, \quad (133)$$

where the stock mean return, μ , and volatility, σ , are deterministic functions of S and the

state variable X , which satisfies

$$dX_t = m(X_t, t)dt + \nu(X_t, t)dw_{xt}. \quad (134)$$

Under appropriate conditions, the stochastic differential equations (133)–(134) 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 (133)–(134). 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

$$dW_t = [rW_t + \theta_t(\mu_t - r)] dt + \theta_t\sigma_t dw_t. \quad (135)$$

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

$$\max_{\theta} E[W_T] - \frac{\gamma}{2} \text{var}[W_T], \quad (136)$$

subject to the dynamic budget constraint (135). In Section 2.2.2, we provide the *time-consistent* solution to this problem via a recursive formulation that employs dynamic programming, while in Section 2.2.4, 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 2.4.3. We further note that even though the mean-variance criterion (136) 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.

2.2.2. 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. Our method is analogous to the original approach of Strotz (1956), recently re-emphasized by Caplin and Leahy (2006), who advocates solving problems with time-inconsistent criteria recursively. The dynamic programming approach here, 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 criteria 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

$$\text{var}_t[W_T] = E_t[\text{var}_{t+\tau}(W_T)] + \text{var}_t[E_{t+\tau}(W_T)], \quad \tau > 0. \quad (137)$$

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 $t + \tau$.¹⁹ 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

$$U_t \equiv E_t[W_T] - \frac{\gamma}{2} \text{var}_t[W_T], \quad (138)$$

subject to the dynamic budget constraint (135). Substituting (137) into (138) and using the law of iterated expectations, we obtain the following recursive representation for the time- t objective function of the mean-variance optimizer:

$$U_t = E_t[U_{t+\tau}] - \frac{\gamma}{2} \text{var}_t[E_{t+\tau}(W_T)]. \quad (139)$$

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

¹⁹Johnsen and Donaldson (1985) provide necessary and sufficient conditions for the time-consistency of intertemporal preferences in a discrete-time setting. These conditions essentially require that initial time-0 preferences can be expressed as an increasing function of preferences at future states/dates, and actions in each state enter separably into the current utility function. The time-inconsistency of our mean-variance criteria can alternatively be demonstrated by explicitly verifying that these conditions fail to hold even in a simple two-period binomial setting.

is not restricted from revising her policy at all times. We elaborate more on the issue of time-consistency in Section 2.2.4.

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 (139) 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

$$E_t[J_{t+\tau}] - \frac{\gamma}{2} \text{var}_t[E_{t+\tau}(W_T)]. \quad (140)$$

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 (140), 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 (140) presented above, and the definition of the value function after some algebra (proof of Lemma 2.1) lead to the following recursive equation for J :

$$J_t = \max_{\theta_s, s \in [t, t+\tau]} E_t[J_{t+\tau}] - \frac{\gamma}{2} \text{var}_t[f_{t+\tau} - f_t + W_{t+\tau}e^{r(T-t-\tau)} - W_t e^{r(T-t)}], \quad (141)$$

subject to the budget constraint (135) and the terminal condition $J_T = W_T$, where f_t is shorthand for $f(W_t, S_t, X_t, t)$ defined as

$$f(W_t, S_t, X_t, t) \equiv E_t[W_T^*] - W_t e^{r(T-t)}, \quad (142)$$

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^* , $s \geq t$.²⁰ The dynamic budget constraint (135) allows us to obtain the following representation for f_t in terms of the optimal stock investment policy θ_s^* :

$$f(W_t, S_t, X_t, t) = E_t \left[\int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds \right]. \quad (143)$$

Going back to (141), 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 2.1. *The value function $J(W_t, S_t, X_t, t)$ of a mean-variance optimizing investor satisfies the following recursive equation:*

$$0 = \max_{\theta_t} E_t [dJ_t] - \frac{\gamma}{2} \text{var}_t [df_t + d(W_t e^{r(T-t)})], \quad (144)$$

subject to $J_T = W_T$ and the budget constraint (135), where f_t is as in (143). Moreover, $J(W_t, S_t, X_t, t)$ is separable in wealth and admits the representation

$$J(W_t, S_t, X_t, t) = W_t e^{r(T-t)} + \tilde{J}(S_t, X_t, t), \quad (145)$$

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.

We note that df_t term in (144) is unaffected by the control θ_t since according to Lemma 2.1, 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 (145) and applying Itô's Lemma to \tilde{J}_t , f_t and $W_t e^{r(T-t)}$, from the HJB equation (144) we obtain

$$0 = \max_{\theta_t} \left\{ \mathcal{D} \tilde{J}_t dt + \theta_t (\mu_t - r) e^{r(T-t)} dt - \frac{\gamma}{2} \text{var}_t \left[\sigma_t S_t \frac{\partial f_t}{\partial S_t} dw_t + \nu_t \frac{\partial f_t}{\partial X_t} dw_{Xt} + \theta_t \sigma_t e^{r(T-t)} dw_t \right] \right\}, \quad (146)$$

²⁰In deriving (141) we use the fact that $\text{var}_t [f_{t+\tau} + W_{t+\tau} e^{r(T-t-\tau)}] = \text{var}_t [f_{t+\tau} - f_t + W_{t+\tau} e^{r(T-t-\tau)} - W_t e^{r(T-t)}]$.

where \mathcal{D} denotes the Dynkin operator.²¹ Computation of the variance term in (146) yields the following PDE for the function \tilde{J}_t :

$$0 = \max_{\theta_t} \left\{ \mathcal{D}\tilde{J}_t + \theta_t(\mu_t - r)e^{r(T-t)} - \frac{\gamma}{2} \left[\sigma_t^2 S_t^2 \left(\frac{\partial f_t}{\partial S_t} \right)^2 + \nu_t^2 \left(\frac{\partial f_t}{\partial X_t} \right)^2 + 2\rho\nu_t\sigma_t S_t \frac{\partial f_t}{\partial S_t} \frac{\partial f_t}{\partial X_t} + \theta_t^2 \sigma_t^2 e^{2r(T-t)} + 2\theta_t\sigma_t \left(\sigma_t S_t \frac{\partial f_t}{\partial S_t} + \rho\nu_t \frac{\partial f_t}{\partial X_t} \right) e^{r(T-t)} \right] \right\}, \quad (147)$$

subject to $\tilde{J}_T = 0$. The HJB equation (147) 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 (147) 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 2.1. *The optimal stock investment policy of a dynamic mean-variance optimizer is given by*

$$\theta_t^* = \frac{\mu_t - r}{\gamma\sigma_t^2} e^{-r(T-t)} - \left(S_t \frac{\partial f_t}{\partial S_t} + \frac{\rho\nu_t}{\sigma_t} \frac{\partial f_t}{\partial X_t} \right) e^{-r(T-t)}, \quad (148)$$

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

$$f(S_t, X_t, t) = E_t \left[\int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds \right]. \quad (149)$$

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,

²¹The Dynkin operator transforms an arbitrary twice continuously differentiable function $F(S_t, X_t, t)$ as follows:

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

following Merton (1971). What is different in our case is that we explicitly identify the hedging demands to be given by the sensitivities 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

$$\theta_{Ht} \equiv - \left(S_t \frac{\partial f_t}{\partial S_t} + \frac{\rho \nu_t}{\sigma_t} \frac{\partial f_t}{\partial X_t} \right) e^{-r(T-t)} = - \frac{\text{cov}_t(dS_t/S_t, df_t)}{\sigma_t^2 dt} e^{-r(T-t)}. \quad (150)$$

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 (148) into (149), we obtain the following representation for f under the original measure P :

$$f(S_t, X_t, t) = E_t \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] - E_t \left[\int_t^T \left(S_s \frac{\partial f_s}{\partial S_s} + \frac{\rho \nu_s}{\sigma_s} \frac{\partial f_s}{\partial X_s} \right) (\mu_s - r) ds \right]. \quad (151)$$

The first component in (151) 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:

$$\begin{aligned} \frac{\partial f_t}{\partial t} + rS_t \frac{\partial f_t}{\partial S_t} + \left(m_t - \rho\nu_t \frac{\mu_t - r}{\sigma_t}\right) \frac{\partial f_t}{\partial X_t} + \frac{1}{2} \left(\sigma_t^2 S_t^2 \frac{\partial^2 f_t}{\partial S_t^2} + \nu_t^2 \frac{\partial^2 f_t}{\partial X_t^2} + 2\rho\nu_t \sigma_t S_t \frac{\partial^2 f_t}{\partial X_t \partial S_t}\right) \\ + \frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma_t}\right)^2 = 0, \end{aligned} \quad (152)$$

with $f_T = 0$. Again, by the Feynman-Kac theorem, (152) admits a unique solution with the following representation:

$$f(S_t, X_t, t) = E_t^* \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds \right], \quad (153)$$

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 drifts²²

$$\frac{dS_t}{S_t} = rdt + \sigma_t dw_t^*, \quad dX_t = \left(m_t - \rho\nu_t \frac{\mu_t - r}{\sigma_t}\right) dt + \nu_t dw_{Xt}^*, \quad (154)$$

and where w_t^* and w_{Xt}^* are Brownian motions under P^* with correlation ρ . Comparing (153) with (151), 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;²³ in the special case of a complete market, the hedge-neutral and risk-neutral measures coincide. Proposition 2.2 summarizes the results above.

²²Since the coefficients assigned to partial derivatives $\partial f_t/\partial S_t$ and $\partial f_t/\partial X_t$ in the PDE (152) represent the drifts of stochastic processes for S and X , it follows that measure P^* modifies the drifts so that S and X satisfy (154).

²³To 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 2.2. *The anticipated portfolio gains, f , can be expressed as*

$$f(S_t, X_t, t) = E_t^* \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right], \quad (155)$$

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

$$dw_t^* = dw_t + \frac{\mu_t - r}{\sigma_t} dt, \quad dw_{X_t}^* = dw_{X_t} + \rho \frac{\mu_t - r}{\sigma_t} dt, \quad (156)$$

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

$$\frac{dP^*}{dP} = e^{-\frac{1}{2} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds - \int_0^T \frac{\mu_s - r}{\sigma_s} dw_s}. \quad (157)$$

Consequently, the optimal investment policy is given by

$$\theta_t^* = \frac{\mu_t - r}{\gamma \sigma_t^2} e^{-r(T-t)} - \frac{1}{\gamma} \left(S_t \frac{\partial E_t^* \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right]}{\partial S_t} + \frac{\rho \nu_t}{\sigma_t} \frac{\partial E_t^* \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right]}{\partial X_t} \right) e^{-r(T-t)}. \quad (158)$$

Proposition 2.2 provides a fully analytical characterization of the optimal investment policy in terms of the model parameters which exists under certain regularity conditions.²⁴ 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 2.3. For economic environments in which explicit computations are not possible, the optimal investment expression (158) can easily be computed numerically by standard Monte Carlo

²⁴In particular, the Proposition proves the existence of the optimal policy θ_t^* satisfying the recursive equations (148)–(149) assuming the exogenous parameters μ_t , σ_t , m_t and ν_t are such that the expectations in (158) are well-defined and twice continuously differentiable (these conditions can explicitly be verified in the specific applications of Section 2.3). Under this assumption f_t is a classical solution of the PDE (152), and hence, f_t and θ_t^* solve equations (148)–(149). 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.

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 (158) 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 2.3.2–2.3.3), 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 2.2.3. The correlation parameter also affects the joint probability distribution of the stock and state variable under which the expressions in (158) are evaluated. This indirect correlation effect will be assessed in the applications studied in Section 2.3. 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 (158)) due to the market price of risk possibly being dependent 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 $T - t$, in general the investor's optimal risky stock investment may increase or decrease as the investment horizon increases. Nevertheless, we see that longer time-horizons decrease the myopic demand in (158) (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 increases, while the investor's myopic demand decreases, her expected portfolio gains are higher, which may increase the hedging demand. For short time-horizons, the former effect dominates, and the hedging demand vanishes as the horizon $T - t$ is shortened. In the applications of Section 2.3, we demonstrate that the optimal risky investment may either increase or decrease as the horizon increases, depending on the specific economic setting.

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 (158) 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. Another common feature is the presence of the multiplicative discounting term in (158), which is also present in the literature studying optimal CARA investments (Merton, 1971; Cox and Huang, 1989).²⁵

Remark 1 (Recovering time-consistent objective functions). It is of interest to see

²⁵Indeed, for CARA preferences with risk aversion parameter γ , it can be demonstrated that $\gamma e^{r(T-t)} = -J_{WW}/J_W$ where J denotes the investor's value function, and hence $\gamma e^{r(T-t)}$ can be interpreted as absolute risk aversion.

whether there are economically meaningful time-consistent objective functions leading to our dynamically optimal investment policy (158). 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:

$$\max_{\theta_t} E_t \left[-\varepsilon_T e^{-\gamma W_T} \right], \quad (159)$$

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 (155), subject to the budget constraint (135). 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 (158). To understand the intuition behind the process ε_t we observe (from the optimal wealth (161)) 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 (159) puts higher weight on those states of the economy in which the optimal wealth process is less volatile along its path. Since $\varepsilon_t > 0$ we may construct a new probability measure Q defined by the Radon-Nikodym derivative $\varepsilon_T/E_t[\varepsilon_T]$ with respect to the original measure P so that the optimization problem (159) is equivalent to

$$\max_{\theta_t} E_t^Q \left[-e^{-\gamma W_T} \right], \quad (160)$$

where $E_t^Q[\cdot]$ denotes the expectation under the beliefs represented by measure Q . Furthermore, the CARA-type dynamic optimization problems, (159) or (160), leading to the same solution as our dynamic mean-variance problem generalizes the well-known equivalence of mean-variance and CARA optimization in a one-period setting with normally distributed stock returns. We finally note that there are other time-consistent, state-dependent objec-

tive functions leading to the optimal policy (158).²⁶

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, as in the literature on consumer behavior under time-inconsistent preferences (e.g., Peleg and Yaari, 1973; 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 (138) over terminal wealth, and chooses a time- t Markovian investment strategy $\theta(W_t, S_t, X_t, t)$ subject to the budget constraint (135). 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^*

²⁶In particular, using the results of Lemma 2.1 and Proposition 2.1, it is straightforward to see that the dynamically optimal policy (26) can be obtained by solving a time-consistent instantaneous mean-variance problem of the form

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

with f_t given by (155), subject to budget constraint (135). Interestingly, this objective becomes myopic in the restrictive special case of deterministic anticipated portfolio gains $f_t, \max_{\theta_t} \theta_t(\mu_t - r)e^{r(T-t)} - (\gamma/2)\theta_t^2 \sigma_t^2 e^{2r(T-t)}$. Moreover, solving the instantaneous problem generally can be demonstrated to be equivalent to a CARA-type problem

$$\max_{\theta_t} E_t[-de^{-\gamma(W_t e^{r(T-t)} + f_t)}],$$

since $dW_t e^{r(T-t)} + df_t$ is locally normally distributed. This is again analogous to the equivalence of the mean-variance and CARA portfolio choice problems in a one-period setting under normality. Furthermore, since $J_t = -\exp(-\gamma(W_t e^{r(T-t)} + f_t))$ is a value function for a CARA-type criterion in (159), it turns out that dynamic mean-variance portfolio choice is equivalent to the portfolio choice with monotonic CARA-type preferences in (159).

solves the dynamic optimization problem (144).

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 (148), 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 (158).

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

Proposition 2.3. *The optimal terminal wealth, its mean, variance and the value function of a dynamic mean-variance optimizer are given by*

$$W_T^* = W_t e^{r(T-t)} + f_t + \frac{1}{\gamma} \int_t^T \frac{\mu_s - r}{\sigma_s} dw_s + \sqrt{1 - \rho^2} \int_t^T \nu_s \frac{\partial f_s}{\partial X_s} d\tilde{w}_s, \quad (161)$$

$$\text{var}_t[W_T^*] = \frac{1}{\gamma^2} E_t \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] + (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 \left(\frac{\partial f_s}{\partial X_s} \right)^2 ds \right], \quad (162)$$

$$E_t[W_T^*] = W_t e^{r(T-t)} + f_t, \quad (163)$$

$$J_t = W_t e^{r(T-t)} + f_t - \frac{1}{2\gamma} E_t \left[\int_t^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] - \frac{\gamma}{2} (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 \left(\frac{\partial f_s}{\partial X_s} \right)^2 ds \right], \quad (164)$$

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 (161)), capturing anticipated bond and stock gains, and risky terms driven by the hedgeable stock uncertainty (third term in (161)) and unhedgeable uncertainty (fourth term in (161)). 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 $\rho^2 = 1$.²⁷

The variance of optimal terminal wealth is determined by the variances of hedgeable (first term in (162)) and unhedgeable uncertainties (second term in (162)). 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 2.3(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.²⁸ This indirect effect can be disentangled in the applications, where the expectation under measure P^* can explicitly be computed.

²⁷As for CARA and quadratic preferences, the optimal mean-variance terminal wealth in (161) may become negative. In future work, it would be of interest to incorporate a non-negativity constraint on terminal wealth into our analysis which would ensure that the investor remains solvent.

²⁸This 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 (154).

The effect of market incompleteness on the mean of terminal wealth enters via the anticipated gains, f . Proposition 2.3(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 2.2.2. 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 2.3, the sign of the hedging demand can readily be identified in particular applications, simplifying the analysis in incomplete markets.

2.2.4. Optimal Pre-commitment Policy

In Section 2.2.2, 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. Thus, if the investor were able to pre-commit, her time-0 value function would have been higher than the one under the time-consistent policy.

The pre-commitment mean-variance problem and its variations have been analyzed in the literature, amongst others, by Bajeux-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

$$\xi_t = \xi_0 e^{-rt - \frac{1}{2} \int_0^t \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds - \int_0^t \frac{\mu_s - r}{\sigma_s} dw_s}. \quad (165)$$

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:

$$\max_{W_T} E_0[W_T] - \frac{\gamma}{2} \text{var}_0[W_T], \quad (166)$$

$$\text{subject to } E_0[\xi_T W_T] \leq W_0. \quad (167)$$

Proposition 2.4 presents the optimal solution to this problem in terms of the state price density.

Proposition 2.4. *The optimal terminal wealth of a mean-variance optimizer under pre-commitment is given by*

$$\hat{W}_T = W_0 e^{rT} + \frac{1}{\gamma} E_0[\xi_T^2] e^{2rT} - \frac{1}{\gamma} \xi_T e^{rT}. \quad (168)$$

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

$$\hat{W}_T = W_0 e^{rT} + \frac{1}{\gamma} e^{(\frac{\mu-r}{\sigma})^2 T} - \frac{1}{\gamma} \xi_T e^{rT}, \quad (169)$$

$$\hat{\theta}_t = \frac{\mu - r}{\gamma \sigma^2} e^{-r(T-t)} \xi_t e^{(\frac{\mu-r}{\sigma})^2 (T-t) + rt}. \quad (170)$$

To facilitate comparisons with the pre-commitment solution above, we also provide the time-consistent solution (Propositions 2.2–2.3). 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

$$W_T^* = W_0 e^{rT} + \frac{1}{\gamma} E_0 \left[\xi_T e^{rT} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right] - \frac{1}{\gamma} \left[\ln \xi_T + rT + \frac{1}{2} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds \right]. \quad (171)$$

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

$$W_T^* = W_0 e^{rT} + \frac{1}{\gamma} \left(\frac{\mu - r}{\sigma} \right)^2 T - \frac{1}{\gamma} \left[\ln \xi_T + rT + \frac{1}{2} \left(\frac{\mu - r}{\sigma} \right)^2 T \right], \quad (172)$$

$$\theta_t^* = \frac{\mu - r}{\gamma \sigma^2} e^{-r(T-t)}. \quad (173)$$

Clearly, the pre-commitment solution ((168)–(170)) and the time-consistent solution ((171)–(173)) 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.²⁹ 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 (169) 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.³⁰

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

²⁹To see this, observe that the expected wealth under pre-commitment (169) grows exponentially with the horizon, while the expected wealth under time-consistency (172) grows linearly. Even though the variance is also higher in the pre-commitment case, it can be verified that the time-0 probability $\text{Prob}_0(\hat{W}_T > W_T^*)$ approaches unity with long horizons. Indeed, from (165), (169) and (172) $\hat{W}_T - W_T^*$ can be written as a function of w_T/\sqrt{T} . Moreover, it can be shown that this function is positive whenever $-0.5((\mu - r)/\sigma)\sqrt{T} \leq w_T/\sqrt{T} \leq 0.5((\mu - r)/\sigma)\sqrt{T}$ and T is sufficiently large. Therefore, since w_T/\sqrt{T} is standard normal, $\hat{W}_T \geq W_T^*$ with probability approaching unity as $T \rightarrow \infty$.

³⁰This 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 (172).

because bad, costly states reduce her expected terminal wealth. To offset this, the investor takes on more risk by increasing her risky investment.

2.3. Applications

This Section provides several applications that illustrate the simplicity and the usefulness of the methodology developed in Section 2.2. In Sections 2.3.1–2.3.3, 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 2.3.2–2.3.3.

2.3.1. 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:

$$\frac{dS_t}{S_t} = \mu dt + \bar{\sigma} S_t^{\alpha/2} dw_t, \quad (174)$$

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 2.2. 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 2.2 presents the optimal investment policy, as well as some of its properties.

Corollary 2.1. *The optimal stock investment policy for the CEV model (174) is given by:*

$$\theta_t^* = \frac{\mu - r}{\gamma \bar{\sigma}^2 S_t^\alpha} e^{-r(T-t)} - \frac{1}{\gamma} \left(\frac{\mu - r}{\bar{\sigma} S_t^{\alpha/2}} \right)^2 \frac{e^{-\alpha r(T-t)} - 1}{r} e^{-r(T-t)}. \quad (175)$$

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 is a decreasing function of the time-horizon, $T - t$, for $\alpha \leq -1$ and $\mu - r > 0$, and is non-monotonic otherwise.*

Corollary 2.1(i) reveals that the sign of the hedging demand (the second term in (175)) 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 (155)) since the latter are positively related to the market price of risk. As discussed in Section 2.2.2, this gives rise to a positive hedging demand. Analogously, the hedging demand is negative for negative elasticity.

Property (ii) of Corollary 2.1 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 can be either decreasing or non-monotonic as the time-horizon increases, due to two effects working in opposite directions. On one hand, the investment is perceived riskier at longer horizons which decreases the demand for stock. On the other hand, the anticipated gains may become higher with longer horizon, inducing larger hedging demands in absolute terms. As a result, hedging demands can be either increasing or decreasing functions of the horizon, depending on their sign. Thus, depending on which effect is stronger for a given horizon, the optimal risky investment may increase or decrease as the horizon increases, which results in a non-monotonic pattern.

2.3.2. Stochastic Volatility

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

$$\frac{dS_t}{S_t} = (r + \delta X_t^{(1+\beta)/2\beta})dt + X_t^{1/2\beta} dw_t, \quad (176)$$

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

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}\sqrt{X_t}dw_{Xt}, \quad (177)$$

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 2.2. 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 (154)) to also follow a mean-reverting, square-root process as in (177). 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 2.2 reports our solution and some of its properties.

Corollary 2.2. *The optimal stock investment policy for the stochastic-volatility model (176)–(177) is given by:*

$$\theta_t^* = \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} e^{-r(T-t)} - \rho\bar{\nu}\delta \left(\frac{1 - e^{-(\lambda+\rho\bar{\nu}\delta)(T-t)}}{\lambda + \rho\bar{\nu}\delta} \right) \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} e^{-r(T-t)}. \quad (178)$$

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 is increasing in the time-horizon $T-t$ for $\lambda+r+\rho\bar{\nu}\delta < 0$, is decreasing for $\rho\bar{\nu}\delta > 0$, and is non-monotonic otherwise;*
- (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 2.2(i) shows that the sign of the hedging demand (second term in (178)) 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 2.2.2, 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 2.2 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.³¹ 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 increases or decreases as

³¹The optimal investment, however, may become negative for negative speed of mean-reversion λ , which corresponds to an explosive process for the state variable.

the time-horizon increases, similarly to the setting in Section 2.3.1. This horizon effect depends on the covariance between the stock returns and state variable per unit of stock volatility, $\rho\bar{\nu}$, 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 2.2(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 2.3, the variance of terminal wealth is lowest in the complete market case ($\rho^2 = 1$), in which perfect hedging is possible, 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 2.1, 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 2.1 presents the percentage hedging demand for varying levels of correlation,

Table 2.1
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					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

speed of mean-reversion and the investor's horizon.³² 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 2.1, 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

³²We 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 .

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.³³ Thus, the hedging demand in our setting is larger in absolute terms than in Chacko and Viceira.

2.3.3. 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:

$$\frac{dS_t}{S_t} = (r + \sigma X_t)dt + \sigma dw_t, \quad (179)$$

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

$$dX_t = \lambda(\bar{X} - X_t)dt + \nu dw_{Xt}, \quad (180)$$

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

³³Chacko 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.

infinite-horizon discrete-time consumption and portfolio choice of an investor with recursive utility and under discrete-time versions of the dynamics (179)–(180), 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 (154) that under this measure, the market price of risk follows a simple mean-reverting process as in (180) 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 2.2 reports our mean-variance solution and some of its unambiguous properties.

Corollary 2.3. *The optimal stock investment policy for the time-varying Gaussian mean returns model (179)–(180) is given by:*

$$\theta_t^* = \frac{X_t}{\gamma\sigma} e^{-r(T-t)} - \frac{\rho\nu}{\gamma\sigma} \left(\lambda \left(\frac{1 - e^{-(\lambda+\rho\nu)(T-t)}}{\lambda + \rho\nu} \right)^2 \bar{X} + \frac{1 - e^{-2(\lambda+\rho\nu)(T-t)}}{\lambda + \rho\nu} X_t \right) e^{-r(T-t)}. \quad (181)$$

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 (181)) 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 2.2(i)).

Table 2.2
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

Corollary 2.3(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).³⁴ 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 2.2

³⁴The 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.

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 2.2.

Inspection of Table 2.2 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 (2002).

2.4. Extensions and Ramifications

In this Section, we demonstrate that the baseline analysis of Section 2.2 can easily be adopted to alternative or richer economic environments. Section 2.4.1 illustrates our methodology in a discrete-time framework, and provides an explicit solution to the stochastic-volatility model in discrete time. Sections 2.4.2 – 2.4.3 demonstrate that the results of Section 2.2 are readily extendable to more realistic environments with stochastic interest rates and with multiple stocks, state variables and sources of uncertainty.

2.4.1. 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.³⁵

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 Δt , $D(0, \Delta t)$. 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

$$\Delta W_t = [rW_t + \theta_t(\mu_t - r)]\Delta t + \theta_t\sigma_t\Delta w_t, \quad (182)$$

where θ_t again denotes the dollar stock investment. The investor maximizes the objective function (138) subject to the dynamic budget constraint (182) for each time $t = 0, \Delta t, \dots, T - \Delta t$. Proposition 2.5 is the discrete-time analogue of Proposition 2.1 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

³⁵Since 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.

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

$$\theta_t^* = \frac{\mu_t - r}{\gamma\sigma_t^2} R^{-(T-\Delta t-t)} - \frac{\text{cov}_t(\Delta S_t/S_t, \Delta f_t)}{\sigma_t^2 \Delta t} R^{-(T-\Delta t-t)}, \quad (183)$$

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

$$f(S_t, X_t, t) = E_t \left[\sum_{s=t}^{T-\Delta t} \theta_s^* (\mu_s - r) R^{(T-\Delta t-s)} \Delta t \right], \quad (184)$$

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

The discrete-time optimal investment policy has the same structure as in Proposition 2.1 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 2.2. Nevertheless, expression (183) can be used to obtain an explicit representation for the optimal policy for specific applications either by solving (183) backwards or by guessing the structure of the solution.

To illustrate an application of Proposition 2.5, we solve the discrete-time versions of the models of Sections 2.3.2–2.3.3. Specifically, the discrete-time dynamics of the stock price and state variable for the stochastic-volatility model are specified as follows:

$$\frac{\Delta S_t}{S_t} = (r + \delta X_t^{(1+\beta)/2\beta}) \Delta t + X_t^{1/2\beta} \Delta w_t, \quad (185)$$

$$\Delta X_t = \lambda(\bar{X} - X_t) \Delta t + \bar{v} \sqrt{X_t} \Delta w_{Xt}, \quad (186)$$

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 (183) 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:

$$\frac{\Delta S_t}{S_t} = (r + \sigma X_t)\Delta t + \sigma \Delta w_t, \quad (187)$$

$$\Delta X_t = \lambda(\bar{X} - X_t)\Delta t + \nu \Delta w_{Xt}, \quad (188)$$

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 (183), as in the previous case, we obtain recursive equations for $g_1(t)$ and $g_2(t)$ which we solve explicitly. Corollary 2.4 reports the results.

Corollary 2.4. *The optimal investment policy for the discrete-time stochastic-volatility model (185)–(186) is given by*

$$\theta_t^* = \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} R^{-(T-\Delta t-t)} - \rho\bar{\nu}\delta \frac{1 - (1 - (\lambda + \rho\bar{\nu}\delta)\Delta t)^{(T-\Delta t-t)/\Delta t}}{\lambda + \rho\bar{\nu}\delta} \frac{\delta}{\gamma} X_t^{(\beta-1)/2\beta} R^{-(T-\Delta t-t)}, \quad (189)$$

and for the discrete-time model with Gaussian mean-returns (187)–(188) is given by

$$\theta_t^* = \frac{X_t}{\gamma\sigma} R^{-(T-\Delta t-t)} - \frac{g_1(t) + g_2(t)X_t}{\gamma\sigma} R^{-(T-\Delta t-t)}, \quad (190)$$

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 2.2 and 2.3. As a result, the comparative statics for (189) and (190) are similar to those in the continuous-time case. We note that in deriving expressions (189) and (190), we do not assume normality of the stochastic processes w and w_X , as in continuous-time.

2.4.2. Multiple Stock Formulation

We now generalize the baseline analysis of Section 2.2 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

$$dW_t = [rW_t + \theta_t^\top (\mu_t - r)]dt + \theta_t^\top \sigma_t dw_t, \quad (191)$$

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 2.2. For each time $t \in [0, T]$, she maximizes the time- t objective function (138) subject to the dynamic budget

constraint (191). As in Section 2.2, the optimal policy is characterized in terms of the anticipated portfolio gains, f , and arises from the HJB equation adjusted for time-inconsistency. Proposition 2.6 generalizes Proposition 2.2 and reports the optimal investment and anticipated gains in terms of the model parameters and the hedge-neutral measure.

Proposition 2.6. *The optimal investment policy in the multiple-stock economy is given by*

$$\theta_t^* = \frac{1}{\gamma} (\sigma_t \sigma_t^\top)^{-1} (\mu_t - r) e^{-r(T-t)} - \left(I_{S_t} \frac{\partial f_t}{\partial S_t^\top} + (\nu_t \rho^\top \sigma_t^{-1})^\top \frac{\partial f_t}{\partial X_t^\top} \right) e^{-r(T-t)}, \quad (192)$$

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 (192) 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 2.6 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 2.2. 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 2.2.3.

2.4.3. 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 2.2. 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

$$dr_t = \mu_r(X_t, r_t, t)dt + \sigma_r(X_t, r_t, t)dw_{rt}, \quad (193)$$

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, $\tilde{\theta}_t \equiv \theta_t^*/B_t$.³⁶ Proposition 2.7 reports our results.

Proposition 2.7. *The optimal investment policy in the economy with stochastic interest rates is given by*

$$\tilde{\theta}_t = \frac{\mu_t - r_t}{\gamma\sigma_t^2} - \left(S_t \frac{\partial f_t}{\partial S_t} + \frac{\rho\nu_t}{\sigma_t} \frac{\partial f_t}{\partial X_t} + \frac{\rho_{rS}\sigma_{rt}}{\sigma_t} \frac{\partial f_t}{\partial r_t} \right), \quad (194)$$

where f_t is as in Proposition 2.2, but with r following (193).

³⁶Otherwise, if the mean-variance criterion is over W_T and the interest rate is stochastic, in contrast to Lemma 2.1 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.

The optimal policy (194) 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 2.3, the optimal 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)

$$dr_t = \lambda_r(\bar{r} - r_t)dt + \sigma_r dw_{rt}. \quad (195)$$

Along the lines of Corollaries 2.1–2.3, 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 (181) in Section 2.3.3, but now additionally allows us to consider comparative statics with respect to the parameters of the interest rate dynamics (195).

2.5. 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. For example, portfolio selection problems in incomplete markets are notoriously hard to solve, and it may be very convenient for future applications to have a reasonable criterion for which the solution can actually be computed. In our ongoing work, we extend the current study to an incomplete market setting where the investment opportunities may also experience jumps and again discover much tractability in the ensuing analysis. We also foresee potential applications in security pricing with incomplete markets, for which the investor preferences are to be accounted for.

2.6. Appendix: Proofs

Proof of Lemma 2.1. We first demonstrate equation (141). Using Itô's Lemma we rewrite the budget constraint (135) as

$$d(W_t e^{r(T-t)}) = \theta_t(\mu_t - r)e^{r(T-t)}dt + \theta_t \sigma_t e^{r(T-t)}dw_t. \quad (196)$$

Integrating (196) from $t + \tau$ to T and assuming that the investor follows the optimal policy θ_s^* from $t + \tau$ onwards we obtain

$$E_{t+\tau}[W_T^*] = W_{t+\tau}e^{r(T-t-\tau)} + f_{t+\tau}, \quad (197)$$

where W_T^* is optimal wealth given the optimal policy θ_s^* , $s \geq t + \tau$, and f_t is given by (143). Substituting (197) into (140) and noting that $W_t e^{r(T-t)}$ and f_t are adapted to the filtration \mathcal{F}_t , we obtain (141).

The HJB equation in differential form (144) follows from equation (141) 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^* , we integrate (196) from t to T and substitute W_T into the time- t objective function:

$$\begin{aligned} E_t[W_T] - \frac{\gamma}{2}\text{var}_t[W_T] &= W_t e^{r(T-t)} + E_t \left[\int_t^T \theta_s(\mu_s - r)e^{r(T-s)}ds \right] \\ &\quad - \frac{\gamma}{2}\text{var}_t \left[\int_t^T \theta_s(\mu_s - r)e^{r(T-s)}ds + \int_t^T \theta_s \sigma_s e^{r(T-s)}dw_s \right]. \end{aligned} \quad (198)$$

It can be observed from (198) 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 (143). The separability of the value function J_t from $W_t e^{r(T-t)}$ follows from (198).

Q.E.D.

Proof of Proposition 2.1. To prove Proposition 2.1, it remains to derive the first order condition for the problem (147).³⁷ The objective function in (147) is quadratic and concave in θ_t , and so the unique optimal policy solves the first order condition:

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

leading to the expression (148).

Q.E.D.

Proof of Proposition 2.2. Under standard conditions, there exists a probability measure P^* under which the function f_t admits the Feynman-Kac representation (153) (Karatzas and Shreve, 1991) and under this measure, the processes S and X satisfy the stochastic differential equations (154). Comparing (154) with (133)–(134), we obtain that measure P^* transforms Brownian motions w_t and w_{Xt} into w_t^* and w_{Xt}^* satisfying (156).

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 (156) 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 (157) obtains by Girsanov's Theorem. Finally, substituting f_t given by (155) into the recursive representation (148), we obtain (158).
Q.E.D.

³⁷The proof of Proposition 2.1 implicitly assumes that $f(S_t, X_t, t)$ is twice differentiable and

$$E_t\left[\int_t^\tau \nu_s^2\left(\frac{\partial f_s}{\partial X_s}\right)^2 + (\sigma_s S_s)^2\left(\frac{\partial f_s}{\partial S_s}\right)^2 ds\right] < \infty.$$

These conditions justify the computation of the variance term in (146) (e.g., Karatzas and Shreve, 1991) which yields the HJB equation (147), and can be verified in the specific applications of Section 2.3 where the process f_t is computed explicitly.

Proof of Proposition 2.3. First, we derive the terminal wealth expression. Substituting the optimal policy θ_t^* from (148) into (196) and rearranging terms we obtain

$$d(W_t e^{r(T-t)}) = -df_t + \frac{\mu_t - r}{\gamma \sigma_t} dw_t + \sqrt{1 - \rho^2} \nu_t \frac{\partial f_t}{\partial X_t} d\tilde{w}_t, \quad (199)$$

where \tilde{w}_t is defined in Proposition 2.3. Integrating (199) from t to T , we obtain (161). 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 (162). 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 (162). Since $(\mu_t - r)/\sigma_t$ is assumed to not depend on S_t , the first term in (162) 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 (155) also depends only on X_t . As a result, the PDE (152) for f_t becomes:

$$\frac{\partial f_t}{\partial t} + \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} \right) \frac{\partial f_t}{\partial X_t} + \frac{\nu_t^2}{2} \frac{\partial^2 f_t}{\partial X_t^2} + \frac{1}{\gamma} \left(\frac{\mu_t - r}{\sigma_t} \right)^2 = 0, \quad (200)$$

with $f_T = 0$. Differentiating (200) with respect to ρ and denoting $\tilde{f}_t \equiv \partial f_t / \partial \rho$, we obtain the equation for \tilde{f}_t :

$$\frac{\partial \tilde{f}_t}{\partial t} + \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} \right) \frac{\partial \tilde{f}_t}{\partial X_t} + \frac{\nu_t^2}{2} \frac{\partial^2 \tilde{f}_t}{\partial X_t^2} - \nu_t \frac{\mu_t - r}{\sigma_t} \frac{\partial f_t}{\partial X_t} = 0, \quad (201)$$

where $\tilde{f}_T = 0$. Applying the Feynman-Kac Theorem to equation (201), using the expression (150) 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 2.4. The optimal pre-commitment terminal wealth \hat{W}_T solves the first order condition

$$1 - \gamma\hat{W}_T + \gamma E_0[\hat{W}_T] - \psi\xi_T = 0, \quad (202)$$

where ψ is the Lagrange multiplier of the static budget constraint (167). Taking time-zero expectation on both sides of (202) yields $\psi = 1/E_0[\xi_T]$, or $\psi = e^{rT}$, since $E_0[\xi_T e^{rT}] = 1$. Substituting ψ back into (202) we obtain

$$\hat{W}_T = \frac{1}{\gamma} \left(1 + \gamma E_0[\hat{W}_T] - \xi_T e^{rT} \right). \quad (203)$$

(203) substituted into the static budget constraint (167) leads to $\gamma E_0[\hat{W}_T] = \gamma W_0 e^{rT} - 1 + E_0[\xi_T^2] e^{2rT}$, which along with (203) yields the optimal terminal wealth (168).

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 (169). To compute the pre-commitment investment policy, $\hat{\theta}_t$, we first consider the optimal time- t wealth:

$$\hat{W}_t = E_t \left[\frac{\xi_T}{\xi_t} \hat{W}_T \right] = a(t) - \frac{1}{\gamma} e^{-(2r - (\mu - r)^2/\sigma^2)(T-t)} e^{rT} \xi_t, \quad (204)$$

where the second equality follows by substituting \hat{W}_T from (169) and evaluating the moments of ξ_T , and $a(t)$ is a deterministic function of time. Applying Itô's Lemma to (204) 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 (135) yields $\hat{\theta}_t$ in (170). *Q.E.D.*

Proof of Corollary 2.1. In the case of a complete market, measure P^* coincides with the risk-neutral one. To compute the optimal investment policy from Proposition 2.2, 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:

$$dS_t^{-\alpha} = \left(-\alpha r S_t^{-\alpha} + \frac{\alpha(1+\alpha)\bar{\sigma}^2}{2} \right) dt - \alpha \bar{\sigma} S_t^{-\alpha} dw_t^*. \quad (205)$$

Integrating (205) from t to s and taking the time- t expectation under the risk-neutral measure on both sides, we obtain the equation for $g(t, s)$:

$$g(t, s) = S_t^{-\alpha} - \int_t^s \left(\alpha r g(t, y) - \frac{\alpha(1+\alpha)\bar{\sigma}^2}{2} \right) dy. \quad (206)$$

Differentiating (206) with respect to s yields the linear differential equation

$$\frac{\partial g(t, s)}{\partial s} = -\alpha r g(t, s) + \frac{\alpha(1+\alpha)\bar{\sigma}^2}{2}, \quad (207)$$

with initial condition $g(t, t) = S_t^{-\alpha}$. The unique solution to equation (207) is given by

$$g(t, s) = S_t^{-\alpha} e^{-\alpha r(s-t)} + (1+\alpha)\bar{\sigma}^2 \frac{1 - e^{-\alpha r(s-t)}}{2r}. \quad (208)$$

Substitution of (208) into the optimal investment policy (158) leads to the θ_t^* reported in Corollary 2.1. We also note that the process for the market price of risk is explosive for $\alpha \leq -1$ since the conditional expectation (208) is unbounded for large horizons. For $-1 < \alpha < 0$, the conditional expectation (208) 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 (175). Property (ii) follows from (175) and the fact that the hedging demand is negative for $\alpha < 0$. Finally, property (iii) can be demonstrated by observing that the derivative of θ_t^* with respect to the time-horizon is negative for $\alpha \leq -1$ and $\mu > r > 0$, and can both be negative or positive otherwise, depending on the horizon, $T - t$.

Q.E.D.

Proof of Corollary 2.2. 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 (154) that the state variable under measure P^* follows the process

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

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

$$E_t^*[X_s] = \frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}\delta} + \left(X_t - \frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}\delta} \right) e^{-(\lambda + \rho\bar{\nu}\delta)(s-t)}. \quad (209)$$

Substituting (209) into (158) 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 (178)) 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 derivative of θ_t^* with respect to the time-horizon is positive for $\lambda + r + \rho\bar{\nu}\delta < 0$, negative for $\rho\bar{\nu}\delta > 0$, and can both be negative or positive otherwise, depending on the horizon.

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 (209), $E_t^*[X_s] = \int_t^s e^{-(\lambda + \rho\bar{\nu}\delta)(y-t)} dy + X_t e^{-(\lambda + \rho\bar{\nu}\delta)(s-t)}$, which is clearly decreasing in correlation ρ . Similarly, using Proposition 2.3, 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 2.3. 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):

$$E_t^*[X_s^2] = \left(\frac{\lambda \bar{X}}{\lambda + \rho\nu} + \left(X_t - \frac{\lambda \bar{X}}{\lambda + \rho\nu} \right) e^{-(\lambda + \rho\nu)(s-t)} \right)^2 + \nu^2 \frac{1 - e^{-2(\lambda + \rho\nu)(s-t)}}{2(\lambda + \rho\nu)}. \quad (210)$$

Substituting (210) into (158) 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 2.2(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 2.5. The proof is similar to the proof of Proposition 2.1. 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 2.4. 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 (185)–(186) into the expression for f_t (184), we obtain:

$$f_t = E_t \left[\sum_{s=t}^{T-\Delta t} \delta X_s g(s) \Delta t \right]. \quad (211)$$

To compute $E_t[X_s]$, we take expectations of both sides of the state variable process, (186), and obtain a difference equation for $E_t[X_s]$:

$$E_t[X_{s+\Delta t}] = \lambda \bar{X} \Delta t + (1 - \lambda \Delta t) E_t[X_s], \quad (212)$$

with initial condition $E_t[X_t] = X_t$. The unique solution to equation (212) is

$$E_t[X_s] = (X_t - \bar{X})(1 - \lambda\Delta t)^{(s-t)/\Delta t} + \bar{X}. \quad (213)$$

Substituting (213) into (211) yields:

$$f_t = d(t) + \delta X_t \sum_{s=t}^{T-\Delta t} g(s)(1 - \lambda\Delta t)^{(s-t)/\Delta t} \Delta t, \quad (214)$$

where $d(t)$ denotes a time-deterministic function. Using (214), we compute Δf_t and substitute it into the recursive expression for the optimal strategy (183). Taking into account the conjecture for θ_t^* , after some algebra we obtain a recursive equation for $g(t)$:

$$g(t) = \delta/\gamma - \rho\bar{\nu}\delta \sum_{s=t+\Delta t}^{T-\Delta t} g(s)(1 - \lambda\Delta t)^{(s-\Delta t-t)/\Delta t} \Delta t. \quad (215)$$

Evaluating (215) at time $t - \Delta t$ and then subtracting it from (215), 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 (187)–(188), 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 (184). Then, substituting f_t into (183) 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 2.4, 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 $\varphi = \text{cov}(\Delta w, \Delta w_X^2)$.³⁸

Q.E.D.

³⁸It can easily be demonstrated that $\varphi = 0$ if Δw and Δw_X are normally distributed.

Proof of Proposition 2.6. The proof is a multi-dimensional version of the proofs for Propositions 2.1–2.2.

Q.E.D.

Proof of Proposition 2.7. The proof is similar to those of Propositions 2.1–2.2, but now accounting for the mean-variance criterion being over W_T/B_T . As in the proof of Lemma 2.1, 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 2.2 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 (194). Employing measure P^* it can be shown that f_t is the same as in Proposition 2.2, but now with stochastic r_t .

Q.E.D.

3. Dynamic Hedging in Incomplete Markets: A Simple Solution

3.1. Introduction

Perfect hedging is a risk management activity that aims to eliminate risk completely. In theory, perfect hedges are possible via dynamic trading in frictionless complete markets and are obtained by standard no-arbitrage methods (e.g., Cvitanic and Zapatero, 2004). In reality, however, “perfect hedges are rare,” as simply put by Hull (2008). Despite the unprecedented development in the menu of financial instruments available, market frictions render markets incomplete, making perfect hedging impossible. Consequently, hedging in incomplete markets has much occupied the profession. The traditional, pragmatic approach is to employ static minimum-variance hedges (e.g., Stulz, 2003; McDonald, 2006; Hull, 2008) or the corresponding myopic hedges that repeat the static ones over time. While intuitive and tractable, these hedges are not necessarily optimal in multi-period settings and may lead to significant welfare losses (e.g., Brandt, 2003). Moreover, they do not generally provide perfect hedges in dynamically complete markets. The alternative route is to consider richer dynamic incomplete-market settings and characterize hedges that maximize a hedger’s preferences or provide the best hedging quality. The latter is measured by various criteria in terms of means and variances of the *hedging error*, as given by the deviation of the hedge from its target value. Despite much work, the literature still lacks tractable dynamic hedges in plausible stochastic environments, with explicit solutions arising in a few settings (typically with constant means and volatilities of pertinent processes).

In this paper, we provide tractable dynamically optimal hedges in a general incomplete-market economy by employing the minimum-variance criterion. We demonstrate that these hedges retain the basic structure of perfect hedges, as well as the intuitive elements of the static minimum-variance hedges. Towards that, we consider a hedger who is concerned with reducing the risk of a non-tradable or illiquid asset, or a contingent claim at some future date. Notable examples include various commodities, human capital, housing, commercial

properties, various financial liabilities, executive stock options. The market is incomplete in that the hedger cannot take an exact offsetting position to the non-tradable asset payoff by dynamically trading in the available securities, a bond and a stock (or futures or any other derivative) that is correlated with the non-tradable. We employ the familiar minimum-variance criterion for the quality of the hedging but considerably differ from the literature in that we account for the time-inconsistency of this criterion and obtain the solution by dynamic programming. We here follow a methodology developed in the context of dynamic mean-variance portfolio choice in Chapter 2. In dynamically complete markets, there is no time-inconsistency issue (unlike the problem in Chapter 2) and our dynamically optimal minimum-variance hedges reduce to perfect hedges, unlike their static or myopic counterparts. In incomplete markets, we show that the variance criterion becomes time-consistent only when the stock has zero risk premium or when considered under any risk-neutral probability measure (which is not unique here). Our dynamically optimal hedge can then alternatively be obtained by minimizing such a criterion under a specific risk-neutral measure.

We obtain a fully analytical characterization of the dynamically optimal minimum-variance hedges in terms of the exogenous model parameters. The complete-market dynamic hedge, obtained by no-arbitrage, is determined by the “Greeks” that quantify the sensitivities of the asset value under the unique risk-neutral measure to the pertinent stochastic variables in the economy. Ours is given by generalized Greeks, still representing the asset value sensitivities to the same variables, but now in terms of an additional parameter accounting for the market incompleteness and where the asset-value is under a specific risk-neutral measure accounting for the hedging costs. The hedges are in terms of the Greeks since, as we demonstrate, a higher variability of asset value implies a lower quality of hedging, and hence the need to account for asset-value sensitivities. We further demonstrate the tractability and practical usefulness of our solution by explicitly computing the hedges for plausible intertemporal economic environments with stochastic market prices of risk and volatilities of non-tradable asset and stock returns.

We next compare the performances of our dynamically optimal hedges with those of the minimum-variance hedges employed in the literature and practice. We quantify the relative performance by the percentage increase or decrease in the expected hedging error variance when the hedger switches from our hedge to the alternative one. Two popular alternatives are the classic static hedge, initially minimizing the hedging error variance and subsequently not readjusting, and the myopic hedge repeating over time the static one with small horizons. These popular hedges are simply driven by the comovement of the stock return and the non-tradable asset payoff over the relevant horizon. Our dynamic hedge inherits this basic structure, but now tracking the comovement between the instantaneous stock return and the non-tradable asset value under our risk-neutral measure, and so additionally capturing the arrival of new information. Consequently, we show that our dynamic hedge typically outperforms the static and myopic ones in plausible intertemporal settings for stock and non-tradable asset dynamics, especially when there is predictability in the non-tradable asset. Only in the special case of random walk processes do the static and myopic hedges coincide with ours. We also compare our hedges with the dynamic hedges considered in the literature that minimize the hedging error variance sitting at an initial date. These hedges, which we refer to as the “pre-commitment” hedges, are generically different from ours since they do not account for the time-inconsistency of the variance criteria and the hedger may deviate from them at later dates unless she can pre-commit to follow them. By definition, a pre-commitment hedge outperforms ours at the initial date. We demonstrate that for a one-year hedging horizon and plausible parameters, it requires less than half a year for our hedge to start outperforming when the stock and the non-tradable asset follow geometric Brownian motions (GBMs).

We generalize our basic framework to the case when the hedger additionally accounts for the mean hedging error, trading it off against the hedging error variance, as commonly considered in the literature under static settings. We also relate this mean-variance hedging to the benchmarking literature in which a money manager’s performance is evaluated relative to that of a benchmark. We show that the dynamic hedge now has an additional speculative

component and additional hedging demands due to the anticipated speculative gains or losses, as in the related literature. We also show that our main baseline results can easily be extended to the case of multiple non-tradable assets and stocks.

The subject of hedging is, of course, prevalent in the literature on derivatives and risk management. Major textbooks, Duffie (1989), Siegel and Siegel (1990), Stulz (2003), Cvitanic and Zapatero (2004), McDonald (2006), Hull (2008), all present the classic static minimum-variance hedging and demonstrate its usefulness for real-life risk management applications. Ederington (1979), Rolfo (1980), Figlewski (1984), Kamara and Siegel (1987), Kerkvliet and Moffett (1991), In and Kim (2006) employ minimum-variance static hedges and evaluate their quality in different empirical applications. Kroner and Sultan (1993), Lioui and Poncet (2000), Brooks, Henry and Persaud (2002) study the performance and economic implications of closely related myopic hedges. In an economy with a static mean-variance hedger, Anderson and Danthine (1980, 1981) study futures hedging and evaluate its impact on production, while Hirshleifer (1988) derives futures risk premia under transaction costs. Roll (1992), Chan, Karceski and Lakonishok (1999), Costa and Paiva (2002), Jorion (2003), Gomez and Zapatero (2003), Cornell and Roll (2005) employ static mean-variance criteria and consider portfolio management with tracking error, deviation from a benchmark, which is just the opposite of hedging error. In the literature above, the hedger either cannot rebalance her portfolio over time or is myopic and looks one period ahead only. This limitation is underscored by Brandt (2003) who demonstrates that when hedging S&P 500 index options under CARA utility, the multi-period hedges can generate substantial welfare gains.

A steadily growing strand of work investigates optimal dynamic hedges consistent with a hedger's utility maximization in typically continuous-time incomplete market settings. Breeden (1984) provides optimal hedging policies with futures in terms of the value function for a general utility function over intertemporal consumption. Stulz (1984) derives explicit optimal hedges with foreign currency forward contracts when the exchange rate follows a GBM and the hedger has logarithmic utility over intertemporal consumption. He

further argues that this hedger behaves like a myopic mean-variance one. Adler and Detemple (1988) consider the hedging of a non-traded cash position for logarithmic utility over terminal wealth and provide an explicit solution in complete markets, and a solution in terms of the value function in incomplete markets. Svensson and Werner (1993), Tepla (2000b) and Henderson (2005) study the optimal hedging of non-tradable income or assets for general utility over intertemporal consumption or terminal wealth. To obtain explicit solutions, these authors specialize to constant relative risk aversion (CARA) preferences, GBM tradable asset prices and an income process following an arithmetic Brownian motion (ABM), while Henderson additionally obtains hedges for GBM and mean-reverting incomes in incomplete and complete markets, respectively. For more general processes or utilities, the solutions in Svensson and Werner and Henderson are typically in terms of value functions, while in Tepla in terms of sensitivities of tradable wealth with respect to asset and state prices. Duffie, Fleming, Soner and Zariphopoulou (1997) and Viceira (2001) consider the hedging of stochastic income with constant relative risk aversion (CRRA) preferences and the tradable asset and income following GBMs and discrete-time lognormal processes, respectively. The former work demonstrates the existence of the solution in a feedback form and derives its asymptotic behavior for large wealth, while the latter work derives a log-linear approximation for the optimal policies in discrete time.

The rapidly growing so-called “mean-variance” hedging literature in dynamic incomplete market settings studies optimal policies based on hedging error means and variances. A large body of literature characterizes these hedges for a quadratic criterion over the hedging error. In the context of futures hedging, Duffie and Richardson (1991) provide explicit optimal hedges that minimize the expected squared error when both the tradable and non-tradable asset prices follow GBMs. Schweizer (1994) and Pham, Rheinlander and Schweizer (1998) in a more general stochastic environment obtain a feedback representation for the optimal policy. Gouieroux, Laurent and Pham (1998) derive hedges in terms of parameters from a specific non-tradable asset payoff decomposition, but are difficult to obtain explicitly. Bertsimas, Kogan and Lo (2001) solve the quadratic hedging problem via dynamic programming

and numerically compute the optimal hedges. Schweizer (2001) provides a comprehensive survey of this literature with further references and notes that finding tractable optimal quadratic hedges is still an open problem. To our best knowledge, with the exception of Duffie and Richardson, there are no works that derive explicit quadratic hedges.

Duffie and Richardson (1991), Schweizer (1994), Musiela and Rutkowski (1998) solve the dynamic minimum-variance hedging problem by reducing it to a quadratic one, thus characterizing the pre-commitment hedges at an initial date from which the hedger may deviate in the future. Duffie and Richardson and Bielecki, Jeanblanc and Rutkowski (2004) also characterize the pre-commitment minimum-variance hedge subject to a constraint on the mean hedging error. This literature, however, lacks explicit results in the case of stochastic mean returns and volatilities, and explicit hedges are only obtained in Duffie and Richardson for GBM asset prices. Duffie and Jackson (1990) derive explicit minimum-variance hedges in futures markets under the special case of martingale futures prices, which makes the hedging problem time-consistent. In the case of mean-variance hedging, by employing backward induction, Anderson and Danthine (1983) obtain hedges in a simple three-period production economy, while Duffie and Jackson (1989) in a two-period binomial model of optimal innovation of futures contracts.

The remainder of the paper is organized as follows. In Section 3.2, we describe the economic setting and determine the dynamically optimal minimum-variance hedges via dynamic programming. We then explicitly compute these hedges in plausible environments with stochastic mean returns and volatilities, and present the time-consistency conditions. In Section 3.3, we compare our dynamically optimal hedge with the pre-commitment, static and myopic hedges, while in Section 3.4, we generalize our baseline model to the case of mean-variance hedging and the case of multiple assets. Section 3.5 concludes. Proofs are in the Appendix.

3.2. Dynamic Minimum-Variance Hedging

3.2.1. Economic Setup

We consider a continuous-time incomplete-market Markovian economy with a finite horizon $[0, T]$. The 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 well-defined and adapted to $\{\mathcal{F}_t, t \in [0, T]\}$, the augmented filtration generated by w and w_X .

An agent in this economy, henceforth the *hedger*, is committed to hold a *non-tradable asset* with payoff X_T at time T . The non-tradable asset can be interpreted in different ways depending on the application. The process X may represent the price of oil, copper or other commodity that the hedger is committed to sell at time T , or may denote the price of a company share that the hedger cannot trade so as to preserve company control. Alternatively, the non-tradable asset may be interpreted as a firm or a project cash flow, the realization of which is defined by the non-tradable state variable X , such as economic conditions, temperature or precipitation level.³⁹ Without loss of generality, we adopt the first interpretation and postulate the price of the non-tradable asset to follow the dynamics

$$\frac{dX_t}{X_t} = m(X_t, t)dt + \nu(X_t, t)dw_{Xt}, \quad (216)$$

where the stochastic mean, m , and volatility, ν , are deterministic functions of X . The risk associated with holding the non-tradable asset can be hedged by continuous trading in two securities, a riskless bond that provides a constant interest rate r and a tradable risky security. Depending on the application, the risky security can be interpreted as a stock, a futures contract or any other derivative security written on the non-tradable asset. Accordingly, the mean and volatility of instantaneous returns on tradable security, which for expositional simplicity we call the *stock*, in general may depend on the non-tradable asset price, X . The dynamics for the stock price, S , is then modeled as

$$\frac{dS_t}{S_t} = \mu(S_t, X_t, t)dt + \sigma(S_t, X_t, t)dw_t, \quad (217)$$

³⁹If the terminal payoff is a non-linear function of some state variable Y , $h(Y_T)$, one can always redefine the state variable to be $X_t = E_t[h(Y_T)]$, so that the terminal payoff is X_T .

where the stochastic mean return, μ , and volatility, σ , are deterministic functions of S and X . We will denote μ_t , σ_t , m_t and ν_t as shorthand for the coefficients in equations (216)–(217).

The hedger chooses a hedging policy, θ , where θ_t denotes the dollar amount invested in the stock at time t , given initial wealth W_0 . The hedger's tradable wealth W then follows the process

$$dW_t = [rW_t + \theta_t(\mu_t - r)] dt + \theta_t \sigma_t dw_t. \quad (218)$$

The market in this economy is incomplete in that it is impossible to hedge perfectly the fluctuations of the non-tradable asset by tradable wealth. Dynamic market completeness obtains only in the special case of perfect correlation between the non-tradable asset and stock returns, $\rho = \pm 1$, in which case the non-tradable asset can be replicated by stock trading and the hedge portfolio uniquely determined by standard no-arbitrage methods. Since perfect hedging is not possible in incomplete markets, the common approach in the literature is to determine a hedging policy according to some criterion that determines the quality of hedging.

The mean-variance hedging literature addresses this for hedging criteria based on the mean and variance of the *hedging error*, $X_T - W_T$. The mean squared error, $E_t(X_T - W_T)^2$, is a commonly employed measure for the quality of hedging from the class of mean-variance criteria (e.g., Duffie and Richardson, 1991; Schweizer, 1994; Gouieroux, Laurent and Pham 1998; Bertsimas, Kogan and Lo, 2001, among others). In general, these quadratic hedges have a complex structure in that they are derived either in a recursive feedback form (e.g., Schweizer, 1994; Pham, Rheinlander and Schweizer, 1998) or depend on parameters from a specific decomposition of the non-tradable asset price X which are difficult to obtain explicitly (e.g., Gouieroux, Laurent and Pham, 1998). Duffie and Richardson provide an explicit quadratic hedge for the special case of both the non-tradable asset and stock prices following GBMs. However, for richer stochastic environments, quadratic hedging has failed to produce tractable, explicit policies.

Another natural criterion for the quality of hedging is the variance of the hedging error, $\text{var}_t[X_T - W_T]$, widely employed in static and myopic settings (analyzed in Sections 3.3.2–3.3.3), as well as dynamic settings (e.g., Duffie and Richardson, 1991; Schweizer, 1994; Musiela and Rutkowski, 1998; Bielecki, Jeanblanc and Rutkowski, 2004, among others). This literature obtains the variance-minimizing policies primarily as a special case of the quadratic hedging problem sitting at an initial date. The time-inconsistency of the variance criterion, however, may induce the hedger to deviate from the initial policy at a later date, as discussed in Section 3.2.4. Moreover, as in the quadratic case, the variance-minimizing policies have not generally been obtained explicitly, with the notable exception being the Duffie and Richardson case of both risky assets following GBMs.

In this paper, we employ the variance-minimizing criterion for the hedger whose problem is

$$\min_{\theta} \text{var}_t[X_T - W_T], \quad (219)$$

subject to the dynamic budget constraint (218). We solve this problem by dynamic programming and hence provide the time-consistent dynamic hedging policy.

3.2.2. Dynamically Optimal Hedging Policy

In this Section, we determine the dynamically optimal minimum variance hedges. The application of dynamic programming, however, is complicated by the fact that the variance criterion is non-linear in the expectation operator and in general not time-consistent. To address these problems, we follow the approach in Chapter 2 developed in the context of dynamic mean-variance portfolio choice and derive a recursive formulation for the hedger's objective function, which yields the appropriate Hamilton-Jacobi-Bellman (HJB) equation of dynamic programming. Proposition 3.1 reports the optimal policy derived from the solution to this equation and the resulting optimal quality of the hedge.

Proposition 3.1. *The optimal hedging policy and the corresponding variance of the hedging*

error are given by

$$\theta_t^* = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial S_t}, \quad (220)$$

$$\text{var}_t[X_T - W_T^*] = (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 X_s^2 \left(\frac{\partial E_s^*[X_T]}{\partial X_s} \right)^2 ds \right], \quad (221)$$

where W_T^* is the terminal tradable wealth under the optimal hedging policy, and $E_t^*[\cdot]$ denotes the expectation under the unique probability measure P^* on which are defined two Brownian motions w_x^* and w^* with correlation ρ such that the processes for the non-tradable asset, X , and stock price, S , are given by

$$\frac{dX_t}{X_t} = \left(m_t - \rho\nu_t \frac{\mu_t - r}{\sigma_t} \right) dt + \nu_t dw_{x_t}^*, \quad \frac{dS_t}{S_t} = rdt + \sigma_t dw_t^*, \quad (222)$$

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

$$\frac{dP^*}{dP} = e^{-\frac{1}{2} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds - \int_0^T \frac{\mu_s - r}{\sigma_s} dw_s}. \quad (223)$$

Proposition 3.1 provides a simple, fully analytical characterization of the optimal hedging policy in terms of the exogenous model parameters and a probability measure P^* (discussed below). We first note that the optimal hedging policy (220) preserves the basic structure of that in complete markets. Indeed, the perfect hedging policy in complete markets (with $\rho = \pm 1$), obtained by standard no-arbitrage methods, is given by

$$\theta_t^* = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^{RN}[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^{RN}[X_T e^{-r(T-t)}]}{\partial S_t}, \quad (224)$$

where $E_t^{RN}[\cdot]$ denotes the expectation under the unique risk-neutral measure and $E_t^{RN}[X_T e^{-r(T-t)}]$ represents the unique no-arbitrage value of the asset payoff X_T . The complete-market dynamic hedge is comprised of the Greeks, given by the sensitivities of the time- t asset value to the non-tradable asset and stock prices (X and S dynamics under the risk-neutral measure are as in (222) with $\rho = \pm 1$). Thus, our dynamic hedge (220) is a simple generalization of the complete-market perfect hedge, with the additional parameter ρ accounting for the market incompleteness and the measure P^* replacing the risk-neutral measure. This is

in stark contrast to the optimal hedging policies obtained in the mean-variance hedging literature which reduce to perfect hedges in complete markets but do not maintain their intuitive structure in incomplete markets. Moreover, as demonstrated in Section 3.2.3, our simple structure allows us to explicitly compute the optimal hedges under various stochastic economic setups.

The probability measure P^* naturally arises in our setting and facilitates much tractability. To highlight the role of this measure, we note the following relation (as derived from Proposition 3.1 in the Appendix) between the expected discounted non-tradable asset payoff, $X_T e^{-r(T-t)}$, under the new and original measures:

$$E_t^*[X_T e^{-r(T-t)}] = E_t[X_T e^{-r(T-t)}] - E_t[W_T^* e^{-r(T-t)} - W_t]. \quad (225)$$

The residual term, $E_t[W_T^* e^{-r(T-t)} - W_t]$, represents the expected discounted gains in tradable wealth that the hedger forgoes in order to hedge the non-tradable asset over the period $[t, T]$, that is, the cost of hedging. So, the right-hand side of (225) represents the expected discounted terminal payoff net of the hedging cost, while the left-hand side the expectation under P^* . In other words, the probability measure P^* incorporates the hedging cost when computing the expected discounted asset payoff. Henceforth, we label P^* as the “hedge-neutral measure” (see Remark 1), and the quantity $E_t^*[X_T e^{-r(T-t)}]$ as the “hedge-neutral value” of the payoff X_T , analogously to the risk-neutral value in the complete-market case. We further note that the hedge-neutral value can also be interpreted as the minimal time- t value of a self-financing minimum-variance hedging portfolio for which the expected hedging error, $E_t[X_T - W_T^*]$, is zero. To demonstrate this interpretation, we observe from (225) that the expected hedging error is zero only if the initial value of the self-financing portfolio equals the expected discounted non-tradable asset payoff under the hedge-neutral measure, that is, $W_t = E_t^*[X_T e^{-r(T-t)}]$. Since the hedge-neutral value is related to the expected hedging error, the hedger guided by the minimum-variance criterion can achieve a better hedging quality by accounting for the sensitivities of the hedge-neutral value. Hence, the hedges are in terms of the hedge-neutral value sensitivities, which we interpret as the delta-hedges, as in the standard analysis of the Greeks.

The quality of the optimal hedge, as measured by the variance of the hedging error (221), also has a simple structure. The hedging error variance is driven by the level of market incompleteness, ρ^2 , and becomes zero in complete markets. Moreover, the quality of the hedge decreases with higher volatility of the non-tradable asset, ν_t , or higher sensitivity of the hedge-neutral value with respect to the asset price, $\partial E_t^*[X_T]/\partial X_t$, since it becomes more difficult to hedge the non-tradable asset.

The optimal hedging policy (220) admits intuitive comparative statics with respect to the model parameters. Assuming for simplicity that the market price of risk, $(\mu_t - r)/\sigma_t$, is driven by the variable X_t only, we see that the total investment in absolute terms, $|\theta_t^*|$, is decreasing in the stock price volatility, σ_t , because higher volatility makes hedging less efficient. The correlation parameter ρ has both a direct and an indirect effect on the magnitude and sign of the hedge. The direct effect implies that the magnitude of the hedge is decreasing in the absolute value of the correlation, $|\rho|$. Intuitively, for higher absolute correlation more wealth is allocated to the stock as the hedge becomes more efficient. This effect is most pronounced in complete markets when $\rho = \pm 1$, and the non-tradable asset can perfectly be hedged. With zero correlation, $\rho = 0$, the direct effect disappears as it becomes impossible to hedge the non-tradable asset. The indirect effect enters via the joint probability distribution of the prices of tradable and non-tradable assets. This latter effect, along with the effects of the non-tradable asset volatility, time horizon and market price of risk, can only be assessed in specific examples for which the optimal hedge can explicitly be computed.

Remark 1 (The hedge-neutral measure). Our hedge-neutral measure P^* is a particular risk-neutral measure, which is not unique in incomplete markets. A similar intuition for P^* with the same label is developed in Chapter 2 in the context of dynamic mean-variance portfolio choice, where this measure is shown to absorb intertemporal hedging demands in such a setting. The measure P^* also turns out to coincide with the so-called “minimal martingale measure” solving $\min_Q E[-\ln(dQ/dP)]$, where dQ/dP denotes the Radon-Nikodym derivative of measure Q with respect to the original measure P . The minimal martingale

measure is argued to arise naturally in the different context of “risk-minimizing hedging,” introduced by Follmer and Sondermann (1986) and Follmer and Schweizer (1991). These works define the cost of hedging as $C_t = W_t - \int_0^t \theta_\tau dS_\tau/S_\tau$ and minimize the risk measure, $E_t[(C_T - C_t)^2]$, with respect to W_τ and θ_τ , for $t \leq \tau \leq T$. In contrast to our work, the resulting hedging policies do not satisfy the budget constraint and require additional zero-mean inflows or outflows to it. As argued by Pham, Rheinlander and Schweizer (1998) in the context of mean-variance hedging a more suitable measure is the “variance-optimal measure” that solves $\min_Q E[(dQ/dP)^2]$. The reason is that in general the optimal policy can be characterized in terms of the variance optimal measure, and only in terms of the minimal martingale measure in the special cases where the two measures coincide under the restrictive conditions of either $\int_0^T (\mu_s - r)/\sigma_s ds$ being deterministic or the stock price, S , not being affected by the state variables.

3.2.3. Applications

In this Section, we demonstrate that in contrast to the extant mean-variance hedging literature, our dynamically optimal minimum-variance hedges can easily be explicitly computed in settings with stochastic means and volatilities. We here interpret the hedging instrument as the stock of a firm that produces the commodity the hedger is committed to hold. It is then plausible that the stock mean return is increasing in the commodity price and the stock volatility decreasing, since the higher commodity price tends to increase the firm cash flows and decrease their risk. Towards this, we consider two examples, each accounting for one of these effects.⁴⁰ In both examples, the non-tradable asset price follows a mean-reverting process, which is consistent with the empirical evidence on oil and other commodity prices. For example, Schwartz (1997) and Schwartz and Smith (2000) provide supporting evidence for Gaussian mean-reverting logarithmic commodity prices, while Dixit and Pindyck (1994) and Pindyck (2004) employ a geometric Ornstein-Uhlenbeck process to model and estimate

⁴⁰A more realistic model would combine the two effects and may include dependence on the state variables that affect both tradable and non-tradable asset prices. In Section 3.4.2 we show that our model can easily be extended to incorporate additional state variables.

oil price dynamics.

In our first example, the non-tradable asset price follows a mean-reverting Ornstein-Uhlenbeck (OU) process:⁴¹

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}dw_{Xt}, \quad (226)$$

with $\lambda > 0$. The stock price has mean returns linear in price X and follows the dynamics considered in Kim and Omberg (1996) in the context of dynamic portfolio choice:

$$\frac{dS_t}{S_t} = (r + \sigma X_t)dt + \sigma dw_t. \quad (227)$$

According to Proposition 3.1, finding the optimal hedging policy amounts to computing the expected non-tradable payoff under the hedge-neutral measure. Since under the hedge-neutral measure the non-tradable asset price, X , also follows an OU process, its first two moments are straightforward to obtain (e.g., Vasicek, 1977). Corollary 3.1 reports the optimal hedging policy and its corresponding quality.

Corollary 3.1. *The optimal hedging policy and the corresponding variance of the hedging error for the mean-reverting Gaussian model (226)–(227) are given by*

$$\theta_t^* = \frac{\rho\bar{\nu}}{\sigma} e^{-(r+\lambda+\rho\bar{\nu})(T-t)}, \quad (228)$$

$$\text{var}_t[X_T - W_T^*] = (1 - \rho^2)\bar{\nu}^2 \frac{1 - e^{-2(\lambda+\rho\bar{\nu})(T-t)}}{2(\lambda + \rho\bar{\nu})}. \quad (229)$$

The optimal hedge is a simple generalization of the complete-market perfect hedge, with $\rho\bar{\nu}$ replacing $\bar{\nu}$ in complete markets to account for the imperfect correlation between the stock and non-tradable asset. This explicit solution also yields further insights that cannot be analyzed in the general framework of Section 3.2.2. In particular, Corollary 3.1 reveals that the sign of the hedge is given by that of the correlation parameter, ρ . When the non-tradable asset and stock prices are positively correlated, only a long position in the stock

⁴¹The OU process allows considerable tractability at the cost of possibly negative prices. Alternatively, the hedging strategies can explicitly be derived in a model with the stock mean return driven by a mean-reverting logarithmic non-tradable asset price, as in Schwartz (1997) and Schwartz and Smith (2000). In this case all prices would remain positive.

can reduce the hedging error variance, and vice versa for negative correlation. Moreover, the absolute value of the hedge and the variance of the hedging error are decreasing in the speed of mean-reversion parameter, λ . This is intuitive since a higher speed of convergence to the mean leads to a lower variance of the non-tradable asset payoff, and hence a smaller hedge. The hedging quality also improves as the degree of market completeness, captured by ρ^2 , increases. Moreover, the hedging quality is higher for a positive correlation than for a negative one of the same magnitude since positively correlated stock better tracks the non-tradable asset price.

The second example considers the case of the stock volatility being decreasing in the non-tradable asset price, which follows a square-root mean-reverting process

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}\sqrt{X_t}dw_{Xt}, \quad (230)$$

with $\lambda > 0$. The stock price follows the stochastic-volatility model employed by Chacko and Viceira (2005) in the context of portfolio choice:

$$\frac{dS_t}{S_t} = \mu dt + \sqrt{\frac{1}{X_t}}dw_t. \quad (231)$$

As in the previous example, the explicit hedge follows easily from Proposition 3.1. Corollary 3.2 presents the optimal hedge along with the associated variance of the hedging error.

Corollary 3.2. *The optimal hedging policy and the corresponding variance of the hedging error for the mean-reverting stochastic-volatility model (230)–(231) are given by*

$$\begin{aligned} \theta_t^* &= \rho\bar{\nu}X_t e^{-(r+\lambda+\rho\bar{\nu}(\mu-r))(T-t)}, \quad (232) \\ \text{var}_t[X_T - W_T^*] &= (1 - \rho^2)\bar{\nu}^2\bar{X} \frac{1 - e^{-2(\lambda+\rho\bar{\nu}(\mu-r))(T-t)}}{2(\lambda + \rho\bar{\nu}(\mu - r))} \\ &+ (1 - \rho^2)\bar{\nu}^2(X_t - \bar{X}) \frac{e^{-\lambda(T-t)} - e^{-2(\lambda+\rho\bar{\nu}(\mu-r))(T-t)}}{\lambda + 2\rho\bar{\nu}(\mu - r)}. \quad (233) \end{aligned}$$

Corollary 3.2 reveals that the absolute value of the hedge is increasing in the non-tradable asset price. This is because a high asset price implies a low stock volatility. Hence, a higher stock holding is required to hedge the non-tradable asset. The sign of the optimal hedge

equals that of the correlation ρ and its absolute value is decreasing in the mean-reversion parameter λ . For the same reason as in the previous example, the hedging quality improves with increased mean-reversion or degree of market completeness.

3.2.4. Time-Consistency Conditions

We here discuss the time-inconsistency of the variance minimization criterion and establish conditions on the economy, albeit restrictive, under which time-consistency obtains. First, we observe that by the law of total variance

$$\text{var}_t[X_T - W_T] = E_t[\text{var}_{t+\tau}(X_T - W_T)] + \text{var}_t[E_{t+\tau}(X_T - W_T)], \quad \tau > 0. \quad (234)$$

Sitting at time t , the hedger minimizes the sum of the expected future $(t + \tau)$ -variance of hedging error and the variance of its future expectation, both of which may depend on future strategies. When the hedger arrives at the future time $t + \tau$, however, she minimizes just the variance at that time, and regrets having taken into account the second term in (234), the time- t variance of future expectation, since it vanishes at time $t + \tau$, and hence the time-inconsistency.

The time-inconsistency issue disappears in complete markets ($\rho = \pm 1$), where the non-tradable asset can perfectly be replicated by dynamic trading, leading to zero hedging error variance. However, it is still possible to have time-consistency of the variance criterion in an incomplete-market economy under certain restrictions, as summarized in Proposition 3.2.

Proposition 3.2. *Assume that the stock risk premium is zero, $\mu_t - r = 0$. Then the variance criterion (219) is time-consistent and the ensuing optimal dynamic minimum-variance hedging policy is given by*

$$\theta_t^* = \frac{\rho \nu_t}{\sigma_t} X_t \frac{\partial E_t[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t[X_T e^{-r(T-t)}]}{\partial S_t}. \quad (235)$$

In an economy with no compensation for risk taking and where the stock is traded only for hedging purposes, the variance criterion becomes time-consistent. The reason is that

with zero stock risk premium, the (discounted) tradable wealth reduces to a martingale and so the hedging costs (second term in (225)) disappear. Consequently, the non-tradable asset, and hence time- t hedge, are not affected by future policies, eliminating the time-inconsistency.⁴² Moreover, we see that the structure of the optimal hedge is as in complete and incomplete markets, but now the original measure acts as the valuating expectation. This optimal hedge generates those obtained by Duffie and Jackson (1990), who consider among other problems, minimum-variance hedging with futures contracts which turns out to be time-consistent. As in Proposition 3.2, it can be shown for their economic setting with martingale futures prices and interest accruing on a futures margin account that the variance criterion is time-consistent and the optimal hedge is given by (235), which generalizes their explicit hedges derived for martingale and geometric Brownian motion non-tradable asset prices.

Proposition 3.2 also allows us to convert the minimum-variance hedging problem considered in Section 3.2.2 to a time-consistent one, as discussed in Corollary 3.3.

Corollary 3.3. *In our incomplete-market economy consider the class of risk-neutral probability measures, P^η , parameterized by η , on which are defined two Brownian motions w_X^η and w^η with correlation ρ such that the processes for the non-tradable asset, X , and stock price, S , are given by*

$$\frac{dX_t}{X_t} = \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} - \sqrt{1 - \rho^2} \eta_t \right) dt + \nu_t dw_{X_t}^\eta, \quad \frac{dS_t}{S_t} = r dt + \sigma_t dw_t^\eta, \quad (236)$$

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

$$\frac{dP^\eta}{dP} = e^{-\frac{1}{2} \int_0^T \left(\left(\frac{\mu_s - r}{\sigma_s} \right)^2 + \eta_s^2 \right) ds - \int_0^T \frac{\mu_s - r}{\sigma_s} dw_s - \int_0^T \eta_s dw_s^\perp}, \quad (237)$$

where w^\perp is a Brownian motion uncorrelated with w and defined by $dw_t^\perp \equiv (dw_{X_t} - \rho dw_t) / \sqrt{1 - \rho^2}$.

The following minimum-variance criteria

$$\text{var}_t^\eta [X_T - W_T], \quad (238)$$

⁴²Formally, the first term in the law of total variance (234) depends only on future policies, while the second term depends only on the time- t policy, θ_t . As a result, the minimization of time- t variance does not lead to any inconsistency.

where the variance is taken under a risk-neutral measure P^η , are time-consistent with the optimal hedge given by

$$\theta_t^\eta = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^\eta[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^\eta[X_T e^{-r(T-t)}]}{\partial S_t}, \quad (239)$$

where $E_t^\eta[\cdot]$ denotes the expectation under P^η . For $\eta = 0$, a risk-neutral measure is hedge-neutral and the optimal hedge (239) equals the dynamically optimal hedge (220).

Corollary 3.3 reveals that a risk-neutral measure adjusts the variance criterion so that it becomes time-consistent. The criterion (238) treats the non-tradable asset and stock price processes as if they were under a risk-neutral measure. Under this measure the stock has mean return equal to the riskless rate r , and hence zero risk premium, which implies time-consistency by Proposition 3.2. The dynamically optimal hedge (220) is then obtained from the time-consistent hedging problem when $\eta_t = 0$.

3.3. Comparison with Pre-commitment, Static and Myopic Hedges

In this Section, we compare our dynamically optimal hedging policy with popular minimum-variance hedging policies employed in the finance literature and practice. First, we consider the policy that minimizes the hedging error variance at an initial date. Second, we look at the classic static hedge that minimizes the hedging error variance at an initial date and does not allow subsequent portfolio rebalancing. Finally, we study the popular myopic hedge that in each period hedges the changes in the non-tradable asset price over the next period.

To assess the relative performance of any given two policies, we compare their hedging error variances. Since the conditional variances in general are stochastic, for tractability we consider a relative performance measure that computes the percentage increase or decrease in the unconditional expected variance when the hedger switches from the dynamically optimal to an alternative hedging policy:

$$\Delta_t = \frac{E_0[\text{var}_t(X_T - W_T^{\text{alternative}})]}{E_0[\text{var}_t(X_T - W_T^*)]} - 1, \quad (240)$$

where $W_T^{\text{alternative}}$ denotes the terminal tradable wealth under the alternative policy, $\theta_t^{\text{alternative}}$,

considered in Sections 3.3.1–3.3.3.⁴³ A positive relative performance measure implies that the quality of the dynamically optimal hedge is on average higher than that of the alternative hedge, in which case we say that the dynamically optimal hedge outperforms the alternative one.

3.3.1. Comparison with Pre-commitment Policy

We here investigate the performance of the dynamically optimal hedging policy as compared with that of the policy that minimizes the hedging error variance at an initial date 0, as considered in the literature (e.g., Duffie and Richardson, 1991; Schweizer, 1994; Musiela and Rutkowski, 1998). As discussed in Section 3.2, the variance-minimizing hedger may find it optimal to deviate from the latter policy at future dates unless she can pre-commit to follow it, and henceforth we refer to it as the *pre-commitment* policy.

To our best knowledge, Duffie and Richardson are the only ones to provide an explicit expression for this policy in the context of hedging with futures contracts and interest accruing on a futures margin account when the futures and non-tradable asset prices follow GBMs. Therefore, we compare the two policies for the case of the non-tradable asset and stock prices following GBMs:

$$\frac{dX_t}{X_t} = mdt + \nu dw_{xt}, \quad (241)$$

$$\frac{dS_t}{S_t} = \mu dt + \sigma dw_t, \quad (242)$$

where m , ν , μ and σ are constants. The dynamically optimal hedge is obtained from Proposition 3.1, while the pre-commitment hedge along the lines of Duffie and Richardson (1991) adapted to our setup. Proposition 3.3 presents the two policies and a simple condition for the dynamically optimal hedge to outperform.⁴⁴

Proposition 3.3. *The dynamically optimal and pre-commitment policies for GBM non-*

⁴³All our relative performance results in Sections 3.3.2–3.3.3 (Propositions 3.4–3.5) remain valid for a more general, conditional relative performance measure given by $\text{var}_t[X_T - W_T^{\text{alternative}}] / \text{var}_t[X_T - W_T^*] - 1$.

⁴⁴Proposition 3.3 does not report the variances of hedging errors under the two policies since this Section focuses on relative rather than individual performances. These variances, however, can be deduced in the proof of Proposition 3.3 in the Appendix.

tradable asset and stock prices (241)–(242) are given by

$$\theta_t^* = \frac{\rho\nu}{\sigma} X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)}, \quad (243)$$

$$\begin{aligned} \theta_t^{commit} &= \frac{\rho\nu}{\sigma} X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} \\ &- \frac{\mu-r}{\sigma^2} \left((X_0 e^{(m-r-\rho\nu(\mu-r)/\sigma)T} - W_0) e^{rt} - (X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} - W_t^{commit}) \right). \end{aligned} \quad (244)$$

Furthermore, $\exists \bar{t} < T$ such that the relative performance $\Delta_t > 0$ for $t > \bar{t}$, i.e., the dynamically optimal hedge outperforms the pre-commitment hedge after a period of time.

Proposition 3.3 reveals that the dynamically optimal and pre-commitment hedges are generically different. While the dynamically optimal hedge is a simple generalization of the complete-market hedge (with ρ substituted in), the pre-commitment hedge inherits an additional stochastic term (second term in (244)). To see why this term arises, we observe that if the hedger follows the dynamically optimal policy from time t on, her expected hedging error is (as shown in the Appendix)

$$E_t[(X_T - W_T^*)e^{-r(T-t)}] = X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} - W_t. \quad (245)$$

Hence, the second term in (244) hedges the deviations of the expected discounted hedging error, $E_t[(X_T - W_T^*)e^{-r(T-t)}]$, from its time-zero value (compounded by a term reflecting accrued interest in $[0, t]$). The hedger tries to keep this deviation close to zero because a high variability in the expected hedging error implies a high time-zero hedging error variance (from the law of total variance (234)). So, when the second term in (244) is positive, the hedger reduces her stock holding, and hence her anticipated tradable wealth, thereby increasing the future expected hedging error making it closer to time-zero hedging error; and vice versa when the second term is negative. The structure of the pre-commitment policy highlights the time-inconsistency of the problem. It shows that sitting at time t , the hedger still behaves so as to maintain a low time-zero rather than time- t hedging error variance.

Proposition 3.3 also reveals that the dynamically optimal hedge always outperforms the pre-commitment one after a certain period of time. Since the pre-commitment hedge minimizes the time-zero hedging error variance, it performs better for small time periods

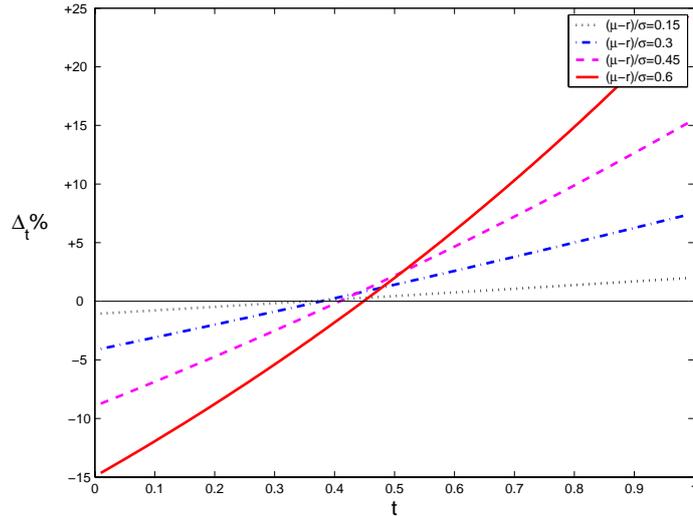


Figure 3.1: Relative Performance of Dynamically Optimal and Pre-commitment Hedges.

The figure plots the relative performance measure Δ_t (equation (240)) as a function of time for varying levels of market price of risk. The dynamically optimal policy outperforms the pre-commitment one whenever $\Delta_t > 0$. Correlation and horizon parameters are assumed to be $\rho = 0.5$ and $T = 1$. The volatility parameter $\nu = 0.36$ is taken from the estimate in Schwartz (1997), based on weekly oil futures price data in January 1990–1995, while market price of risk range of $[0.15, 0.6]$ is consistent with the estimates in Mehra and Prescott (1985), Cogley and Sargent (2008), and others.

t . However, at later dates, the dynamically optimal hedge performs better since the time-inconsistency makes the pre-commitment hedge suboptimal. In the case when the non-tradable asset and stock prices follow GBMs, the relative performance measure, Δ_t , can explicitly be computed (as reported in the proof of Proposition 3.3). Conveniently, this measure depends only on the correlation parameter ρ , non-tradable asset volatility ν , market price of risk $(\mu - r)/\sigma$, and the hedging horizon $T - t$. Since the relative performance measure turns out to not be sensitive to ρ and ν , we focus below on its behavior with respect to $T - t$ and $(\mu - r)/\sigma$.

We consider a specific example in which the non-tradable asset is oil and the stock represents the stock of an oil producing company. The GBM model for oil prices is a special case of those in Gibson and Schwartz (1990) and Schwartz (1997) when the convenience yield is assumed constant rather than mean-reverting. Figure 3.1 plots the relative performance measure Δ_t over time for different market prices of risk. It demonstrates that for plausible

parameters and one-year hedging horizon, the dynamically optimal policy starts outperforming the pre-commitment policy from mid-year on. Moreover, for lower market price of risk, the relative performance measure gets closer to zero, reflecting the fact that the difference between the two policies is reduced (as observed from the expressions (243)–(244)).

3.3.2. Comparison with Static Policy

We now examine the classic *static* hedging problem in finance where an initial hedge, chosen to minimize the hedging error variance, is not readjusted throughout the hedging period. Due to its tractability and intuitive appeal, the static minimum-variance hedge is widely used by both practitioners and academics. The classic theory of the static hedge and its real-life applications are presented in all the prominent textbooks in derivatives and risk-management (e.g., Duffie, 1989; Siegel and Siegel, 1990; Stultz, 2003; Cvitanic and Zapatero, 2004; McDonald, 2006; Hull, 2008), as well as being employed in empirical works (e.g., Ederington, 1979; Rolfo, 1980; Figlewski, 1984; Kamara and Siegel, 1987; Kerkvliet and Moffett, 1991; In and Kim, 2006). As discussed in Section 3.4.1, a generalization of the static hedge to static mean-variance hedge (incorporating additionally the mean in the hedging criterion) is also widely employed in the literature.

A static hedger minimizes the variance of the hedging error at the initial date 0, subject to the static budget constraint

$$W_T = W_0 e^{rT} + \theta_0 (S_T/S_0 - e^{rT}), \quad (246)$$

and holds the initially chosen hedge, θ_0^{static}/S_0 , in units of stock, throughout the hedging horizon. The solution to this problem is easily obtained and the time- t static hedge is given by⁴⁵

$$\theta_t^{static} = \frac{\text{cov}_0(S_T/S_t, X_T)}{\text{var}_0[S_T/S_t]}. \quad (247)$$

The static hedge is simply driven by the comovement of the stock return and the non-tradable asset payoff over the remaining hedging period. The hedge is positive when the

⁴⁵Since the hedger holds the same number of units of stock over the horizon, $\theta_t^{static} = (\theta_0^{static}/S_0)S_t$.

stock is positively correlated with the asset payoff since then the stock better tracks the asset payoff over the period. We observe that our optimal dynamic hedge (220) can equivalently be rewritten as

$$\theta_t^* = \frac{\text{cov}_t(dS_t/S_t, dE_t^*[X_T e^{-r(T-t)}])}{\sigma_t^2}. \quad (248)$$

Clearly, the dynamic hedge inherits the basic intuitive structure of the static hedge, but now tracking comovement between the instantaneous stock return and the change in the hedge-neutral asset payoff value, and so capturing the arrival of new information. The dynamic hedge is positive for positive correlation between the stock and hedge-neutral value since then the stock trading better replicates the non-tradable payoff.

One important difference between the static and dynamically optimal hedges is that the static hedge in general does not provide a perfect hedge, even in dynamically complete markets when $\rho^2 = 1$ (with one notable exception as discussed below), in contrast to the dynamic one. This is because the static hedge cannot adjust to the arrival of new information as it does not rebalance the initially chosen policy. Consequently, the dynamic hedge always outperforms the static one when the market is close to being complete. We next compare the two hedges and their performances under popular price dynamics for which the relative performance measure Δ_t (expression (240)) can explicitly be computed. In addition to considering the non-tradable and stock prices following GBMs (241)–(242), we also study the cases of their following ABMs

$$dX_t = \bar{m}dt + \bar{\nu}dw_{xt}, \quad (249)$$

$$dS_t = \bar{\mu}dt + \bar{\sigma}dw_t, \quad (250)$$

with \bar{m} , $\bar{\nu}$, $\bar{\mu}$, $\bar{\sigma}$ constants, as well as the non-tradable asset following an OU process

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}dw_{xt}, \quad (251)$$

with \bar{X} and $\lambda > 0$ constants. Proposition 3.4 reports the results.

Proposition 3.4. *The dynamically optimal and static hedges and their relative performances under various non-tradable asset and stock price processes are as given in Table 3.1.*

Table 3.1
Optimal Dynamic and Static Hedges and Their Relative Performances

The table reports the dynamically optimal and static hedges and the sign of their relative performance measure Δ_t (equation (240)) when non-tradable asset and stock prices follow various stochastic processes. We say that the dynamically optimal hedge outperforms the static one when Δ_t is positive, and underperforms when Δ_t is negative. ABM, GBM and OU denote arithmetic Brownian motion (equations (249)–(250)), geometric Brownian motion (equations (241)–(242)) and Ornstein-Uhlenbeck mean-reverting (equation (251)) processes, respectively. In all cases, we assume $\rho \neq 0$, since otherwise the stock cannot hedge the non-traded asset and all the hedges are trivially zero.

Optimal Hedges		Performance	Processes	
dynamic θ_t^*	static θ_t^{static}	sign Δ_t	asset X	stock S
$\frac{\rho\bar{v}S_t}{\sigma}$	$\frac{\rho\bar{v}S_t}{\sigma}$	0	ABM	ABM
$\frac{\rho\nu X_t}{\sigma} e^{(m-r-\rho\nu\frac{\mu-r}{\sigma})(T-t)}$	$\frac{X_0 S_t}{S_0} \frac{e^{\rho\nu\sigma T} - 1}{e^{\sigma^2 T} - 1} e^{(m-\mu)T}$	$\left\{ \begin{array}{l} + \quad \rho > 0 \\ +/- \quad \rho < 0 \end{array} \right.$	GBM	GBM
$\frac{\rho\bar{v}S_t}{\sigma} \frac{\lambda e^{-\lambda(T-t)} + r e^{r(T-t)}}{\lambda + r}$	$\frac{\rho\bar{v}S_t}{\sigma} \frac{1 - e^{-\lambda T}}{\lambda T}$	+	OU	ABM
$\frac{\rho\bar{v}}{\sigma} e^{-(\lambda+r)(T-t)}$	$\frac{\rho\bar{v}\sigma S_t}{\lambda S_0} \frac{1 - e^{-\lambda T}}{e^{\sigma^2 T} - 1} e^{-\mu T}$	+	OU	GBM

The dynamically optimal and static hedges coincide in the special case of the non-tradable asset and stock prices both following ABMs. This is because, with random walk prices, the non-tradable asset and stock price variances and covariances are deterministic and hence the new information released over time does not help predict them better than the information available at the initial date, and therefore the hedging problem is effectively static by its nature. The two policies are considerably different, however, in the other settings where the dynamic hedge typically outperforms the static one. In particular, when the non-tradable asset and stock prices follow GBMs, the dynamically optimal policy outperforms the static one when the correlation parameter ρ is positive. With a positive asset-stock correlation, the stock process better imitates the fluctuations in the non-tradable asset price, which improves the quality of hedging. When the correlation is negative, the dynamically optimal hedge always outperforms after a certain period of time (as demonstrated in the Appendix) but may underperform in the beginning if the stock market price of risk is implausibly high. Finally, when the non-tradable asset price follows an OU process, the dynamic hedge always outperforms. With the predictability in the non-tradable asset price

present, the dynamic hedge better accounts for the arrival of new information over time, and hence performs better.

3.3.3. Comparison with Myopic Policy

Finally, we compare the dynamically optimal and myopic hedges. At each point in time, a *myopic* hedger hedges the instantaneous changes in the non-tradable asset price via the instantaneous changes in tradable wealth. Hence, the myopic hedge can be viewed as the static hedge over an infinitesimally small hedging horizon, repeated over time. The myopic hedge retains the tractability of the static hedge which makes it appealing for practitioners and academics (e.g., Kroner and Sultan, 1993; Lioui and Poncet, 2000; Brooks, Henry and Persaud, 2002).

The myopic hedger at each point of time minimizes the variance of the hedging error over the next instant

$$\min_{\theta_t} \text{var}_t[dX_t - dW_t], \quad (252)$$

subject to the budget constraint (218). The variance of this instantaneous hedging error can explicitly be computed to be given by a quadratic function of a hedging policy. The minimization of this variance leads to the following explicit expression for the optimal myopic policy:

$$\theta_t^{myopic} = \frac{\rho \nu_t}{\sigma_t} X_t. \quad (253)$$

The myopic hedge is simply the instantaneous version of the static hedge and is in general different from the dynamically optimal hedge (220). In particular, the myopic hedge ignores the potential impact of mean-returns on the hedging error variance since the first term in the asset dynamics (216) is conditionally riskless over next instant. As a result, the myopic policy in general does not provide a perfect hedge even in dynamically complete markets, just like the static one. Consequently, it underperforms the dynamically optimal hedge when the market is close to being complete. As in the previous Section, we compare the myopic and dynamically optimal hedges for popular price processes, including GBMs (241)–(242),

Table 3.2
Dynamically Optimal and Myopic Policies and Their Relative Performances

The table reports the dynamically optimal and myopic hedges and the sign of their relative performance measure Δ_t (equation (240)) when non-tradable asset and stock prices follow various stochastic processes. We say that the dynamically optimal hedge outperforms the myopic one when Δ_t is positive, and underperforms when Δ_t is negative. ABM, GBM and OU denote arithmetic Brownian motion (equations (249)–(250)), geometric Brownian motion (equations (241)–(242)) and Ornstein-Uhlenbeck mean-reverting (equation (251)) processes, respectively. In all cases, we assume $\rho \neq 0$, since otherwise the stock cannot hedge the non-traded asset and all the hedges are trivially zero.

Optimal Hedges		Performance	Processes	
dynamic θ_t^*	myopic θ_t^{myopic}	sign Δ_t	asset X	stock S
$\frac{\rho\bar{v}S_t}{\sigma}$	$\frac{\rho\bar{v}S_t}{\sigma}$	0	ABM	ABM
$\frac{\rho\nu X_t}{\sigma} e^{(m-r-\rho\nu\frac{\mu-r}{\sigma})(T-t)}$	$\frac{\rho\nu X_t}{\sigma}$	$\left\{ \begin{array}{l} 0 \quad \frac{m-r}{\nu} = \rho\frac{\mu-r}{\sigma} \\ + \quad \rho > 0, \quad \frac{m-r}{\nu} > \rho\frac{\mu-r}{\sigma} \\ +/- \quad \rho < 0 \text{ or } \frac{m-r}{\nu} < \rho\frac{\mu-r}{\sigma} \end{array} \right.$	GBM	GBM
$\frac{\rho\bar{v}S_t}{\sigma} \frac{\lambda e^{-\lambda(T-t)} + r e^{r(T-t)}}{\lambda+r}$	$\frac{\rho\bar{v}S_t}{\sigma}$	+	OU	ABM
$\frac{\rho\bar{v}}{\sigma} e^{-(\lambda+r)(T-t)}$	$\frac{\rho\bar{v}}{\sigma}$	+	OU	GBM

ABMs (249)–(250), and OU (251). Proposition 3.5 reports the two hedges under these settings, as well as their relative performances.

Proposition 3.5. *The dynamically optimal and myopic hedges and their relative performances under various non-tradable asset and stock price processes are as given in Table 3.2.*

The myopic and dynamically optimal hedges coincide under the random walk environment of ABM since the hedging problem is effectively static by its nature, as discussed in Section 3.3.2. In the other environments, the two hedges generically differ, with the dynamically optimal hedge typically outperforming the myopic hedge. With predictable OU non-tradable asset prices, the dynamically optimal policy better incorporates the arrival of new information and hence outperforms the myopic one, as in static case of Section 3.3.2. When the asset and stock both follow GBMs, the two hedges coincide in the very special case of the non-tradable asset market price of risk equalling that of the stock (adjusted by

correlation ρ). The reason is that in this case the tradable wealth better tracks the non-tradable asset price since the myopic hedge not only minimizes the instantaneous hedging error variance but also matches the risk premia on the non-tradable asset and tradable wealth.⁴⁶ The dynamically optimal hedge, however, outperforms for positive asset-stock correlation and relatively high asset market price of risk, and otherwise can outperform or underperform. As an example, consider the case of hedging gas prices that follow GBM with parameters $m = 0.836$ and $\nu = 0.59$ (approximated from OU gas log-prices estimated in Jalliet, Ronn and Tompaidis (2004)). In this case, the dynamically optimal policy outperforms the myopic one for positive correlation and plausible stock market prices of risk of $[0.15, 0.6]$.

3.4. Extensions

We now generalize the results on minimum-variance hedging derived in Section 3.2 along two dimensions. First, we consider a more general model in which the hedger is guided by a linear mean-variance criterion over the hedging error. Second, we demonstrate that the minimum-variance hedging model can easily be extended to a richer environment with multiple non-tradable assets and stocks.

3.4.1. Mean-Variance Hedging and Benchmark Tracking

We here consider a hedger who also accounts for the mean hedging error, and trades it off against the hedging error variance. Such a mean-variance hedging criterion is commonly employed in a variety of, primarily static, settings (e.g., Anderson and Danthine, 1980, 1981, 1983; Hirshleifer, 1988; Duffie, 1989; Duffie and Jackson, 1989). Our analysis in this Section is also related to the literature on portfolio management with benchmarking. In this literature, money managers are evaluated relative to a benchmark portfolio and are concerned about their tracking error, defined as the deviation of a manager's performance from that of the benchmark. The mean-variance tracking error model amounts to mean-variance hedging

⁴⁶More generally, it can be shown that the dynamically optimal and myopic hedges coincide whenever $(m_t - r)/\nu_t = \rho(\mu_t - r)/\sigma_t$. The intuition is the same as in the case of GBM asset and stock prices.

if we relabel the non-tradable asset X as the benchmark portfolio and observe that tracking error is the negative of hedging error. Roll (1992), Jorion (2003), Gomez and Zapatero (2003), Cornell and Roll (2005) discuss the implications of such benchmarking on portfolio efficiency and asset pricing. Chan, Karceski and Lakonishok (1999) and Costa and Paiva (2002) discuss the implications of estimation risk and robust portfolio selection with benchmarking. These works all employ a static mean-variance framework by either minimizing the tracking error variance for a given mean, or maximizing the tracking error mean for a given variance.

A dynamic mean-variance hedger chooses an optimal hedge, trading-off lower variance against higher mean of hedging error, by solving the dynamic problem

$$\max_{\theta_t} E_t[X_T - W_T] - \frac{\gamma}{2} \text{var}_t[X_T - W_T], \quad (254)$$

subject to the budget constraint (218), where the parameter γ captures the hedger's attitudes towards risk. The optimal quality of the hedge is measured by the value function J_t , given by the criterion in (254) evaluated at the optimal policy. As in Section 3.2 we consider the time-consistent solution to problem (254) obtained by dynamic programming. Proposition 3.6 reports the dynamically optimal hedging policy along with the value function.

Proposition 3.6. *The dynamically optimal mean-variance hedge, θ_t^* , and the corresponding value function, J_t , are given by*

$$\begin{aligned} \theta_t^* &= \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial S_t} - \frac{\mu_t - r}{\gamma\sigma_t^2} e^{-r(T-t)} \\ &+ \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*\left[\left(\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds\right) e^{-r(T-t)}\right]}{\partial X_t} + S_t \frac{\partial E_t^*\left[\left(\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds\right) e^{-r(T-t)}\right]}{\partial S_t}, \end{aligned} \quad (255)$$

$$\begin{aligned} J_t &= -\frac{\gamma}{2} (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 X_s^2 \left(\frac{\partial E_s^*[X_T + \int_s^T \frac{1}{\gamma} \left(\frac{\mu_\tau - r}{\sigma_\tau}\right)^2 d\tau]}{\partial X_s} \right)^2 ds \right] \\ &+ E_t^*[X_T e^{-r(T-t)}] - W_t e^{r(T-t)} + \frac{1}{2} E_t^* \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds \right]. \end{aligned} \quad (256)$$

Proposition 3.6 reveals that the dynamically optimal mean-variance hedge is comprised of three types of terms. The first two terms in (255) comprise the variance-minimizing

hedge of Section 3.2, reflecting the hedger's aversion towards hedging error variance. The third term is the speculative demand, as referred to in the related works (e.g., Anderson and Danthine, 1980, 1981; Duffie, 1989), and arises due to the hedger's desire for high mean hedging error. Finally, the last two terms in (255) are the intertemporal hedging demands, familiar in the portfolio choice literature. These demands arise due to the fluctuations in the non-tradable asset and stock mean returns and volatilities, and in our framework are simply given by the sensitivities of the hedge-neutral value of anticipated speculative gains.

The optimal hedge (255) can explicitly be computed for specific stochastic environments, as in the case of the minimum-variance hedge. However, in this case, the computations are more involved, and the hedge depends on the hedger-specific parameter γ . Moreover, in contrast to the minimum-variance hedge, the dynamically optimal mean-variance hedge, in general, differs from its associated pre-commitment one even in complete markets. Furthermore, even though the hedging problem can be reduced to one with a time-consistent criterion under some conditions as in Section 3.2.4, the solution from such a criterion does not, in general, coincide with the actual one (255), unlike in the minimum-variance case. The value function (256) that measures the quality of the optimal hedge implies a better hedge with a higher value. However, it can be verified that unlike the minimum-variance hedge, the optimal mean-variance hedge does not provide a perfect hedge (i.e., having zero hedging error variance) even in complete markets because the hedger forgoes lower hedging error variance for higher mean.

3.4.2. Multi-dimensional case

We now demonstrate that the results of Section 3.2 can be extended to the case with multiple non-tradable assets and stocks. We consider an economy in which uncertainty is generated by two multi-dimensional Brownian motions $w_X = (w_{X1}, \dots, w_{XN})^\top$ and $w = (w_1, \dots, w_K)^\top$. By ρ we denote the $N \times K$ correlation matrix with elements $\rho = \{\rho_{nk}\}$ representing correlations between the Brownian motions w_{Xn} and w_k .

There are N non-tradable assets whose prices, $X = (X_1, \dots, X_N)^\top$, follow dynamics

$$\frac{dX_{it}}{X_{it}} = m_i(X_t, t)dt + \nu_i(X_t, t)^\top dw_{Xt}, \quad i = 1, \dots, N, \quad (257)$$

where m_i and ν_i are deterministic functions of X . We let $m = (m_1, \dots, m_N)^\top$ and $\nu = (\nu_1, \dots, \nu_N)^\top$ denote the vector of mean returns and the volatility matrix whose elements $\nu = \{\nu_{ni}\}$ represent covariances between the non-tradable asset returns and Brownian motion w_X . At future date T , the hedger is committed to hold a portfolio of non-tradable assets with payoff $\phi^\top X_T$, where $\phi = (\phi_1, \dots, \phi_N)^\top$ denotes the vector of units held in assets. An asset that is not held by the hedger ($\phi_i = 0$) may still affect the dynamics of the assets held and can be relabeled to be a state variable, such as economic conditions, temperature or precipitation level.

The risk associated with the portfolio of non-tradable assets can be hedged by trading in a riskless bond with constant interest rate r and K tradable securities with prices $S = (S_1, \dots, S_K)^\top$ that follow the dynamics

$$\frac{dS_{jt}}{S_{jt}} = \mu_j(S_t, X_t, t)dt + \sigma_j(S_t, X_t, t)^\top dw_t, \quad j = 1, \dots, K, \quad (258)$$

where μ_i and σ_i are deterministic functions of S and we let $\mu = (\mu_1, \dots, \mu_K)^\top$ and $\sigma = (\sigma_1, \dots, \sigma_K)^\top$ denote the vector of mean returns and the volatility matrix of stock returns, assumed invertible, respectively. The hedger chooses a hedging policy, $\theta = (\theta_1, \dots, \theta_K)$, where θ_t denotes the vector of dollar amounts invested in stocks at time t . The tradable wealth W then follows the process

$$dW_t = [rW_t + \theta_t^\top (\mu_t - r)]dt + \theta_t^\top \sigma_t dw_t. \quad (259)$$

The hedger's dynamic optimization problem is as in Section 3.2. At each time t , she minimizes the variance of her hedging error, $\phi^\top X_T - W_T$, subject to the budget constraint (259). The optimal policy is then derived by dynamic programming as in Section 3.2. Proposition 3.7 reports the dynamically optimal hedge and its associated quality.

Proposition 3.7. *The optimal hedging policy and the corresponding variance of hedging*

error are given by

$$\theta_t^* = (\nu_t \rho^\top \sigma_t^{-1})^\top I_{X_t} \frac{\partial E_t^*[\phi^\top X_T e^{-r(T-t)}]}{\partial X_t^\top} + I_{S_t} \frac{\partial E_t^*[\phi^\top X_T e^{-r(T-t)}]}{\partial S_t^\top} \quad (260)$$

$$\text{var}_t[\phi^\top X_T - W_T^*] = E_t \left[\int_t^T \frac{\partial E_s^*[\phi^\top X_T]}{\partial X_s} \nu_s I_{X_s} (I - \rho^\top \rho) I_{X_s} \nu_s^\top \frac{\partial E_s^*[\phi^\top X_T]}{\partial X_s^\top} ds \right], \quad (261)$$

where I_{X_t} and I_{S_t} are square matrices with the main diagonals X_{1t}, \dots, X_{Nt} and S_{1t}, \dots, S_{Kt} , respectively, I a $K \times K$ identity matrix, and $E_t^*[\cdot]$ denotes the expectation under the unique hedge-neutral measure P^* on which are defined N -dimensional Brownian motion w_x^* and K -dimensional Brownian motion w^* with correlation ρ such that the process for the non-tradable assets, X , and stock prices, S , are given by

$$\begin{aligned} \frac{dX_{it}}{X_{it}} &= (m_{it} - \nu_{it}^\top \rho^\top \sigma_t^{-1} (\mu_t - r)) dt + \nu_{it}^\top dw_{xt}^*, \quad i = 1, \dots, N, \\ \frac{dS_{jt}}{S_{jt}} &= r dt + \sigma_{jt}^\top dw_t^*, \quad j = 1, \dots, K, \end{aligned}$$

and the P^* -measure 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 dynamically optimal hedge (260) has the same structure as in the case of the single non-tradable asset and stock, but now additionally incorporates the effects of cross-correlations. This hedge can explicitly be computed for various stochastic investment opportunities, leading to a rich set of comparative statics. The expression (261) for the optimal hedging error variance reveals that the dynamically optimal hedge provides a perfect hedge when $\rho^\top \rho = I$, which generalizes the market completeness condition of Section 3.2.

3.5. Conclusion

This work tackles the problem of dynamic hedging in incomplete markets and provides tractable optimal hedges according to the traditional minimum-variance criterion over the hedging error. The optimal hedges are shown to retain both the simple structure of complete-market hedges and the intuitive features of static hedges, and are in terms of the familiar

Greeks, widely employed in risk management applications. Moreover, in contrast to the existing literature, the hedges are derived via dynamic programming and hence are time-consistent. The dynamically optimal hedges are shown to outperform the static and myopic ones in plausible stochastic environments, coinciding with them only in the simple case of both risky assets following ABMs. They also outperform the pre-commitment hedges after a period of time, as demonstrated in the case of assets following GBMs. Due to its tractability, the baseline analysis can easily be extended in various directions, as shown in the paper.

3.6. Appendix: Proofs

Proof of Proposition 3.1. We obtain the optimal hedge (220) by following the methodology in Chapter 2 and applying dynamic programming to the value function J_t , defined as

$$J(X_t, S_t, W_t, t) \equiv \text{var}_t[X_T - W_T^*]. \quad (262)$$

Suppose, the hedger rebalances the portfolio over time intervals τ . The law of total variance (234) substituted into (262) yields a recursive representation for the value function:

$$J_t = \min_{\theta_t} E_t[J_{t+\tau}] + \text{var}_t[E_{t+\tau}(X_T - W_T)]. \quad (263)$$

We next substitute W_T in (263) by its integral form

$$W_T = W_t e^{r(T-t)} + \int_t^T \theta_s (\mu_s - r) e^{r(T-s)} ds + \int_t^T \theta_s \sigma_s e^{r(T-s)} dw_s, \quad (264)$$

obtained from the budget constraint (218), and take into account that optimal hedges θ_s^* , $s \in [t+\tau, T]$, are already known at time- t from backward induction. Letting the time interval τ shrink to zero and manipulating (263), we obtain the continuous-time HJB equation

$$0 = \min_{\theta_t} E_t[dJ_t] + \text{var}_t[dG_t - d(W_t e^{r(T-t)})], \quad (265)$$

with the terminal condition $J_T = 0$, where G_t is defined by

$$G(X_t, S_t, W_t, t) \equiv E_t[X_T - \int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds]. \quad (266)$$

We note that θ_t^* , J_t and G_t do not depend on wealth W_t . To verify this, we substitute W_T in (264) into the variance criterion and observe that the variance criterion is not affected by W_t , and hence θ_t^* , J_t and G_t depend only on X_t , S_t and t . Applying Itô's lemma to J_t , G_t and $W_t e^{r(T-t)}$, substituting them into (265) and computing the variance term, we obtain the equation

$$\begin{aligned} 0 = \mathcal{D}J_t &+ \nu_t^2 X_t^2 \left(\frac{\partial G_t}{\partial X_t} \right)^2 + 2\rho \nu_t \sigma_t X_t S_t \frac{\partial G_t}{\partial X_t} \frac{\partial G_t}{\partial S_t} + \sigma_t^2 S_t^2 \left(\frac{\partial G_t}{\partial S_t} \right)^2 \\ &+ \min_{\theta_t} \left[\theta_t^2 \sigma_t^2 e^{2r(T-t)} - 2\theta_t \sigma_t \left(\rho \nu_t X_t \frac{\partial G_t}{\partial X_t} + \sigma_t S_t \frac{\partial G_t}{\partial S_t} \right) e^{r(T-t)} \right], \end{aligned} \quad (267)$$

subject to $J_T = 0$. The minimization in (267) has a unique solution

$$\theta_t^* = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial G_t}{\partial X_t} e^{-r(T-t)} + S_t \frac{\partial G_t}{\partial S_t} e^{-r(T-t)}. \quad (268)$$

Substituting (268) back into (267), we obtain the following PDE for the value function

$$\mathcal{D}J_t + (1 - \rho^2) \left(\nu_t X_t \frac{\partial G_t}{\partial X_t} \right)^2 = 0, \quad (269)$$

with the terminal condition $J_T = 0$. The Feynman-Kac solution (Karatzas and Shreve, 1991) to equation (269) is then given by

$$J_t = (1 - \rho^2) E_t \left[\int_t^T \left(\nu_s X_s \frac{\partial G_s}{\partial X_s} \right)^2 ds \right]. \quad (270)$$

To complete the proof it remains to determine the process G_t in terms of the exogenous model parameters. By applying the Feynman-Kac theorem to (266), we obtain a PDE for G_t . Substituting θ_t^* from (268) into this PDE, we obtain the equation

$$\frac{\partial G_t}{\partial t} + \left(m_t - \rho\nu_t \frac{\mu_t - r}{\sigma_t} \right) X_t \frac{\partial G_t}{\partial X_t} + r S_t \frac{\partial G_t}{\partial S_t} + \frac{1}{2} \left(\nu_t^2 X_t^2 \frac{\partial^2 G_t}{\partial X_t^2} + 2\rho\nu_t \sigma_t X_t S_t \frac{\partial^2 G_t}{\partial X_t \partial S_t} + \sigma_t^2 S_t^2 \frac{\partial^2 G_t}{\partial S_t^2} \right) = 0,$$

with the terminal condition $G_T = X_T$. Its Feynman-Kac solution is then given by

$$G_t = E_t^*[X_T], \quad (271)$$

where the expectation is under the unique probability measure P^* on which are defined two Brownian motions w_X^* and w^* such that under P^* the asset X and stock S follow the processes given in (222). Substituting (271) into (268) and (270) yields the optimal hedge (220) and the hedging error variance (221), respectively. To find the Radon-Nikodym derivative dP^*/dP , we decompose the Brownian motion w_X as $dw_{Xt} = \rho dw_t + \sqrt{(1 - \rho^2)} dw_t^\perp$, where $w_t^\perp \equiv (w_{Xt} - \rho w_t) / \sqrt{(1 - \rho^2)}$ is a Brownian motion uncorrelated with w_t . Applying the Girsanov's theorem (Karatzas and Shreve, 1991) to the two-dimensional Brownian motion $(w_t, w_t^\perp)^\top$ yields the Radon-Nikodym derivative (223).

Finally, we derive the representation (225) for $E_t^*[X_T e^{-r(T-t)}]$ by first taking the expectation of (264)

$$E_t \left[W_T^* - W_t e^{r(T-t)} \right] = E_t \left[\int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds \right], \quad (272)$$

and then substituting (271) and (272) into (266).

Proof of Corollary 3.1. Under the probability measure P^* , the process (226) becomes

$$dX_t = (\lambda + \rho\bar{\nu}) \left(\frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}} - X_t \right) dt + \bar{\nu} dw_{Xt}^*, \quad (273)$$

for which the conditional moments are well-known (e.g., Vasicek, 1977), yielding

$$E_t^*[X_T] = \frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}} + \left(X_t - \frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}} \right) e^{-(\lambda + \rho\bar{\nu})(T-t)}.$$

Substituting this into the expressions in Proposition 3.1 yields the desired expressions (228)–(229).

Q.E.D.

Proof of Corollary 3.2. Under the probability measure P^* , the process (230) follows dynamics

$$dX_t = \left(\lambda + \rho\bar{\nu}(\mu - r) \right) \left(\frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}(\mu - r)} - X_t \right) dt + \bar{\nu}\sqrt{X_t} dw_{Xt}^*. \quad (274)$$

The conditional expectation of X_T is well-known (e.g., Cox, Ingersoll, and Ross, 1985) to be

$$E_t^*[X_T] = \frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}(\mu - r)} + \left(X_t - \frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}(\mu - r)} \right) e^{-(\lambda + \rho\bar{\nu}(\mu - r))(T-t)}.$$

Substituting this into the expressions in Proposition 3.1 yields (232)–(233).

Q.E.D.

Proof of Proposition 3.2. First, we derive a variation of the law of total variance. From the law of total variance (234) with an infinitesimally small interval τ , we obtain the following equality in differential form:

$$0 = E_t \left[d \text{var}_s(X_T - W_T) + \text{var}_s(dE_s[X_T - W_T]) \right]. \quad (275)$$

Integrating (275) from t to T yields

$$\text{var}_t[X_T - W_T] = E_t \left[\int_t^T \frac{\text{var}_s(dE_s[X_T - W_T])}{ds} ds \right]. \quad (276)$$

From the assumption $\mu_t - r = 0$ and the integrated budget constraint (264), it follows that $E_t[W_T] = W_t e^{r(T-t)}$. Hence, by Itô's lemma

$$dE_t[X_T - W_T] = (\dots)dt + \nu_t X_t \frac{\partial E_t[X_T]}{\partial X_t} dw_{Xt} + \sigma_t S_t \frac{\partial E_t[X_T]}{\partial S_t} dw_t - \theta_t \sigma_t e^{r(T-t)} dw_t. \quad (277)$$

Substituting (277) into (276) and computing $\text{var}_s(dE_t[X_T - W_T])$, we obtain:

$$\text{var}_t[X_T - W_T] = E_t \left[\int_t^T \left(\theta_s \sigma_s e^{r(T-s)} - \rho \nu_s X_s \frac{\partial E_s[X_T]}{\partial X_s} - \sigma_s S_s \frac{\partial E_s[X_T]}{\partial S_s} \right)^2 + (1 - \rho^2) \left(\frac{\partial E_s[X_T]}{\partial X_s} \right)^2 ds \right]. \quad (278)$$

Minimizing the expression under the integral in (278) gives the global minimum to the variance criterion, yielding the hedge (235). Finally, we observe that for $\mu_t - r = 0$, the dynamically optimal hedge (220) coincides with the hedge (235) since the Radon-Nikodym derivative (223) equals unity, and hence the variance criterion is time-consistent *Q.E.D.*

Proof of Corollary 3.3. The hedging criterion (238) can be represented in integral form (276) in which all the expectations and variances are under the measure P^η (223). By definition of a risk-neutral measure P^η , the stock mean return equals r , and hence $E_t^\eta[W_T] = W_t e^{r(T-t)}$. Then, along the same lines as in the proof of Proposition 3.2, replacing at each step $E_t[\cdot]$ and $\text{var}_t[\cdot]$ by $E_t^\eta[\cdot]$ and $\text{var}_t^\eta[\cdot]$, respectively, it can be shown that the criterion (238) is time-consistent and the solution is given by (220). *Q.E.D.*

Proof of Proposition 3.3. We first compute the optimal hedges and hedging error variances, and then derive the properties of the performance measure Δ_t . From Proposition 3.1, under the measure P^* the process X is a GBM with mean return $(m - \rho\nu(\mu - r)/\sigma)$ and volatility ν , which then yields

$$E_t^*[X_T] = X_t e^{(m - \rho\nu(\mu - r)/\sigma)(T-t)}. \quad (279)$$

Substituting (279) into Proposition 3.1, we obtain the dynamically optimal hedge (243) and the associated hedging error variance

$$\text{var}_t[X_T - W_T^*] = (1 - \rho^2) \nu^2 X_t^2 e^{2(m - \rho\nu\frac{\mu - r}{\sigma})(T-t)} \frac{e^{(\nu^2 + 2\rho\nu\frac{\mu - r}{\sigma})(T-t)} - 1}{\nu^2 + 2\rho\nu\frac{\mu - r}{\sigma}}. \quad (280)$$

The optimal pre-commitment hedge (244) for the case of $r = 0$ and $W_0 = 0$ has been obtained by Duffie and Richardson (1991) in the context of futures hedging.⁴⁷ To obtain

⁴⁷For the case of $r > 0$, Duffie and Richardson provide the optimal pre-commitment hedge assuming interest accrues to a futures margin account, and so such a hedge will be different from that in our economic setting.

it for our case of $r > 0$ and $W_0 > 0$, we observe that the budget constraint (218) can equivalently be rewritten as

$$d\tilde{W}_t = \theta_t \tilde{\mu}_t dt + \theta_t \tilde{\sigma}_t dw_t, \quad (281)$$

where $\tilde{W}_t = W_t e^{r(T-t)} - W_0 e^{rT}$, $\tilde{\mu}_t = (\mu - r)e^{r(T-t)}$, $\tilde{\sigma}_t = \sigma e^{r(T-t)}$. The hedging problem with the budget constraint (281) reduces to the case with $r = 0$ and $\tilde{W}_0 = 0$, and hence the pre-commitment hedge (244) is easily obtained from the solution in Duffie and Richardson.

We next determine $\text{var}_t[X_T - W_T^{\text{commit}}]$ by deriving the first and second moments of an auxiliary process H_t which coincides with the hedging error at $t = T$ and is defined as

$$H_t \equiv X_t e^{(m-r-\rho\nu\frac{\mu-r}{\sigma})(T-t)} - W_t^{\text{commit}} e^{r(T-t)}. \quad (282)$$

Substituting the pre-commitment hedge (244) into the budget constraint (218) and applying Itô's lemma to H_t we obtain:

$$dH_t = \left(\frac{\mu-r}{\sigma}\right)^2 (H_0 - H_t) dt + \frac{\mu-r}{\sigma} (H_0 - H_t) dw_t + \sqrt{1-\rho^2} \nu X_t e^{(m-\rho\nu\frac{\mu-r}{\sigma})(T-t)} dw_t^\perp. \quad (283)$$

Integrating (283) from t to τ and taking the time- t expectation on both sides yields a simple linear integral equation for $E_t[H_\tau]$, the unique solution to which is given by

$$E_t[H_\tau] = H_0 + (H_t - H_0) e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (\tau-t)}.$$

To find the second moment of H_t , we apply Itô's lemma to $(H_t - H_0)^2$:

$$d(H_t - H_0)^2 = -\left(\left(\frac{\mu-r}{\sigma}\right)^2 (H_t - H_0)^2 - (1-\rho^2)\nu^2 X_t^2 e^{2(m-\rho\nu\frac{\mu-r}{\sigma})(T-t)}\right) dt + (\dots) dw_t + (\dots) dw_t^\perp.$$

Integrating both sides from t to τ and then taking the time- t expectation we obtain $E_t[(H_\tau - H_0)^2]$ as the solution to a linear integral equation given by

$$\begin{aligned} E_t[(H_\tau - H_0)^2] &= (H_t - H_0)^2 e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (\tau-t)} \\ &+ (1-\rho^2)\nu^2 X_t^2 e^{2(m-\rho\nu\frac{\mu-r}{\sigma})(\tau-t)} \frac{e^{(\nu^2+2\rho\nu\frac{\mu-r}{\sigma})(\tau-t)} - e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (\tau-t)}}{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + \left(\frac{\mu-r}{\sigma}\right)^2}. \end{aligned}$$

Given the first two moments of H_t and taking into account that $H_T = X_T - W_T^{\text{commit}}$, we obtain:

$$\begin{aligned} \text{var}_t[X_T - W_T^{\text{commit}}] &= (H_t - H_0)^2 e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (T-t)} \left(1 - e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (T-t)}\right) \\ &+ (1-\rho^2)\nu^2 X_t^2 e^{2(m-\rho\nu\frac{\mu-r}{\sigma})(T-t)} \frac{e^{(\nu^2+2\rho\nu\frac{\mu-r}{\sigma})(T-t)} - e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (T-t)}}{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + \left(\frac{\mu-r}{\sigma}\right)^2}. \end{aligned} \quad (284)$$

Since the second moments of X_t and H_t are determined explicitly, it is straightforward to explicitly compute $E_0[\text{var}_t(X_T - W_T^*)]$ and $E_0[\text{var}_t(X_T - W_T^{\text{commit}})]$. The relative performance measure (240) is then given by:

$$\begin{aligned} \Delta_t &= \frac{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma}}{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + (\frac{\mu-r}{\sigma})^2} e^{-(\frac{\mu-r}{\sigma})^2(T-t)} \left(\frac{e^{(\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + (\frac{\mu-r}{\sigma})^2)(T-t)} - 1}{e^{(\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma})(T-t)} - 1} \right. \\ &\quad \left. + \frac{1 - e^{-(\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + (\frac{\mu-r}{\sigma})^2)t}}{e^{(\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma})(T-t)} - 1} \left(1 - e^{-(\frac{\mu-r}{\sigma})^2(T-t)} \right) \right) - 1. \end{aligned} \quad (285)$$

Letting t go to T in (285), it is easy to show that $\Delta_T > 0$, and hence there exists a time \bar{t} such that $\Delta_t > 0$ whenever $t > \bar{t}$.

Finally, we derive relation (245) by rearranging terms in (225) and substituting $\bar{Q}.E.D.$

Proof of Proposition 3.4. First, we consider the case when both risky assets follow ABMs (249)–(250). From Proposition 3.1, the process for X under the measure P^* is given by

$$dX_t = (\bar{m} - \rho\bar{\nu}\frac{\bar{\mu} - rS_t}{\bar{\sigma}})dt + \bar{\nu}dw_{X_t}^*.$$

Integrating from t to T and taking the expectation $E_t^*[\cdot]$ on both sides we obtain:

$$E_t^*[X_T] = X_t + m(T-t) - \rho\bar{\nu}\frac{\bar{\mu}(T-t) - S_t(e^{r(T-t)} - 1)}{\bar{\sigma}}.$$

Substituting this into Proposition 3.1 yields the optimal hedge reported in Table 3.1. Since X and S follow ABMs, $\text{cov}_0(X_T, S_T) = \rho\bar{\nu}\bar{\sigma}T$ and $\text{var}_0(S_T) = \bar{\sigma}^2T$. Substituting these into θ_t^{static} in (247), we obtain the static hedge, which coincides with the dynamic one, and hence $\Delta_t = 0$.

When the risky asset prices follow GBMs (241)–(242), the dynamically optimal hedge and its corresponding hedging error variance are given by (243) and (280), respectively. The static hedge reported in Table 3.1 is obtained from (247) by observing that since X_t^2 , S_t^2 and X_tS_t follow GBMs,

$$\text{var}_t[X_T] = X_t^2 e^{2m(T-t)} (e^{\nu^2(T-t)} - 1), \quad \text{var}_t[S_T] = S_t^2 e^{2\mu(T-t)} (e^{\sigma^2(T-t)} - 1), \quad (286)$$

$$\text{cov}_t(X_T, S_T) = X_t S_t e^{2(m+\mu)(T-t)} (e^{\rho\nu\sigma(T-t)} - 1). \quad (287)$$

Substituting W_T^{static} from the static budget constraint (246) into the hedging error variance we obtain

$$\text{var}_t[X_T - W_T^{static}] = \text{var}_t[X_T] - 2\frac{\theta_0^{static}}{S_0} \text{cov}_t(X_T, S_T) + \left(\frac{\theta_0^{static}}{S_0}\right)^2 \text{var}_t[S_T]. \quad (288)$$

We now show that for $\rho > 0$ the performance measure (240) is positive in this GBMs case. We note that the static hedging error variance (288) is a quadratic function of θ_0 , the minimization of which along with the expressions (286)–(287) gives the lower bound for the static hedging error variance:

$$\text{var}_t[X_T - W_T^{static}] \geq X_t^2 e^{2m(T-t)} \left(e^{\nu^2(T-t)} - 1 - \frac{(e^{\rho\nu\sigma(T-t)} - 1)^2}{e^{\sigma^2(T-t)} - 1} \right). \quad (289)$$

We next rewrite the dynamically optimal hedging error variance (280) in integral form and find its upper bound for $\rho > 0$ as:

$$\begin{aligned} \text{var}_t[X_T - W_T^*] &= X_t^2 (1 - \rho^2) \nu^2 e^{2m(T-t)} \int_t^T e^{-2\rho\nu\frac{\mu-r}{\sigma}(T-s)} e^{\nu^2(s-t)} ds \\ &\leq X_t^2 (1 - \rho^2) \nu^2 e^{2m(T-t)} \int_t^T e^{\nu^2(s-t)} ds = X_t^2 (1 - \rho^2) e^{2m(T-t)} (e^{\nu^2(T-t)} - 1). \end{aligned} \quad (290)$$

A sufficient condition for the dynamically optimal variance to be lower than the static one is that the upper bound in (290) is below the lower bound in (289), which is equivalent to

$$\left(\frac{e^{\rho\nu\sigma(T-t)} - 1}{\rho\nu\sigma} \right)^2 \leq \left(\frac{e^{\sigma^2(T-t)} - 1}{\sigma^2} \right) \left(\frac{e^{\nu^2(T-t)} - 1}{\nu^2} \right). \quad (291)$$

To show that inequality holds, we rewrite its left-hand side as a squared integral, estimate it from above and then apply the Cauchy-Schwartz inequality:

$$\left(\int_t^T e^{\rho\nu\sigma(T-s)} ds \right)^2 \leq \left(\int_t^T e^{(\frac{\nu^2}{2} + \frac{\sigma^2}{2})(T-s)} ds \right)^2 \leq \left(\int_t^T e^{\sigma^2(T-s)} ds \right) \left(\int_t^T e^{\nu^2(T-s)} ds \right). \quad (292)$$

Computing the integrals in (292) we obtain inequality (291), and hence $\Delta_t > 0$.

For $\rho < 0$ in the case of GBMs, we demonstrate that θ_t^* still outperforms after a certain period of time \bar{t} . Substituting the dynamically optimal and static hedging error variances, (280) and (288), into the performance measure (240), and taking limit as t goes to T we obtain:

$$\Delta_T = \frac{E_0[\nu^2 X_T^2 - 2\rho\nu\sigma X_T S_T \frac{\theta_0^{static}}{S_0} + \sigma^2 S_T^2 (\frac{\theta_0^{static}}{S_0})^2]}{(1 - \rho^2) \nu^2 E_0[X_T^2]} - 1 \equiv \frac{E_0[(\rho\nu X_T - S_T \frac{\theta_0^{static}}{S_0})^2]}{(1 - \rho^2) \nu^2 E_0[X_T^2]}.$$

Since $\Delta_T > 0$ there exists \bar{t} such that $\Delta_t > 0$ whenever $t > \bar{t}$. For some parameter values of ρ , σ , ν and T , the performance measure Δ_t can become negative but only for implausibly large $(\mu - r)/\sigma$.⁴⁸

The remainder of the results for the case when the asset X follows an OU process and the stock S follows either an ABM or a GBM can be obtained along the lines of above. First, we compute the optimal hedges and corresponding hedging error variances by applying Proposition 3.1. Then, we characterize the static hedges by computing relevant moments for the processes X_t , S_t and $X_t S_t$. Finally, we obtain a lower bound for the static hedging error variance as above, and compare it with the dynamically optimal one or its upper bound.

Q.E.D.

Proof of Proposition 3.5. The dynamically optimal hedges reported in Table 3.2 are the same as in Table 3.1, while the myopic hedges are immediate from the expression (253). Thus, it remains to compare the relative performances.

When the risky assets follow ABMs, the two hedges coincide and hence $\Delta_t = 0$. Turning to the case when both X and S follow GBMs, we derive the myopic hedging error variance using the expanded law of total variance (276) and compare it with the dynamically optimal one. From the budget constraint in integral form (264) and the expression for the myopic hedge we obtain:

$$E_t[X_T - W_T^{myopic}] = X_t e^{m(T-t)} \left(1 - \rho \nu \frac{\mu - r}{\sigma} \int_t^T e^{-(m-r)(T-s)} ds \right) - W_t e^{r(T-t)}.$$

Applying Itô's lemma, we derive $dE_t[X_T - W_T^{myopic}]$, substitute it into the law of total variance (276), and after some algebra determine the myopic hedging error variance:

$$\begin{aligned} \text{var}_t[X_T - W_T^{myopic}] &= X_t^2 \nu^2 e^{2m(T-t)} \int_t^T e^{\nu^2(s-t)} \left((1 - \rho^2) \left(1 - \rho \nu \frac{\mu - r}{\sigma} \frac{1 - e^{-(m-r)(T-s)}}{m-r} \right)^2 \right. \\ &\quad \left. + \rho^2 \nu^2 \left(\frac{m-r}{\nu} - \rho \frac{\mu - r}{\sigma} \right)^2 \left(\frac{1 - e^{-(m-r)(T-s)}}{m-r} \right)^2 \right) ds. \end{aligned} \quad (293)$$

⁴⁸For example, if $\rho = -0.2$, $\nu = 0.36$, $\sigma = 0.16$ and $T = 1$, the static policy outperforms at time 0 only for $(\mu - r)/\sigma > 0.6$. If the parameter ρ increases in absolute value, the lower boundary for the market price of risk also increases.

We now show that if $(m - r)/\nu > \rho(\mu - r)/\sigma$, the dynamically optimal hedge outperforms the myopic one. Comparing the dynamically optimal and myopic hedging error variances given by (290) and (293) we observe that a sufficient condition for the dynamically optimal hedge to outperform the myopic one is

$$1 - \rho\nu \frac{\mu - r}{\sigma} \frac{1 - e^{-(m-r)(T-t)}}{m - r} > e^{-\rho\nu \frac{\mu-r}{\sigma}(T-t)}.$$

This inequality can equivalently be rewritten as

$$\int_t^T e^{-\rho\nu \frac{\mu-r}{\sigma}(T-s)} ds \geq \int_t^T e^{-(m-r)(T-s)} ds,$$

which holds whenever condition $(m - r)/\nu > \rho(\mu - r)/\sigma$ is satisfied. If this condition is violated, the dynamically optimal hedge can outperform or underperform.

The remainder of the relative performance results for the case when the asset X follows an OU process while the stock S is an ABM or a GBM are obtained similarly. The dynamically optimal and myopic hedging error variances are obtained from the expressions (221) and (276). It is then directly observed that the myopic hedging error variance exceeds the dynamically optimal one for all parameters, and hence $\Delta_t > 0$. *Q.E.D.*

Proof of Proposition 3.6. The proof is similar to the proof of Proposition 3.1. The hedging problem is solved via dynamic programming and the value function is defined as:

$$J(X_t, S_t, W_t, t) = E_t[X_T - W_T^*] - \frac{\gamma}{2} \text{var}_t[X_T - W_T^*]. \quad (294)$$

Applying the law of total variance along the same steps as in the proof of Proposition 3.1, we obtain an HJB equation. To solve this equation, substituting the budget constraint in integral form (264) into the hedger's objective (254), we show that the objective is linear in W_t and hence θ_t^* and G_t do not depend on W_t . In contrast to the minimum-variance case, the value function linearly depends on $W_t e^{r(T-t)}$ and can be represented as:

$$J(X_t, S_t, W_t, t) = W_t e^{r(T-t)} + \hat{J}(X_t, S_t, t).$$

Applying Itô's lemma to the processes \hat{J}_t , G_t and $W_t e^{r(T-t)}$ we obtain a PDE for the value function and the optimal hedge in a recursive form. The optimal hedge in terms of exogenous

parameters is then obtained by applying the Feynman-Kac theorem, as in Proposition 3.1. Solving the PDE for J_t , we obtain the value function (256). *Q.E.D.*

Proof of Proposition 3.7. Proposition 3.7 is a multidimensional version of Proposition 3.1 and can be proven along the same lines. First, using the law of total variance, we derive an HJB equation and then the optimal hedge in a recursive form. Then, applying the Feynman-Kac theorem we find the optimal hedge in terms of exogenous parameters. Finally, solving the HJB PDE for the value function, we obtain the hedging error variance in closed form. *Q.E.D.*

References

- Acharya, V.V., and L.H. Pedersen, 2005, "Asset Pricing with Liquidity Risk," *Journal of Financial Economics*, 77, 375-410.
- Adler, M., and J.B. Detemple, 1988, "On the Optimal Hedge of a Nontraded Cash Position," *Journal of Finance*, 43, 143-153.
- Ait-Sahalia, Y., and M.W. Brandt, 2001, "Variable Selection for Portfolio Choice," *Journal of Finance*, 56, 1297-1351.
- Anderson, R.W., and J.-P. Danthine, 1980, "Hedging and Joint Production: Theory and Illustrations," *Journal of Finance*, 35, 487-498.
- Anderson, R.W., and J.-P. Danthine, 1981, "Cross Hedging," *Journal of Political Economy*, 89, 1182-1196.
- Anderson, R.W., and J.-P. Danthine, 1983, "The Time Pattern of Hedging and the Volatility of Futures Prices," *Review of Economic Studies*, 50, 249-266.
- 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.
- Bansal, R., and A. Yaron, 2004, "Risk for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481-1509.
- Barberis, N., 2000, "Investing for the Long Run When Returns Are Predictable," *Journal of Finance*, 55, 225-264.
- Basak, S., 2000, "A Model of Dynamic Equilibrium Asset Pricing with Heterogeneous Beliefs and Extraneous Risk," *Journal of Economic Dynamics and Control*, 24, 63-95.

-
- Basak, S., 2005, "Asset Pricing with Heterogeneous Beliefs," *Journal of Banking and Finance*, 29, 2849-2881.
- Basak, S., and B. Croitoru, 2000, "Equilibrium Mispricing in a Capital Market with Portfolio Constraints," *Review of Financial Studies*, 13, 715-748.
- Basak, S., and B. Croitoru, 2006, "On the Role of Arbitrageurs in Rational Markets," *Journal of Financial Economics*, 81, 143-173.
- Basak, S., and D. Cuoco, 1998, "An Equilibrium Model with Restricted Stock Market Participation," *Review of Financial Studies*, 11, 309-341.
- Beckers, S., 1980, "The Constant Elasticity of Variance Model and Its Implications for Option Pricing," *Journal of Finance*, 35, 661-673.
- Berrada, T., 2009, "Bounded Rationality and Asset Pricing with Intermediate Consumption," *Review of Finance*, forthcoming.
- Bertsimas, D., L. Kogan, and A. Lo, 2001, "Hedging Derivative Securities and Incomplete Markets: An ϵ -Arbitrage Approach," *Operations Research*, 49, 372-397.
- Bhamra, H.S., 2007, "Stock Market Liberalization and the Cost of Capital in Emerging Markets," Working Paper, University of British Columbia.
- Bhamra, H.S., and R. Uppal, 2009, "The Effect of Introducing a Non-Redundant Derivative on the Volatility of Stock-Market Returns When Agents Differ in Risk Aversion," *Review of Financial Studies*, 22, 2303-2330.
- Bielecki, T.R., M. Jeanblanc, and M. Rutkowski, 2004, "Mean-Variance Hedging of Defaultable Claims," Working Paper, Illinois Institute of Technology.
- 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., 2003, "Hedging Demands in Hedging Contingent Claims," *Review of Economics and Statistics*, 85, 119-140.
- Brandt, M.W., 2009, "Portfolio Choice Problems," in Ait-Sahalia, Y., and L.P. Hansen (eds.), *Handbook of Financial Econometrics*.
- 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.
- Breeden, D.T., 1984, "Futures Markets and Commodity Options: Hedging and Optimality in Incomplete markets," *Journal of Economic Theory*, 32, 275-300.
- Brennan, M.J., and Y. Xia, 2002, "Dynamic Asset Allocation under Inflation," *Journal of Finance*, 57, 1201-1238.
- Brooks, C., O.T. Henry, and G. Persaud, 2002, "The Effect of Asymmetries on Optimal Hedge Ratios," *Journal of Business*, 75, 333-352.
- Campbell, J.Y., 2003, "Consumption-Based Asset Pricing," *Handbook of Economics and Finance*, G. Constantinides, M. Harris, and R. Stulz (eds.), North Holland, Amsterdam.
- Campbell, J.Y., 2006, "Household Finance," *Journal of Finance*, 61, 1553-1604.
- Campbell, J.Y., and J.H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205-251.

-
- Campbell, J.Y., K. Serfaty-de Medeiros, and L.M. Viceira, 2009, "Global Currency Hedging," *Journal of Finance*, forthcoming.
- 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., and L.M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.
- Caplin, A., and J. Leahy, 2006, "The Recursive Approach to Time Inconsistency," *Journal of Economic Theory*, 131, 134-156.
- 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.
- Chan, L.K.C., J. Karceski, and J. Lakonishok, 1999, "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model," *Review of Financial Studies*, 12, 937-974.
- Chan, Y.L., and L. Kogan, 2002, "Catching Up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices," *Journal of Political Economy*, 110, 1255-1285.
- Chien, Y., H. Cole, and H. Lustig, 2008, "A Multiplier Approach to Understanding the Macro Implications of Household Finance," Working Paper, Purdue University.
- Cochrane, J.H., 2008, "A Mean Variance Benchmark for Intertemporal Portfolio Theory," Working Paper, University of Chicago.
- Cogley, T., and T.J. Sargent, 2008, "The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?" *Journal of Monetary Economics*, 55, 454-476.
- Cornell, B., and R. Roll, 2005, "A Delegated-Agent Asset-Pricing Model," *Financial Analysts Journal*, 61, 57-69.
- Costa, O.L.V., and A.C. Paiva, 2002, "Robust Portfolio Selection Using Linear-Matrix Inequalities," *Journal of Economic Dynamics and Control*, 26, 889-909.

- 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.
- Cuoco, D., 1997, "Optimal Consumption and Equilibrium Prices with Portfolio Constraints and Stochastic Income," *Journal of Economic Theory*, 72, 33-73.
- Cuoco, D., and H. He, 2001, "Dynamic Aggregation and Computation of Equilibria in Finite-Dimensional Economies with Incomplete Markets," *Annals of Economics and Finance*, 2, 265-296.
- 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., and I. Karatzas, 1992, "Convex Duality in Constrained Optimization," *Annals of Applied Probability*, 2, 767-818.
- Cvitanic, J., A. Lazrak, and T. Wang, 2008, "Implications of the Sharpe Ratio as a Performance Measure in Multi-Period Settings," *Journal of Economic Dynamics and Control*, 32, 1622-1649.
- 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.
- Detemple, J., and S. Murthy, 1997, "Equilibrium Asset Prices and No-Arbitrage with Portfolio Constraints," *Review of Financial Studies*, 10, 1133-1174.

-
- Dixit, R.K., and R.S. Pindyck, 1994, *Investment under Uncertainty*, Princeton University Press.
- Duffie, D., 1989, *Futures Markets*, Prentice Hall.
- Duffie, D., W. Fleming, H.M. Soner, and T. Zariphopoulou, 1997, "Hedging in Incomplete Markets with HARA Utility," *Journal of Economic Dynamics and Control*, 21, 753-782.
- Duffie, D., and M.O. Jackson, 1989, "Optimal Innovation of Futures Contracts," *Review of Financial Studies*, 2, 275-296.
- Duffie, D., and M.O. Jackson, 1990, "Optimal Hedging and Equilibrium in a Dynamic Futures Market," *Journal of Economic Dynamics and Control*, 14, 21-33.
- Duffie, D., and H.R. Richardson, 1991, "Mean-Variance Hedging in Continuous Time," *Annals of Applied Probability*, 1, 1-15.
- Duffie, D., and L. Epstein, 1992, "Stochastic Differential Utility," *Econometrica*, 60, 353-394.
- Dumas, B., 1989, "Two-Person Dynamic Equilibrium in the Capital Market," *Review of Financial Studies*, 2, 157-188.
- Dumas, B., A. Kurshev, and R. Uppal, 2009, "Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility," *Journal of Finance*, 64, 579-629.
- Dumas, B., and A. Lyasoff, 2008, "Incomplete market Equilibria Solved Recursively on an Event Tree," Working Paper, Swiss Finance Institute.
- Dumas, B., and P. Maenhout, 2002, "A Central-Planning Approach to Dynamic Incomplete-Market Equilibrium," Working Paper, INSEAD.
- Ederington, L.H., 1979, "The Hedging Performance of the New Futures Markets," *Journal of Finance*, 34, 157-170.
- Ekeland, I., and A. Lazrak, 2008, "Equilibrium Policies when Preferences are Time Inconsistent," Working Paper, University of British Columbia.

-
- Ferson, W.E, and C.R. Harvey, 1991, "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99, 385-415.
- Figlewski, S., 1984, "Hedging Performance and Basis Risk in Stock Index Futures," *Journal of Finance*, 39, 657-669.
- Follmer, H., and M. Schweizer, 1991, "Hedging Contingent Claims under Incomplete Information," in Davis, M.H.A., and R.J. Elliott (eds.) *Applied Stochastic Analysis*, Gordon and Breach, 389-414.
- 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.
- 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.
- Gallmeyer, M., and B. Hollifield, 2008, "An Examination of Heterogeneous Beliefs with a Short-Sale Constraint in a Dynamic Equilibrium," *Review of Finance*, 12, 323-364.
- Gibson, R., and E.S. Schwartz, 1990, "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims," *Journal of Finance*, 45, 959-976.
- Gomes, F., and A. Michaelides, 2008, "Asset Pricing with Limited Risk Sharing and Heterogeneous Agents," *Review of Financial Studies*, 21, 416-448.
- Gomez, J.-P., and F. Zapatero, 2003, "Asset Pricing Implications of Benchmarking: a Two-Factor CAPM," *European Journal of Finance*, 9, 343-357.
- Gourieroux C., L.P. Laurent, and H. Pham, 1998, "Mean-Variance Hedging and Numeraire," *Mathematical Finance*, 8, 179-200.
- Grenadier, S.R., and N. Wang, 2007, "Investment under Uncertainty and Time-Inconsistent Preferences," *Journal of Financial Economics*, 84, 2-39.
- Guvenen, F., 2006, "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective," *Journal of Monetary Economics*, 53,

1451-1472.

- Harris, C., and D. Laibson, 2001, "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 69, 935-957.
- He, Z., and A. Krishnamurthy, 2008, "Intermediary Asset Prices," Working Paper, Northwestern University.
- Henderson, V., 2005, "Explicit Solutions to an Optimal Portfolio Choice Problem with Stochastic Income," *Journal of Economic Dynamics and Control*, 29, 1237-1266.
- 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.
- Hirshleifer, D., 1988, "Residual Risk, trading Costs, and Commodity Futures Risk Premia," *Review of Financial Studies*, 1, 173-193.
- Hong, H., J.D. Kubik, and J.C. Stein, 2004, "Social Interaction and Stock Market Participation," *Journal of Finance*, 59, 137-163.
- 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.
- Huang, C.-F., and H. Pages, 1992, "Optimal Consumption and Portfolio Policies with an Infinite Horizon: Existence and Convergence," *Annals of Applied Probability*, 2, 36-64.
- Hugonnier, J., 2008, "Bubbles and Multiplicity of Equilibria under Portfolio Constraints," Working Paper, University of Lausanne.
- Hull, J.C., 2008, *Options, Futures and Other Derivatives*, Seventh Edition, Prentice Hall.
- In, F., and S. Kim, 2006, "The Hedge Ratio and the Empirical Relationship between the Stock and Futures Markets," *Journal of Business*, 79, 799-820.
- Jagannathan, R., and T. Ma, 2003, "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraint Helps," *Journal of Finance*, 58, 1651-1683.

-
- Jaillet, P., E.I. Ronn, and S. Tompaidis, 2004, "Valuation of Commodity-Based Swing Options," *Management Science*, 50, 909-921.
- Jarrow, R., 1980, "Heterogeneous Expectations, Restrictions on Short Sales, and Equilibrium Asset Prices," *Journal of Finance*, 25, 1105-1113.
- Johnsen, T.H., and J.B. Donaldson, 1985, "The Structure of Intertemporal Preferences under Uncertainty and Time Consistent Plans," *Econometrica*, 53, 1451-1458.
- Jorion, F., 2003, "Portfolio Optimization with Tracking-Error Constraints," *Financial Analysts Journal*, 59, 70-82.
- Kamara, A., and A.F. Siegel, 1987, "Optimal Hedging in Futures Markets with Multiple Delivery Specification," *Journal of Finance*, 42, 1007-1021.
- 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.
- Kerkvliet, J., and M.H. Moffett, 1991, "The Hedging of an Uncertain Future Foreign Currency Cash Flow," *Journal of Financial and Quantitative Analysis*, 26, 565-578.
- Kihlstrom, R., 2008, "Risk Aversion and the Elasticity of Substitution in General Dynamic Portfolio Theory: Consistent Planning by Forward Looking, Expected Utility Maximizing Investor," *Journal of Mathematical Economics*, forthcoming.
- Kim, T.S., and E. Omberg, 1996, "Dynamic Nonmyopic Portfolio Behavior," *Review of Financial Studies*, 9, 141-161.
- Kogan, L., I. Makarov, and R. Uppal, 2007, "The Equity Premium and the Riskfree Rate in an Economy with Borrowing Constraints," *Mathematics and Financial Economics*, 1, 1-19.

-
- Kogan, L., S.A. Ross, J. Wang, and M.M. Westerfield, "The Price Impact and Survival of Irrational Traders," *Journal of Finance*, 61, 195-229.
- Korn, R., and H. Kraft, 2004, "On the Stability of Continuous-Time Portfolio Problems with Stochastic Opportunity Sets," *Mathematical Finance*, 14, 403-414.
- Kroner, K.F., and J. Sultan, 1993, "Time-Varying Distributions and Dynamic Hedging with Foreign Currency Futures," *Journal of Financial and Quantitative Analysis*, 28, 535-551.
- 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.
- Lioui, A., and P. Poncet, 2000, "The Minimum Variance Hedge Ratio Under Stochastic Interest Rates," *Management Science*, 46, 658-668.
- 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.
- Ljungqvist, L., and T.J. Sargent, 2004, *Recursive Macroeconomic Theory*, MIT Press.
- Longstaff, F.A., 2009, "Asset Pricing in Markets with Illiquid Assets," *American Economic Review*, forthcoming.
- Maenhout, P., 2006, "Robust Portfolio Rules and Detection-Error Probabilities for a Mean-Reverting Risk Premium," *Journal of Economic Theory*, 128, 136-163.
- Malloy, C.J., T.J. Moskowitz, and A. Vissing-Jorgensen, 2009, "Long-Run Stockholder Consumption Risk and Asset Returns," *Journal of Finance*, forthcoming.

-
- Mankiw, N.G., and S.P. Zeldes, 1991, "The Consumption of Stockholders and Non-stockholders," *Journal of Financial Economics*, 29, 97-112.
- Markowitz, H.M., 1952, "Portfolio Selection," *Journal of Finance*, 7, 77-91.
- McDonald, R., 2006, *Derivatives Markets*, Second Edition, Addison Wesley.
- Mehra, R., and E.C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145-161.
- 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.
- Musiela, M., and M. Rutkowski, 1998, *Martingale Methods in Financial Modelling*, Springer-Verlag.
- Pavlova, A., and R. Rigobon, 2008, "The Role of Portfolio Constraints in the International Propagation of Shocks," *Review of Economic Studies*, 75, 1215-1256.
- Peleg, B., and M.E. Yaari, 1973, "On the Existence of Consistent Course of Action when Tastes are Changing," *Review of Economic Studies*, 40, 391-401.
- Pham, H., T. Rheinlander, and M. Schweizer, 1998, "Mean-Variance Hedging for Continuous Processes: New Proofs and Examples," *Finance and Stochastics*, 2, 173-198.
- Pindyck, R.S., 2004, "Volatility and Commodity Price Dynamics," *Journal of Futures Markets*, 24, 1029-1047.
- Rolfo, J., 1980, "Optimal Hedging under Price and Quantity Uncertainty: The Case of a Cocoa Producer," *Journal of Political Economy*, 88, 100-116.
- Roll, R., 1992, "A Mean-Variance Analysis of Tracking Error," *Journal of Portfolio Management*, 18, 13-23.

-
- Samuelson, W., and R. Zeckhauser, 1988, "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1, 7-59.
- 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.
- Schornick, A.V., 2009, "Competing Capital Constraints," Working Paper, INSEAD.
- Schroder, M., 1989, "Computing the Constant Elasticity of Variance Option Pricing Formula," *Journal of Finance*, 44, 211-219.
- Schwartz, E.S., 1997, "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging," *Journal of Finance*, 52, 923-973.
- Schwartz, E.S., and J.E. Smith, 2000, "Short-Term Variations and Long-Term Dynamics in Commodity Prices," *Management Science*, 46, 893-911.
- Schweizer, M., 1992, "Mean-Variance Hedging for General Claims," *Annals of Probability*, 2, 171-179.
- Schweizer, M., 1994, "Approximating Random Variables by Stochastic Integrals," *Annals of Probability*, 22, 1536-1575.
- Schweizer, M., 1999, "A Minimality Property of the Minimal Martingale Measure," *Statistics and Probability Letters*, 42, 27-31.
- Schweizer, M., 2001, "A Guided Tour through Quadratic Hedging Approaches," in Jouini, J., J. Cvitanic, and M. Musiela (eds.), *Option Pricing, Interest Rates and Risk Management*, Cambridge University Press, 538-574.
- Schwert, W., 1989, "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance*, 44, 1115-1153.
- Siegel, D.R., and D.F. Siegel, 1990, *The Futures Markets: Arbitrage, Risk Management and Portfolio Choice*, McGraw Hill.
- Srinivas, P.S., E. Whitehouse, and J. Yermo, 2000, "Regulating Private Pension Funds' Structure, Performance and Investments: Cross-Country Evidence," Working Paper, World Bank.

-
- Strotz, R.H., 1956, "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165-180.
- Stulz, R.M., 1984, "Optimal Hedging Policies," *Journal of Financial and Quantitative Analysis*, 19, 127-140.
- Stulz, R.M., 2003, *Risk Management and Derivatives*, First Edition, South Western College Publishing.
- Svensson, L.E.O, and I.M. Werner, 1993, "Nontraded Assets in Incomplete Markets: Pricing and Portfolio Choice," *European Economic Review*, 37, 1149-1168.
- Tepla, L., 2000a, "Optimal Portfolio Policies with Borrowing and Shortsale Constraints," *Journal of Economic Dynamics and Control*, 24, 1623-1639.
- Tepla, L., 2000b, "Optimal Hedging and Valuation of Non-traded Assets," *European Economic Review*, 4, 231-251.
- Vasicek, O., 1977, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177-188.
- Viceira, L.M., 2001, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income," *Journal of Finance*, 56, 433-470.
- 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.
- Wang, J., 1996, "The Term Structure of Interest rates in a Pure Exchange Economy with Heterogeneous Investors," *Journal of Financial Economics*, 41, 75-110.
- Weiss, N., 2005, *A Course in Probability*, Addison-Wesley.
- Wu, T., 2008, "An Equilibrium Model with Buy and Hold Investors", Working Paper, Illinois Institute of Technology.
- Yan, H., 2008, "Natural Selection in Financial Markets: Does it Work?" *Management Science*, 54, 1935-1950.

- 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.