

Dynamic Hedging in Incomplete Markets: A Simple Solution*

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

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

This revision: November 2008

*We are grateful to Mike Chernov, Francisco Gomes and the seminar participants at London Business School for helpful comments. All errors are our responsibility.

Dynamic Hedging in Incomplete Markets: A Simple Solution

Abstract

Despite much work on hedging in incomplete markets, the literature still lacks tractable dynamic hedges in plausible environments. In this article, we provide a simple solution to this problem in a general incomplete-market economy in which a hedger, guided by the traditional minimum-variance criterion, aims at reducing the risk of a non-tradable asset. We derive fully analytical optimal 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, as well as retaining the intuitive features of their static counterparts. We obtain our time-consistent hedges by dynamic programming, while the extant literature characterizes either static or myopic hedges, or dynamic ones that minimize the variance criterion at an initial date and from which the hedger may deviate unless she can pre-commit to follow them. We demonstrate that our dynamically optimal hedges typically outperform their static and myopic counterparts under plausible economic environments. We also show that our results can be applied to portfolio management with tracking-error.

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

Keywords: Hedging, Incomplete Markets, Minimum-Variance Criterion, Risk Management, Stochastic Economies, Time-Consistency.

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 Basak and Chabakauri (2008). In dynamically complete markets, there is no time-inconsistency issue (unlike the problem in Basak and Chabakauri) 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 (2000) 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. Gourieroux, 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 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, we compare our dynamically optimal hedge with the pre-commitment, static and myopic hedges, while in Section 4, we generalize our baseline model to the case of mean-variance hedging and the case of multiple assets. Section 5 concludes. Proofs are in the Appendix.

2. Dynamic Minimum-Variance Hedging

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.¹ Without loss of generality, we adopt the first interpretation and postulate the price of the non-tradable asset to follow the dynamics

$$\frac{dX_t}{X_t} = m(X_t, t)dt + \nu(X_t, t)dw_{xt}, \quad (1)$$

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

$$\frac{dS_t}{S_t} = \mu(S_t, X_t, t)dt + \sigma(S_t, X_t, t)dw_t, \quad (2)$$

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 (1)–(2).

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

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

¹If 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 .

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; Gourieroux, Laurent and Pham 1998; Betsimas, 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., Gourieroux, 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 3.2-3.3), 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 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

$$\min_{\theta} \text{var}_t[X_T - W_T], \tag{4}$$

subject to the dynamic budget constraint (3). We solve this problem by dynamic programming and hence provide the time-consistent dynamic hedging policy.

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 Basak and Chabakauri (2008) 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 1 reports the optimal policy derived from the solution to this equation and the resulting optimal quality of the hedge.

Proposition 1. *The optimal hedging policy and the corresponding variance of the hedging error are given by*

$$\theta_t^* = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial S_t}, \quad (5)$$

$$\text{var}_t[X_T - W_T^*] = (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 X_s^2 \left(\frac{\partial E_s^*[X_T]}{\partial X_s} \right)^2 ds \right], \quad (6)$$

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

$$\frac{dX_t}{X_t} = \left(m_t - \rho\nu_t \frac{\mu_t - r}{\sigma_t} \right) dt + \nu_t dw_{Xt}^*, \quad \frac{dS_t}{S_t} = r dt + \sigma_t dw_t^*, \quad (7)$$

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

$$\frac{dP^*}{dP} = e^{-\frac{1}{2} \int_0^T \left(\frac{\mu_s - r}{\sigma_s} \right)^2 ds - \int_0^T \frac{\mu_s - r}{\sigma_s} dw_s}. \quad (8)$$

Proposition 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 (5) 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

$$\theta_t^* = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^{RN}[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^{RN}[X_T e^{-r(T-t)}]}{\partial S_t}, \quad (9)$$

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 (7) with $\rho = \pm 1$). Thus, our dynamic hedge (5) 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 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 1 in the Appendix) between the expected discounted non-tradable asset payoff, $X_T e^{-r(T-t)}$, under the new and original measures:

$$E_t^*[X_T e^{-r(T-t)}] = E_t[X_T e^{-r(T-t)}] - E_t[W_T^* e^{-r(T-t)} - W_t]. \quad (10)$$

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 (10) 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 (10) 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 (6), 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 (5) 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 Basak and Chabakauri (2008) 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.

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.² 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

²A 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 4.2 we show that our model can easily be extended to incorporate additional state variables.

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 oil price dynamics.

In our first example, the non-tradable asset price follows a mean-reverting Ornstein-Uhlenbeck (OU) process:³

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}dw_{Xt}, \quad (11)$$

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:

$$\frac{dS_t}{S_t} = (r + \sigma X_t)dt + \sigma dw_t. \quad (12)$$

According to Proposition 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 1 reports the optimal hedging policy and its corresponding quality.

Corollary 1. *The optimal hedging policy and the corresponding variance of the hedging error for the mean-reverting Gaussian model (11)–(12) are given by*

$$\theta_t^* = \frac{\rho\bar{\nu}}{\sigma}e^{-(r+\lambda+\rho\bar{\nu})(T-t)}, \quad (13)$$

$$\text{var}_t[X_T - W_T^*] = (1 - \rho^2)\bar{\nu}^2 \frac{1 - e^{-2(\lambda+\rho\bar{\nu})(T-t)}}{2(\lambda + \rho\bar{\nu})}. \quad (14)$$

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 2.2. In particular, Corollary 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 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.

³The 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.

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

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}\sqrt{X_t}dw_{xt}, \quad (15)$$

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

$$\frac{dS_t}{S_t} = \mu dt + \sqrt{\frac{1}{X_t}}dw_t. \quad (16)$$

As in the previous example, the explicit hedge follows easily from Proposition 1. Corollary 2 presents the optimal hedge along with the associated variance of the hedging error.

Corollary 2. *The optimal hedging policy and the corresponding variance of the hedging error for the mean-reverting stochastic-volatility model (15)–(16) are given by*

$$\theta_t^* = \rho\bar{\nu}X_t e^{-(r+\lambda+\rho\bar{\nu}(\mu-r))(T-t)}, \quad (17)$$

$$\begin{aligned} \text{var}_t[X_T - W_T^*] &= (1 - \rho^2)\bar{\nu}^2\bar{X} \frac{1 - e^{-2(\lambda+\rho\bar{\nu}(\mu-r))(T-t)}}{2(\lambda + \rho\bar{\nu}(\mu - r))} \\ &+ (1 - \rho^2)\bar{\nu}^2(X_t - \bar{X}) \frac{e^{-\lambda(T-t)} - e^{-2(\lambda+\rho\bar{\nu}(\mu-r))(T-t)}}{\lambda + 2\rho\bar{\nu}(\mu - r)}. \end{aligned} \quad (18)$$

Corollary 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.

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

$$\text{var}_t[X_T - W_T] = E_t[\text{var}_{t+\tau}(X_T - W_T)] + \text{var}_t[E_{t+\tau}(X_T - W_T)], \quad \tau > 0. \quad (19)$$

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 (19), 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 2.

Proposition 2. *Assume that the stock risk premium is zero, $\mu_t - r = 0$. Then the variance criterion (4) is time-consistent and the ensuing optimal dynamic minimum-variance hedging policy is given by*

$$\theta_t^* = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t[X_T e^{-r(T-t)}]}{\partial S_t}. \quad (20)$$

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 (10)) disappear. Consequently, the non-tradable asset, and hence time- t hedge, are not affected by future policies, eliminating the time-inconsistency.⁴ 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 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 (20), which generalizes their explicit hedges derived for martingale and geometric Brownian motion non-tradable asset prices.

Proposition 2 also allows us to convert the minimum-variance hedging problem considered in Section 2.2 to a time-consistent one, as discussed in Corollary 3.

Corollary 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*

$$\frac{dX_t}{X_t} = \left(m_t - \rho\nu_t \frac{\mu_t - r}{\sigma_t} - \sqrt{1 - \rho^2} \eta_t \right) dt + \nu_t dw_{X_t}^\eta, \quad \frac{dS_t}{S_t} = r dt + \sigma_t dw_t^\eta, \quad (21)$$

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

$$\frac{dP^\eta}{dP} = e^{-\frac{1}{2} \int_0^T \left(\left(\frac{\mu_s - r}{\sigma_s} \right)^2 + \eta_s^2 \right) ds - \int_0^T \frac{\mu_s - r}{\sigma_s} dw_s - \int_0^T \eta_s dw_s^\perp}, \quad (22)$$

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

⁴Formally, the first term in the law of total variance (19) 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.

The following minimum-variance criteria

$$\text{var}_t^\eta[X_T - W_T], \quad (23)$$

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

$$\theta_t^\eta = \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^\eta[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^\eta[X_T e^{-r(T-t)}]}{\partial S_t}, \quad (24)$$

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 (24) equals the dynamically optimal hedge (5).

Corollary 3 reveals that a risk-neutral measure adjusts the variance criterion so that it becomes time-consistent. The criterion (23) 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 2. The dynamically optimal hedge (5) is then obtained from the time-consistent hedging problem when $\eta_t = 0$.

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:

$$\Delta_t = \frac{E_0[\text{var}_t(X_T - W_T^{\text{alternative}})]}{E_0[\text{var}_t(X_T - W_T^*)]} - 1, \quad (25)$$

where $W_T^{\text{alternative}}$ denotes the terminal tradable wealth under the alternative policy, $\theta_t^{\text{alternative}}$, considered in Subsections 3.1–3.3.⁵ 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.

⁵All our relative performance results in Sections 3.2–3.3 (Propositions 4–5) remain valid for a more general, conditional relative performance measure given by $\text{var}_t[X_T - W_T^{\text{alternative}}] / \text{var}_t[X_T - W_T^*] - 1$.

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 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:

$$\frac{dX_t}{X_t} = mdt + \nu dw_{Xt}, \quad (26)$$

$$\frac{dS_t}{S_t} = \mu dt + \sigma dw_t, \quad (27)$$

where m , ν , μ and σ are constants. The dynamically optimal hedge is obtained from Proposition 1, while the pre-commitment hedge along the lines of Duffie and Richardson (1991) adapted to our setup. Proposition 3 presents the two policies and a simple condition for the dynamically optimal hedge to outperform.⁶

Proposition 3. *The dynamically optimal and pre-commitment policies for GBM non-tradable asset and stock prices (26)–(27) are given by*

$$\theta_t^* = \frac{\rho\nu}{\sigma} X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)}, \quad (28)$$

$$\begin{aligned} \theta_t^{commit} &= \frac{\rho\nu}{\sigma} X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} \\ &\quad - \frac{\mu-r}{\sigma^2} \left((X_0 e^{(m-r-\rho\nu(\mu-r)/\sigma)T} - W_0) e^{rt} - (X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} - W_t^{commit}) \right). \end{aligned} \quad (29)$$

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 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 (29)). To see why this term arises, we observe that if the hedger follows

⁶Proposition 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 in the Appendix.




Figure 1: Relative Performance of Dynamically Optimal and Pre-commitment Hedges.

The figure plots the relative performance measure Δ_t (equation (25)) 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.

the dynamically optimal policy from time t on, her expected hedging error is (as shown in the Appendix)

$$E_t[(X_T - W_T^*)e^{-r(T-t)}] = X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} - W_t. \quad (30)$$

Hence, the second term in (29) 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 (19)). So, when the second term in (29) 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 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 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). 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 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 (28)–(29)).

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 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

$$W_T = W_0 e^{rT} + \theta_0 (S_T/S_0 - e^{rT}), \quad (31)$$

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 by⁷

$$\theta_t^{static} = \frac{\text{cov}_0(S_T/S_t, X_T)}{\text{var}_0[S_T/S_t]}. \quad (32)$$

⁷Since the hedger holds the same number of units of stock over the horizon, $\theta_t^{static} = (\theta_0^{static}/S_0)S_t$.

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 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 (5) can equivalently be rewritten as

$$\theta_t^* = \frac{\text{cov}_t(dS_t/S_t, dE_t^*[X_T e^{-r(T-t)}])}{\sigma_t^2}. \quad (33)$$

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 (25)) can explicitly be computed. In addition to considering the non-tradable and stock prices following GBMs (26)–(27), we also study the cases of their following ABMs

$$dX_t = \bar{m}dt + \bar{\nu}dw_{xt}, \quad (34)$$

$$dS_t = \bar{\mu}dt + \bar{\sigma}dw_t, \quad (35)$$

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

$$dX_t = \lambda(\bar{X} - X_t)dt + \bar{\nu}dw_{xt}, \quad (36)$$

with \bar{X} and $\lambda > 0$ constants. Proposition 4 reports the results.

Proposition 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 1.*

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

Table 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 (25)) 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 (34)–(35)), geometric Brownian motion (equations (26)–(27)) and Ornstein-Uhlenbeck mean-reverting (equation (36)) 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
$\frac{\rho\nu X_t}{\sigma} e^{(m-r-\rho\nu\frac{\mu-r}{\sigma})(T-t)}$	$\frac{X_0 S_t}{S_0} \frac{e^{\rho\nu\sigma T}-1}{e^{\sigma^2 T}-1} e^{(m-\mu)T}$	$\begin{cases} + & \rho > 0 \\ +/- & \rho < 0 \end{cases}$	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

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. 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

$$\min_{\theta_t} \text{var}_t[dX_t - dW_t], \tag{37}$$

Table 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 (25)) 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 (34)–(35)), geometric Brownian motion (equations (26)–(27)) and Ornstein-Uhlenbeck mean-reverting (equation (36)) 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
$\frac{\rho \bar{\nu} S_t}{\sigma}$	$\frac{\rho \bar{\nu} S_t}{\sigma}$	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
$\frac{\rho \bar{\nu} S_t}{\sigma} \frac{\lambda e^{-\lambda(T-t)} + r e^{r(T-t)}}{\lambda + r}$	$\frac{\rho \bar{\nu} S_t}{\sigma}$	+	OU	ABM
$\frac{\rho \bar{\nu}}{\sigma} e^{-(\lambda+r)(T-t)}$	$\frac{\rho \bar{\nu}}{\sigma}$	+	OU	GBM

subject to the budget constraint (3). 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:

$$\theta_t^{myopic} = \frac{\rho \nu t}{\sigma_t} X_t. \quad (38)$$

The myopic hedge is simply the instantaneous version of the static hedge and is in general different from the dynamically optimal hedge (5). 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 (1) 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 Subsection, we compare the myopic and dynamically optimal hedges for popular price processes, including GBMs (26)–(27), ABMs (34)–(35), and OU (36). Proposition 5 reports the two hedges under these settings, as well as their relative performances.

Proposition 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 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 3.2. 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 Subsection 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.⁸ 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 $[0.15, 0.6]$.

4. Extensions

We now generalize the results on minimum-variance hedging derived in Section 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.

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 Subsection 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 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

⁸More 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.

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

$$\max_{\theta_t} E_t[X_T - W_T] - \frac{\gamma}{2} \text{var}_t[X_T - W_T], \quad (39)$$

subject to the budget constraint (3), 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 (39) evaluated at the optimal policy. As in Section 2 we consider the time-consistent solution to problem (39) obtained by dynamic programming. Proposition 6 reports the dynamically optimal hedging policy along with the value function.

Proposition 6. *The dynamically optimal mean-variance hedge, θ_t^* , and the corresponding value function, J_t , are given by*

$$\begin{aligned} \theta_t^* &= \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial S_t} - \frac{\mu_t - r}{\gamma\sigma_t^2} e^{-r(T-t)} \\ &+ \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*\left[\left(\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds\right) e^{-r(T-t)}\right]}{\partial X_t} + S_t \frac{\partial E_t^*\left[\left(\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds\right) e^{-r(T-t)}\right]}{\partial S_t}, \end{aligned} \quad (40)$$

$$\begin{aligned} J_t &= -\frac{\gamma}{2}(1 - \rho^2) E_t \left[\int_t^T \nu_s^2 X_s^2 \left(\frac{\partial E_s^*[X_T + \int_s^T \frac{1}{\gamma} \left(\frac{\mu_\tau - r}{\sigma_\tau}\right)^2 d\tau]}{\partial X_s} \right)^2 ds \right] \\ &+ E_t^*[X_T e^{-r(T-t)}] - W_t e^{r(T-t)} + \frac{1}{2} E_t^* \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds \right]. \end{aligned} \quad (41)$$

Proposition 6 reveals that the dynamically optimal mean-variance hedge is comprised of three types of terms. The first two terms in (40) comprise the variance-minimizing hedge of Section 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 (40) 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 (40) 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 2.4, the solution from such a criterion does not, in general, coincide with the actual one (40), unlike in the minimum-variance case. The value function (41) 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.

4.2. Multi-dimensional case

We now demonstrate that the results of Section 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 $w_X = (w_{X_1}, \dots, w_{X_N})^\top$ 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_{X_n} and w_k .

There are N non-tradable assets whose prices, $X = (X_1, \dots, X_N)^\top$, follow dynamics

$$\frac{dX_{it}}{X_{it}} = m_i(X_t, t)dt + \nu_i(X_t, t)^\top dw_{Xt}, \quad i = 1, \dots, N, \quad (42)$$

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

$$\frac{dS_{jt}}{S_{jt}} = \mu_j(S_t, X_t, t)dt + \sigma_j(S_t, X_t, t)^\top dw_t, \quad j = 1, \dots, K, \quad (43)$$

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

$$dW_t = [rW_t + \theta_t^\top (\mu_t - r)]dt + \theta_t^\top \sigma_t dw_t. \quad (44)$$

The hedger's dynamic optimization problem is as in Section 2. At each time t , she minimizes the variance of her hedging error, $\phi^\top X_T - W_T$, subject to the budget constraint (44). The

optimal policy is then derived by dynamic programming as in Section 2. Proposition 7 reports the dynamically optimal hedge and its associated quality.

Proposition 7. *The optimal hedging policy and the corresponding variance of hedging error are given by*

$$\theta_t^* = (\nu_t \rho^\top \sigma_t^{-1})^\top I_{X_t} \frac{\partial E_t^*[\phi^\top X_T e^{-r(T-t)}]}{\partial X_t^\top} + I_{S_t} \frac{\partial E_t^*[\phi^\top X_T e^{-r(T-t)}]}{\partial S_t^\top}, \quad (45)$$

$$\text{var}_t[\phi^\top X_T - W_T^*] = E_t \left[\int_t^T \frac{\partial E_s^*[\phi^\top X_T]}{\partial X_s} \nu_s I_{X_s} (I - \rho^\top \rho) I_{X_s} \nu_s^\top \frac{\partial E_s^*[\phi^\top X_T]}{\partial X_s^\top} ds \right], \quad (46)$$

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

$$\begin{aligned} \frac{dX_{it}}{X_{it}} &= \left(m_{it} - \nu_{it}^\top \rho^\top \sigma_t^{-1} (\mu_t - r) \right) dt + \nu_{it}^\top dw_{Xt}^*, \quad i = 1, \dots, N, \\ \frac{dS_{jt}}{S_{jt}} &= r dt + \sigma_{jt}^\top dw_t^*, \quad j = 1, \dots, K, \end{aligned}$$

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 (45) 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 (46) 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 2.

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.

Appendix: Proofs

Proof of Proposition 1. We obtain the optimal hedge (5) by following the methodology in Basak and Chabakauri (2008) and applying dynamic programming to the value function J_t , defined as

$$J(X_t, S_t, W_t, t) \equiv \text{var}_t[X_T - W_T^*]. \quad (\text{A.1})$$

Suppose, the hedger rebalances the portfolio over time intervals τ . The law of total variance (19) substituted into (A.1) yields a recursive representation for the value function:

$$J_t = \min_{\theta_t} E_t[J_{t+\tau}] + \text{var}_t[E_{t+\tau}(X_T - W_T)]. \quad (\text{A.2})$$

We next substitute W_T in (A.2) by its integral form

$$W_T = W_t e^{r(T-t)} + \int_t^T \theta_s (\mu_s - r) e^{r(T-s)} ds + \int_t^T \theta_s \sigma_s e^{r(T-s)} dw_s, \quad (\text{A.3})$$

obtained from the budget constraint (3), 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 (A.2), we obtain the continuous-time HJB equation

$$0 = \min_{\theta_t} E_t[dJ_t] + \text{var}_t[dG_t - d(W_t e^{r(T-t)})], \quad (\text{A.4})$$

with the terminal condition $J_T = 0$, where G_t is defined by

$$G(X_t, S_t, W_t, t) \equiv E_t[X_T - \int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds]. \quad (\text{A.5})$$

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 (A.3) 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 (A.4) and computing the variance term, we obtain the equation

$$\begin{aligned} 0 = \mathcal{D}J_t &+ \nu_t^2 X_t^2 \left(\frac{\partial G_t}{\partial X_t} \right)^2 + 2\rho \nu_t \sigma_t X_t S_t \frac{\partial G_t}{\partial X_t} \frac{\partial G_t}{\partial S_t} + \sigma_t^2 S_t^2 \left(\frac{\partial G_t}{\partial S_t} \right)^2 \\ &+ \min_{\theta_t} \left[\theta_t^2 \sigma_t^2 e^{2r(T-t)} - 2\theta_t \sigma_t \left(\rho \nu_t X_t \frac{\partial G_t}{\partial X_t} + \sigma_t S_t \frac{\partial G_t}{\partial S_t} \right) e^{r(T-t)} \right], \end{aligned} \quad (\text{A.6})$$

subject to $J_T = 0$. The minimization in (A.6) has a unique solution

$$\theta_t^* = \frac{\rho \nu_t}{\sigma_t} X_t \frac{\partial G_t}{\partial X_t} e^{-r(T-t)} + S_t \frac{\partial G_t}{\partial S_t} e^{-r(T-t)}. \quad (\text{A.7})$$

Substituting (A.7) back into (A.6), we obtain the following PDE for the value function

$$\mathcal{D}J_t + (1 - \rho^2) \left(\nu_t X_t \frac{\partial G_t}{\partial X_t} \right)^2 = 0, \quad (\text{A.8})$$

with the terminal condition $J_T = 0$. The Feynman-Kac solution (Karatzas and Shreve, 1991) to equation (A.8) is then given by

$$J_t = (1 - \rho^2) E_t \left[\int_t^T \left(\nu_s X_s \frac{\partial G_s}{\partial X_s} \right)^2 ds \right]. \quad (\text{A.9})$$

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 (A.5), we obtain a PDE for G_t . Substituting θ_t^* from (A.7) 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

$$G_t = E_t^*[X_T], \quad (\text{A.10})$$

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 (7). Substituting (A.10) into (A.7) and (A.9) yields the optimal hedge (5) and the hedging error variance (6), 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 (8).

Finally, we derive the representation (10) for $E_t^*[X_T e^{-r(T-t)}]$ by first taking the expectation of (A.3)

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

and then substituting (A.10) and (A.11) into (A.5).

Proof of Corollary 1. Under the probability measure P^* , the process (11) becomes

$$dX_t = (\lambda + \rho \bar{\nu}) \left(\frac{\lambda \bar{X}}{\lambda + \rho \bar{\nu}} - X_t \right) dt + \bar{\nu} dw_{Xt}^*, \quad (\text{A.12})$$

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 1 yields the desired expressions (13)–(14).

Q.E.D.

Proof of Corollary 2. Under the probability measure P^* , the process (15) follows dynamics

$$dX_t = \left(\lambda + \rho\bar{\nu}(\mu - r) \right) \left(\frac{\lambda\bar{X}}{\lambda + \rho\bar{\nu}(\mu - r)} - X_t \right) dt + \bar{\nu}\sqrt{X_t}dw_{xt}^*. \quad (\text{A.13})$$

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 1 yields (17)–(18).

Q.E.D.

Proof of Proposition 2. First, we derive a variation of the law of total variance. From the law of total variance (19) with an infinitesimally small interval τ , we obtain the following equality in differential form:

$$0 = E_t \left[d \text{var}_s(X_T - W_T) + \text{var}_s(dE_s[X_T - W_T]) \right]. \quad (\text{A.14})$$

Integrating (A.14) from t to T yields

$$\text{var}_t[X_T - W_T] = E_t \left[\int_t^T \frac{\text{var}_s(dE_s[X_T - W_T])}{ds} ds \right]. \quad (\text{A.15})$$

From the assumption $\mu_t - r = 0$ and the integrated budget constraint (A.3), it follows that $E_t[W_T] = W_t e^{r(T-t)}$. Hence, by Itô's lemma

$$dE_t[X_T - W_T] = (\dots)dt + \nu_t X_t \frac{\partial E_t[X_T]}{\partial X_t} dw_{xt} + \sigma_t S_t \frac{\partial E_t[X_T]}{\partial S_t} dw_t - \theta_t \sigma_t e^{r(T-t)} dw_t. \quad (\text{A.16})$$

Substituting (A.16) into (A.15) and computing $\text{var}_s(dE_t[X_T - W_T])$, we obtain:

$$\text{var}_t[X_T - W_T] = E_t \left[\int_t^T \left(\theta_s \sigma_s e^{r(T-s)} - \rho \nu_s X_s \frac{\partial E_s[X_T]}{\partial X_s} - \sigma_s S_s \frac{\partial E_s[X_T]}{\partial S_s} \right)^2 + (1 - \rho^2) \left(\frac{\partial E_s[X_T]}{\partial X_s} \right)^2 ds \right]. \quad (\text{A.17})$$

Minimizing the expression under the integral in (A.17) gives the global minimum to the variance criterion, yielding the hedge (20). Finally, we observe that for $\mu_t - r = 0$, the dynamically optimal hedge (5) coincides with the hedge (20) since the Radon-Nikodym derivative (8) equals unity, and hence the the variance criterion is time-consistent.

Q.E.D.

Proof of Corollary 3. The hedging criterion (23) can be represented in integral form (A.15) in which all the expectations and variances are under the measure P^η (8). 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 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 (23) is time-consistent and the solution is given by (5).

Q.E.D.

Proof of Proposition 3. We first compute the optimal hedges and hedging error variances, and then derive the properties of the performance measure Δ_t . From Proposition 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

$$E_t^*[X_T] = X_t e^{(m - \rho\nu(\mu - r)/\sigma)(T-t)}. \quad (\text{A.18})$$

Substituting (A.18) into Proposition 1, we obtain the dynamically optimal hedge (28) and the associated hedging error variance

$$\text{var}_t[X_T - W_T^*] = (1 - \rho^2)\nu^2 X_t^2 e^{2(m - \rho\nu\frac{\mu-r}{\sigma})(T-t)} \frac{e^{(\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma})(T-t)} - 1}{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma}}. \quad (\text{A.19})$$

The optimal pre-commitment hedge (29) for the case of $r = 0$ and $W_0 = 0$ has been obtained by Duffie and Richardson (1991) in the context of futures hedging.⁹ To obtain it for our case of $r > 0$ and $W_0 > 0$, we observe that the budget constraint (3) can equivalently be rewritten as

$$d\tilde{W}_t = \theta_t \tilde{\mu}_t dt + \theta_t \tilde{\sigma}_t dw_t, \quad (\text{A.20})$$

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 (A.20) reduces to the case with $r = 0$ and $\tilde{W}_0 = 0$, and hence the pre-commitment hedge (29) 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

$$H_t \equiv X_t e^{(m - r - \rho\nu\frac{\mu-r}{\sigma})(T-t)} - W_t^{\text{commit}} e^{r(T-t)}. \quad (\text{A.21})$$

Substituting the pre-commitment hedge (29) into the budget constraint (3) and applying Itô's lemma to H_t we obtain:

$$dH_t = \left(\frac{\mu - r}{\sigma}\right)^2 (H_0 - H_t) dt + \frac{\mu - r}{\sigma} (H_0 - H_t) dw_t + \sqrt{1 - \rho^2} \nu X_t e^{(m - \rho\nu\frac{\mu-r}{\sigma})(T-t)} dw_t^\perp. \quad (\text{A.22})$$

Integrating (A.22) 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

$$\begin{aligned} E_t[(H_\tau - H_0)^2] &= (H_t - H_0)^2 e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (\tau-t)} \\ &+ (1 - \rho^2)\nu^2 X_t^2 e^{2(m - \rho\nu\frac{\mu-r}{\sigma})(\tau-t)} \frac{e^{(\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma})(\tau-t)} - e^{-\left(\frac{\mu-r}{\sigma}\right)^2 (\tau-t)}}{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + \left(\frac{\mu-r}{\sigma}\right)^2}. \end{aligned}$$

⁹For 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.

Given the first two moments of H_t and taking into account that $H_T = X_T - W_T^{commit}$, we obtain:

$$\begin{aligned} \text{var}_t[X_T - W_T^{commit}] &= (H_t - H_0)^2 e^{-\left(\frac{\mu-r}{\sigma}\right)^2(T-t)} \left(1 - e^{-\left(\frac{\mu-r}{\sigma}\right)^2(T-t)}\right) \\ &+ (1 - \rho^2) \nu^2 X_t^2 e^{2(m-\rho\nu\frac{\mu-r}{\sigma})(T-t)} \frac{e^{(\nu^2+2\rho\nu\frac{\mu-r}{\sigma})(T-t)} - e^{-\left(\frac{\mu-r}{\sigma}\right)^2(T-t)}}{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + \left(\frac{\mu-r}{\sigma}\right)^2}. \end{aligned} \quad (\text{A.23})$$

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 $E_0[\text{var}_t(X_T - W_T^{commit})]$. The relative performance measure (25) is then given by:

$$\begin{aligned} \Delta_t &= \frac{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma}}{\nu^2 + 2\rho\nu\frac{\mu-r}{\sigma} + \left(\frac{\mu-r}{\sigma}\right)^2} e^{-\left(\frac{\mu-r}{\sigma}\right)^2(T-t)} \left(\frac{e^{(\nu^2+2\rho\nu\frac{\mu-r}{\sigma} + \left(\frac{\mu-r}{\sigma}\right)^2)(T-t)} - 1}{e^{(\nu^2+2\rho\nu\frac{\mu-r}{\sigma})(T-t)} - 1} \right. \\ &+ \left. \frac{1 - e^{-(\nu^2+2\rho\nu\frac{\mu-r}{\sigma} + \left(\frac{\mu-r}{\sigma}\right)^2)t}}{e^{(\nu^2+2\rho\nu\frac{\mu-r}{\sigma})(T-t)} - 1} \left(1 - e^{-\left(\frac{\mu-r}{\sigma}\right)^2(T-t)}\right) \right) - 1. \end{aligned} \quad (\text{A.24})$$

Letting t go to T in (A.24), 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 (30) by rearranging terms in (10) and substituting (A.18). *Q.E.D.*

Proof of Proposition 4. First, we consider the case when both risky assets follow ABMs (34)–(35). From Proposition 1, the process for X under the measure P^* is given by

$$dX_t = (\bar{m} - \rho\bar{\nu}\frac{\bar{\mu} - rS_t}{\bar{\sigma}})dt + \bar{\nu}dw_{Xt}^*.$$

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 1 yields the optimal hedge reported in Table 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 (32), we obtain the static hedge, which coincides with the dynamic one, and hence $\Delta_t = 0$.

When the risky asset prices follow GBMs (26)–(27), the dynamically optimal hedge and its corresponding hedging error variance are given by (28) and (A.19), respectively. The static hedge reported in Table 1 is obtained from (32) by observing that since X_t^2 , S_t^2 and X_tS_t follow GBMs,

$$\text{var}_t[X_T] = X_t e^{2m(T-t)} (e^{\nu^2(T-t)} - 1), \quad \text{var}_t[S_T] = S_t e^{2\mu(T-t)} (e^{\sigma^2(T-t)} - 1), \quad (\text{A.25})$$

$$\text{cov}_t(X_T, S_T) = X_t S_t e^{2(m+\mu)(T-t)} (e^{\rho\nu\sigma(T-t)} - 1). \quad (\text{A.26})$$

Substituting W_T^{static} from the static budget constraint (31) into the hedging error variance we obtain

$$\text{var}_t[X_T - W_T^{static}] = \text{var}_t[X_T] - 2\frac{\theta_0^{static}}{S_0} \text{cov}_t(X_T, S_T) + \left(\frac{\theta_0^{static}}{S_0}\right)^2 \text{var}_t[S_T]. \quad (\text{A.27})$$

We now show that for $\rho > 0$ the performance measure (25) is positive in this GBMs case. We note that the static hedging error variance (A.27) is a quadratic function of θ_0 , the minimization of which along with the expressions (A.25)–(A.26) gives the lower bound for the static hedging error variance:

$$\text{var}_t[X_T - W_T^{static}] \geq X_t^2 e^{2m(T-t)} \left(e^{\nu^2(T-t)} - 1 - \frac{(e^{\rho\nu\sigma(T-t)} - 1)^2}{e^{\sigma^2(T-t)} - 1} \right). \quad (\text{A.28})$$

We next rewrite the dynamically optimal hedging error variance (A.19) in integral form and find its upper bound for $\rho > 0$ as:

$$\begin{aligned} \text{var}_t[X_T - W_T^*] &= X_t^2 (1 - \rho^2) \nu^2 e^{2m(T-t)} \int_t^T e^{-2\rho\nu\frac{\mu-r}{\sigma}(T-s)} e^{\nu^2(s-t)} ds \\ &\leq X_t^2 (1 - \rho^2) \nu^2 e^{2m(T-t)} \int_t^T e^{\nu^2(s-t)} ds = X_t^2 (1 - \rho^2) e^{2m(T-t)} (e^{\nu^2(T-t)} - 1). \end{aligned} \quad (\text{A.29})$$

A sufficient condition for the dynamically optimal variance to be lower than the static one is that the upper bound in (A.29) is below the lower bound in (A.28), which is equivalent to

$$\left(\frac{e^{\rho\nu\sigma(T-t)} - 1}{\rho\nu\sigma} \right)^2 \leq \left(\frac{e^{\sigma^2(T-t)} - 1}{\sigma^2} \right) \left(\frac{e^{\nu^2(T-t)} - 1}{\nu^2} \right). \quad (\text{A.30})$$

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:

$$\left(\int_t^T e^{\rho\nu\sigma(T-s)} ds \right)^2 \leq \left(\int_t^T e^{(\frac{\nu^2}{2} + \frac{\sigma^2}{2})(T-s)} ds \right)^2 \leq \left(\int_t^T e^{\sigma^2(T-s)} ds \right) \left(\int_t^T e^{\nu^2(T-s)} ds \right). \quad (\text{A.31})$$

Computing the integrals in (A.31) we obtain inequality (A.30), 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, (A.19) and (A.27), into the performance measure (25), 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 $(\mu - r)/\sigma$.¹⁰

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 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.*

¹⁰For 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.

Proof of Proposition 5. The dynamically optimal hedges reported in Table 2 are the same as in Table 1, while the myopic hedges are immediate from the expression (38). 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 (A.15) and compare it with the dynamically optimal one. From the budget constraint in integral form (A.3) 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 (A.15), and after some algebra determine the myopic hedging error variance:

$$\begin{aligned} \text{var}_t[X_T - W_T^{myopic}] &= X_t^2 \nu^2 e^{2m(T-t)} \int_t^T e^{\nu^2(s-t)} \left((1 - \rho^2) \left(1 - \rho \nu \frac{\mu - r}{\sigma} \frac{1 - e^{-(m-r)(T-s)}}{m-r} \right)^2 \right. \\ &\quad \left. + \rho^2 \nu^2 \left(\frac{m-r}{\nu} - \rho \frac{\mu - r}{\sigma} \right)^2 \left(\frac{1 - e^{-(m-r)(T-s)}}{m-r} \right)^2 \right) ds. \end{aligned} \quad (\text{A.32})$$

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 (A.29) and (A.32) 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 (6) and (A.15). 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 6. The proof is similar to the proof of Proposition 1. The hedging problem is solved via dynamic programming and the value function is defined as:

$$J(X_t, S_t, W_t, t) = E_t[X_T - W_T^*] - \frac{\gamma}{2} \text{var}_t[X_T - W_T^*]. \quad (\text{A.33})$$

Applying the law of total variance along the same steps as in the proof of Proposition 1, we obtain an HJB equation. To solve this equation, substituting the budget constraint in integral form (A.3) into the hedger's objective (39), 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 1. Solving the PDE for J_t , we obtain the value function (41). *Q.E.D.*

Proof of Proposition 7. Proposition 7 is a multidimensional version of Proposition 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

- Adler, M., and J.B. Detemple, 1988, "On the Optimal Hedge of a Nontraded Cash Position," *Journal of Finance*, 43, 143-153.
- 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.
- Basak, S., and G. Chabakauri, 2008, "Dynamic Mean-Variance Asset Allocation," Working Paper, London Business School.
- Bertsimas, D., L. Kogan, and A. Lo, 2001, "Hedging Derivative Securities and Incomplete Markets: An ϵ -Arbitrage Approach," *Operations Research*, 49, 372-397.
- Bielecki, T.R., M. Jeanblanc, and M. Rutkowski, 2004, "Mean-Variance Hedging of Defaultable Claims," Working Paper, Illinois Institute of Technology.
- Brandt, M.W., 2003, "Hedging Demands in Hedging Contingent Claims," *Review of Economics and Statistics*, 85, 119-140.
- Breeden, D.T., 1984, "Futures Markets and Commodity Options: Hedging and Optimality in Incomplete markets," *Journal of Economic Theory*, 32, 275-300.
- Brooks, C., O.T. Henry, and G. Persaud, 2002, "The Effect of Asymmetries on Optimal Hedge Ratios," *Journal of Business*, 75, 333-352.
- 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.
- 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., J. Ingersoll, and S. Ross, 1985, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385-408.
- Cvitanic, J., and F. Zapatero, 2004, *Introduction to the Economics and Mathematics of Financial Markets*, First Edition, The MIT Press.
- 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.
- Ederington, L.H., 1979, "The Hedging Performance of the New Futures Markets," *Journal of Finance*, 34, 157-170.
- 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.
- Gibson, R., and E.S. Schwartz, 1990, "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims," *Journal of Finance*, 45, 959-976.
- 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.
- Henderson, V., 2005, "Explicit Solutions to an Optimal Portfolio Choice Problem with Stochastic Income," *Journal of Economic Dynamics and Control*, 29, 1237-1266.
- Hirshleifer, D., 1988, "Residual Risk, trading Costs, and Commodity Futures Risk Premia," *Review of Financial Studies*, 1, 173-193.
- 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.
- Jaillet, P., E.I. Ronn, and S. Tompaidis, 2004, "Valuation of Commodity-Based Swing Options," *Management Science*, 50, 909-921.
- 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., and S.E. Shreve, 1991, *Brownian Motion and Stochastic Calculus*, Second Edition, Springer-Verlag.
- 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.
- Kim, T.S., and E. Omberg, 1996, "Dynamic Nonmyopic Portfolio Behavior," *Review of Financial Studies*, 9, 141-161.

- 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.
- Lioui, A., and P. Poncet, 2000, "The Minimum Variance Hedge Ratio Under Stochastic Interest Rates," *Management Science*, 46, 658-668.
- 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.
- Musiela, M., and M. Rutkowski, 1998, *Martingale Methods in Financial Modelling*, Springer-Verlag.
- 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.
- 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., 1994, "Approximating Random Variables by Stochastic Integrals," *Annals of Probability*, 22, 1536-1575.
- 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.
- Siegel, D.R., and D.F. Siegel, 1990, *The Futures Markets: Arbitrage, Risk Management and Portfolio Choice*, McGraw Hill.
- 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., 2000, "Optimal Hedging and Valuation of Non-traded Assets," *European Economic Review*, 4, 231-251.
- Viceira, L.M., 2001, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income," *Journal of Finance*, 56, 433-470.
- Vasicek, O., 1977, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177-188.