

**Dynamic Hedging in Incomplete Markets:
A Simple Solution**

By

**Suleyman Basak
Georgy Chabakauri**

**THE PAUL WOOLLEY CENTRE
WORKING PAPER SERIES NO 23**

**FINANCIAL MARKETS GROUP
DISCUSSION PAPER NO 680**

May 2011

Suleyman Basak is the Class of 2008 Term Chair Professor of Finance at London Business School. Prior to working at London Business School, he was at the Wharton School of the University of Pennsylvania, held a visiting position at the Graduate School of Business at the University of Chicago, and acted as a consultant to Goldman, Sachs & Co. He is listed in Who's Who in the World and Who's Who in Finance and Industry. His research focuses on asset pricing, risk management, market imperfections, international finance, financial innovation, and includes topics such as portfolio insurance, VaR-based risk management, credit risk, tax arbitrage and mispricing in financial markets. Dr. Georgy Chabakauri joined the Finance Department at the London School of Economics and Political Science as a Lecturer in Finance in September, 2009. His research interests include asset pricing in the presence of market imperfections, portfolio choice theory and risk management. He holds PhD in Finance from London Business School (2009), PhD in Mathematics from Moscow State University (2004), and Master of Arts in Economics from New Economic School (2003). Any opinions expressed here are those of the authors and not necessarily those of the FMG. The research findings reported in this paper are the result of the independent research of the authors and do not necessarily reflect the views of the LSE.

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 School of Economics
Department of Finance and FMG
Houghton Street
London WC2A 2AE
United Kingdom
Tel: (44) 20 7107 5374
Fax: (44) 20 7849 4647
G.Chabakauri@lse.ac.uk

This revision: May 2011

*We are grateful to Pietro Veronesi (the editor) and an anonymous referee for valuable suggestions, and to Ulf Axelson, Mike Chernov, Francisco Gomes, Alfredo Ibanez, Stavros Panageas, Marcel Rindisbacher and the seminar participants at European Finance Association Meetings, Society of Economic Dynamics Meetings, London Business School and London School of Economics 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 or a contingent claim. 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 apply our results to the discrete hedging problem of derivatives when trading occurs infrequently. We determine the corresponding optimal hedge and replicating portfolio value, and show that they have structure similar to their complete-market counterparts and reduce to generalized Black-Scholes expressions when specialized to the Black-Scholes setting. We also generalize our results to richer settings to study dynamic hedging with Poisson jumps, stochastic correlation and portfolio management with benchmarking.

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

Keywords: Hedging, incomplete markets, minimum-variance criterion, risk management, time-consistency, discrete hedging, derivatives, benchmarking, correlation risk, Poisson jumps.

1. Introduction

Perfect hedging is a risk management activity that aims at eliminating 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., Cvitanić 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 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 solve the hedging problem by employing the traditional minimum-variance criterion and provide tractable dynamically optimal hedges in a general incomplete-market economy. 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, structured products. 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 (2010). 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 also contrast our hedges to those considered in the literature that minimize the hedging error variance sitting at an initial date. We demonstrate that 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.

We next apply our results to study the long-standing problem of discrete hedging in finance, namely the replication and valuation of derivatives when trading occurs infrequently at discrete periods of time, rendering markets incomplete. Although this problem has received considerable attention, primarily within the Black and Scholes (1973) model context (e.g., Boyle and Emanuel, 1980; Leland, 1985; Bossaerts and Hillion, 1997, 2003, among others), analytical optimal hedges are generally not available.¹ We here provide, in a general discrete-time setting, tractable analytical characterizations of the optimal minimum-variance hedges and the corresponding values of replicating portfolios, which have structure analogous to those in complete markets. This reinforces our earlier results that dynamically optimal minimum-variance hedges extend naturally no-arbitrage-based risk management in complete markets to incomplete markets. Specializing our economy to the Black-Scholes setting, we derive closed-form expressions for the optimal discrete hedge and the replication value of a European call option, which to our knowledge are the first in the literature. These expressions involve a generalized Black-Scholes formula with risk adjustments that account for the inability to trade frequently. We then explore the implications of stock-return predictability. We demonstrate that optimal discrete hedges and replication values are well approximated by their Black-Scholes counterparts absent predictability, consistently with the literature. With predictability, however, infrequent trading considerably impairs the ability of Black-Scholes hedges to approximate discrete hedges and the Black-Scholes prices to predict the market prices of options.

We generalize our results to richer settings to further study dynamic hedging in incomplete markets. We first extend our baseline analysis 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

¹This problem has also been employed in various applications, including in evaluating the performance of competing option pricing models, in developing tests for the existence and the sign of the volatility risk premium, in explaining option bid-ask spreads (e.g., Jameson and Wilhelm, 1992; Bakshi, Cao and Chen, 1997; Bakshi and Kapadia, 2003).

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. We further demonstrate that our main baseline results can easily be extended to feature multiple non-tradable assets and stocks. We here provide an application to study dynamic hedging in the presence of correlation risk by developing a multi-asset model with a flexible correlation structure, which also nests popular stochastic volatility models as special cases. We derive the optimal hedges in closed form, identify the economic channels through which time-varying correlations matter, and disentangle the effects of correlation and volatility risks. Finally, we generalize our basic framework and provide analytical optimal hedges when the assets follow processes with Poisson jumps. We demonstrate that the optimal hedges can easily be computed in specific stochastic settings with jumps.

1.1. Further Discussion and Related Literature

The subject of hedging is, of course, prevalent in the literature on derivatives and risk management. In economic contexts for which there is primarily a pure hedging motive, the goal of minimizing tracking error variance is a natural criterion, as adopted in much work in finance (references discussed below). Such motives arise for market-makers of derivatives who provide immediacy, permitting buyers and sellers to trade whenever they wish. Market-makers select their inventory based on anticipated customer order flow and do not act on personal preference, nor do they expect to profit by speculating, in contrast to proprietary traders (McDonald, 2006). More generally, many intermediary companies in financial services take positions to satisfy client demand for derivatives. Since such positions face possible losses, the companies then aim to eliminate or reduce the risks from derivatives positions sold to clients as much as possible (Cvitanić and Zapatero, 2004). In the context of corporate hedging, such aims to eliminate risk may also be reasonable for a firm in avoiding costly financial distress or mitigating managerial incentives to decline profitable projects if unhedged project risks adversely affect managerial compensation or job security (Duffie and Richardson, 1991). Beyond finance, the minimum-variance criterion also arises as a reasonable objective in various contexts in economics. In particular, when the aim of a government or a central bank is stability and to decrease uncertainty in the economy, such a goal is widely adopted for inflation and output targeting in monetary economics (references discussed below).

Intuitively, the idea of a minimum-variance hedge is to choose a position much like a delta neutral hedge in complete markets, except that not all risk can be eliminated in incomplete markets. So, the best one can hope is to eliminate as much risk as possible so as to replicate target payoff as closely as possible. One simple approximate approach to pure hedging and replication of derivative payoffs in the absence of speculative motives adopted in the literature is to apply complete-market deltas in incomplete-market settings (e.g., Boyle and Emanuel, 1980; Bakshi, Cao, and Chen, 1997; Bertsimas, Kogan and Lo, 2000; Bakshi and Kapadia, 2003; Hull, 2008). However, in contrast to minimum-variance hedges, complete-market hedges being applied in incomplete-market settings do not optimally account for market incompleteness, and hence are sub-optimal. Consequently, the

goal of reducing tracking error variance, which allows replication of target payoffs as closely as possible, can be viewed as a natural generalization of no-arbitrage-based hedging and replication in complete markets to incomplete markets. This observation is supported by the main findings in this paper (as discussed above in results), which characterizes the dynamically optimal time-consistent minimum-variance hedges, in contrast to the existing literature which employs either static, myopic or dynamically time-inconsistent minimum-variance hedges.

All major derivatives and risk management textbooks, Duffie (1989), Siegel and Siegel (1990), Stulz (2003), Cvitanić and Zapatero (2004), McDonald (2006), Hull (2008), present the classic static hedging problem by employing the minimum-variance criterion, and demonstrate its usefulness for real-life risk management applications. Here, the standard cross-hedging problem of basis risk, using a futures contract on an asset to hedge another asset, is also presented. This is a special case of our general hedging problem in incomplete markets, where our tradable asset is the futures contract and our hedging error variance is the basis risk. 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. In a further application, Bakshi, Cao and Chen (1997) use the minimum-variance objective in evaluating the empirical performance of various option pricing and hedging models. Kroner and Sultan (1993), Lioui and Poncet (2000), Brooks, Henry and Persaud (2002) study the performance and economic implications of closely related myopic minimum-variance hedges.

Furthermore, in monetary economics, Sargent and Wallace (1975), Bernanke and Woodford (1997), Svensson (2000), and Evans and McGough (2007) employ the minimum-variance criterion to examine inflation and output targeting. In fiscal theory, Cochrane (2001) appeals to such a criterion for the analysis of long-term government debt and its impact on inflation. Building on the minimum-variance problem in finance, Anderson and Danthine (1980, 1981) study futures hedging and Hirshleifer (1988) derives futures risk premia under transaction costs by considering a static mean-variance hedger. Roll (1992), Chan, Karceski and Lakonishok (1999), Costa and Paiva (2002), Jorion (2003), Gomez and Zapatero (2003), Cornell and Roll (2005) employ the static mean-variance criteria to study portfolio management with tracking error. In this literature, the hedger either cannot rebalance or is myopic. 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.

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 geometric Brownian motions (GBMs). Schweizer (1994) and Pham, Rheinlander and Schweizer (1998) in a more general stochastic environment obtain a recursive feedback representation for the optimal policy. Gouriéroux, Laurent and Pham (1998) derive hedges which are 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.²

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. In this literature, however, the variance-minimizing hedges have not generally been obtained explicitly, with the notable exception of the Duffie and Richardson case of both risky assets following GBMs. 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.

Finally, another strand of work investigates optimal dynamic hedges consistent with a hedger's utility maximization in typically continuous-time incomplete market settings. However, this work has had mixed success in terms of tractability. Breeden (1984) provides optimal hedges with futures in terms of the value function for a general utility function over intertemporal consumption. Stultz (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. 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 only. Svensson and Werner (1993), Tepla (2000) and Henderson (2005) study the optimal hedging of non-tradable income or assets. They obtain explicit solutions in incomplete markets with constant relative risk aversion (CARA) hedgers, GBM tradable asset prices and an income process following a GBM or arithmetic Brownian motion (ABM).³ 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 derives the solution in a feedback form, while the latter work derives a log-linear approximation for the optimal hedges. Detemple and Sundaresan (1999) study the optimal hedging of non-tradable assets numerically in a setting with CRRA utility, binomial asset

²While the mean-variance hedging literature primarily deals with criteria that treat gains and losses symmetrically, it is also reasonable to consider criteria that put more weight on losses than gains. One such criterion considered in the literature is the minimization of expected losses (e.g., Cvitanić, 2000). However, the loss minimization is highly intractable and no closed-form solutions for the optimal hedges have yet been identified in the literature.

³For more general processes or utilities these works derive solutions only in terms of value functions or sensitivities of tradable wealth with respect to asset and state prices.

processes and portfolio constraints.

The paper is organized as follows. In Section 2, we determine the dynamically optimal minimum-variance hedges via dynamic programming, explicitly compute these hedges in plausible environments, and present the time-consistency conditions. In Section 3, we apply our results to study the discrete hedging problem, while in Section 4 we generalize our baseline analysis to incorporate mean-variance hedging, multiple assets, and asset prices with jumps. 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 derivative security, real option, 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. We 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 related to 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

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

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

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

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, defined over the *hedging error* $X_T - W_T$, that determines the quality of hedging.

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

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

subject to the dynamic budget constraint (3). The variance-minimizing criterion has widely been employed in the literature in various economic contexts (e.g., risk management, evaluation of option pricing models, monetary economics), as discussed in Section 1.1. However, in a dynamic setting, the variance-minimizing hedges have not been obtained explicitly in general stochastic environments. Moreover, the literature only characterizes the pre-commitment dynamic hedges that minimize the variance criterion sitting at an initial date but become suboptimal at future dates due to the time-inconsistency of the variance criterion, as discussed in Section 2.4. We solve problem (4) 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 (2010) 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

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

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

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

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 $\rho = \pm 1$). No-arbitrage eliminates the uncertainty in hedging error in complete markets while the minimum-variance criterion reduces it as much as possible when the market is incomplete. 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. When explicit solutions are not available the optimal hedges can be expressed in terms of Malliavin derivatives, which can easily be computed via Monte Carlo simulation following the approach in Detemple, Garcia and Rindisbacher (2003) and Detemple and Rindisbacher (2005).

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.⁴

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. Therefore, it can conveniently be interpreted as a measure of market incompleteness. 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 optimal policy is decreasing in the absolute value of the correlation, $|\rho|$. Intuitively, for higher absolute correlation more wealth is allocated to the stock as the policy 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.

⁴We note that the hedge-neutral value is not an indifference price studied in the literature (e.g. Detemple and Sundaresan, 1999; Henderson, 2005) and defined as the minimal increment to tradable wealth W_t sufficient to induce the investor to give up the non-tradable asset. The utility function-based indifference pricing is vacuous in our framework since the variance criterion is not interpreted as a utility function, and furthermore, the time- t value function (6) does not depend on tradable wealth W_t .

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 (2010) 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.

Remark 2 (Comparison with static policy). We here compare our dynamically optimal hedge (5) with the classic *static* hedge widely employed 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 in empirical works (e.g., Ederington, 1979; Rolfo, 1980; Figlewski, 1984; Kamara and Siegel, 1987; Kerkvliet and Moffett, 1991; In and Kim, 2006). A static hedge θ_0^{static} minimizes the hedging error variance at the initial date, and holds the number of units of stock θ_0^{static}/S_0 constant throughout the hedging horizon. It can be demonstrated that our dynamically optimal and static hedges can be written as follows:

$$\theta_t^* = \frac{\text{cov}_t(dS_t/S_t, dE_t^*[X_T e^{-r(T-t)}])}{\sigma_t^2}, \quad \theta_t^{static} = \frac{\text{cov}_0(S_T/S_t, X_T)}{\text{var}_0[S_T/S_t]}.$$

Clearly, the dynamic hedge inherits the basic intuitive structure of the static hedge, but now tracks the comovement between the instantaneous stock return and the change in the hedge-neutral asset payoff value, and so additionally captures the arrival of new information. One important difference between the two hedges is that the static hedge in general does not provide a perfect hedge, even in dynamically complete markets when $\rho^2 = 1$, in contrast to the dynamic one.⁵

⁵Similarly, our dynamically optimal hedge can be compared with the myopic minimum-variance hedge, which can be viewed as the static hedge over an infinitesimally small hedging horizon, repeated over time. The myopic hedge at each point in time minimizes the variance of the hedging error over next instant, $dX_t - dW_t$, and is given by $\theta_t^{myopic} = \rho \nu_t X_t / \sigma_t$. We note that the myopic hedge ignores the potential impact of mean-returns on the hedging error variance, and in general does not provide a perfect hedge even in complete markets.

2.3. Examples

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, and consider two examples, one with stochastic stock mean return and another with stochastic stock return volatility alone.⁶ We consider the case when the terminal payoff cannot be perfectly hedged by futures contracts either because there are no futures maturing at time T or available contracts are only for similar commodities, and hence do not provide perfect hedges. In both examples, the non-tradable asset price follows a mean-reverting process, which is consistent with the empirical evidence on oil and other commodity prices. For example, Schwartz (1997) and Schwartz and Smith (2000) provide supporting evidence for Gaussian mean-reverting logarithmic commodity prices, while Dixit and Pindyck (1994) and Pindyck (2004) employ a geometric Ornstein-Uhlenbeck process to model and estimate oil price dynamics.

In our first example, the non-tradable asset price follows a mean-reverting Ornstein-Uhlenbeck (OU) process:⁷

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

with $\kappa > 0$. 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*

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

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

⁶A 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.

⁷The 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 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, κ . This is intuitive since a higher speed of convergence to the mean leads to a lower variance of the non-tradable asset payoff, and hence a smaller hedge. The hedging quality also improves as the degree of market completeness, captured by ρ^2 , increases. Moreover, the hedging quality is higher for a positive correlation than for a negative one of the same magnitude since positively correlated stock better tracks the non-tradable asset price.

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

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

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

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

In this setting, the non-tradable asset can also be interpreted as a derivative security tracking the stock return volatility. 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*

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

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

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 κ . 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.

We note that due to the time-inconsistency of the variance criterion our dynamically optimal hedges can be interpreted as the outcome of the sub-game perfect Nash equilibria in an intra-personal game in which current and future selves of a hedger are different players. The time- t self cannot commit future selves to follow the time- t hedging error variance minimizing strategies and hence chooses the time- t strategy as a best response to the strategies of future selves. This game-theoretic interpretation is commonly emphasized in the growing literature that employs time-inconsistent preferences in economic settings with non-exponential discounting and changing preferences (e.g., Peleg and Yaari, 1973; Harris and Laibson, 2001; Grenadier and Wang, 2007).

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.⁸ 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

⁸Formally, 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.

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*

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

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

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

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

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 $\eta_t = 0$.

Remark 3 (Comparison with Duffie and Richardson pre-commitment hedge). As discussed above, the hedge 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),

becomes suboptimal at future dates, unless the hedger can pre-commit to follow it. Accordingly, we refer to it as the *pre-commitment* hedge. We next compare the dynamically optimal hedge (5) with the pre-commitment hedge θ_t^{commit} , as derived in Duffie and Richardson (1991) in a setting with the risky asset dynamics (1)–(2) being GBMs with $m_t = m$, $\nu_t = \nu$, $\mu_t = \mu$, and $\sigma_t = \sigma$.⁹ As demonstrated in the Appendix, in this setting the two hedges are given by:

$$\theta_t^* = \frac{\rho\nu}{\sigma} X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)}, \quad (25)$$

$$\begin{aligned} \theta_t^{commit} &= \frac{\rho\nu}{\sigma} X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} \\ &- \frac{\mu-r}{\sigma^2} \left((X_0 e^{(m-r-\rho\nu(\mu-r)/\sigma)T} - W_0) e^{rt} - (X_t e^{(m-r-\rho\nu(\mu-r)/\sigma)(T-t)} - W_t^{commit}) \right). \end{aligned} \quad (26)$$

In contrast to the dynamically optimal hedge being a simple generalization of the complete-market hedge, the pre-commitment hedge inherits an additional stochastic term. As shown in the Appendix, this additional term allows the hedger to maintain low time-zero hedging error variance. Therefore, sitting at time t , the pre-committed hedger behaves so as to maintain a low time-zero rather than time- t hedging error variance. It can also be demonstrated that the dynamically optimal hedge achieves a lower hedging error variance than the pre-commitment hedge, after a certain period of time.

3. Discrete Hedging

In this Section, we apply the methodology and results of Section 2 to solve a long-standing problem in finance: the replication and valuation of a derivative security when trading can only occur at discrete points in time. With complete financial markets and stock prices following continuous-time processes, the payoffs of derivative securities can be replicated via continuous trading. However, the discreteness of stock price observations, market frictions, and market closures, render continuous hedging and perfect replication infeasible (e.g., Figlewski, 1989). Consequently, in reality hedgers can adjust their portfolios only infrequently. Besides being more accurate descriptions of reality, as argued in Campbell, Lo, and MacKinlay (1997), discrete time models are easier to implement empirically.

We here derive optimal discrete-time hedges that minimize the hedging error variance, and hence replicate derivative payoffs as closely as possible. Moreover, for the case of a European call and GBM stock price, we obtain a closed-form generalized Black-Scholes formula for the value of the best replicating portfolio and an explicit hedge that is the discrete-time analogue of the standard Black-Scholes perfect hedge. Both expressions involve risk adjustments, directly depending on the stock risk premium, accounting for the inability to hedge perfectly. Finally, we study discrete hedging under return predictability and demonstrate that while optimal discrete hedges and the values of replicating portfolios are close to those in the classical Black-Scholes economy absent

⁹To 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, which can easily be adapted to our setup.

predictability, they deviate considerably from their Black-Scholes counterparts with predictability. Intuitively, our results demonstrate that return predictability improves the ability to predict the moneyness of options, and hence may significantly affect their values.

3.1. Optimal Discrete Hedge and Replication Value

The problem of discrete-time hedging of options has received considerable attention in financial economics, primarily in the context of the Black and Scholes (1973) call option pricing model for which closed-form hedges and prices are available with continuous trading. In seminal works, Boyle and Emanuel (1980) study the properties of hedging error distributions when Black-Scholes hedges are implemented discretely in time, while Leland (1985) investigates such discrete hedging with transaction costs and provides an approximate solution for the call price and hedging policy by adjusting the volatility parameter in the Black-Scholes formula. Wilmott (2006) summarizes further attempts in this direction. Bossaerts and Hillion (1997, 2003) obtain optimal discrete-time hedges from local parametric estimation and demonstrate that Black-Scholes hedges being applied over discrete intervals of time may generate larger hedging errors than theirs. Bertsimas, Kogan, and Lo (2000, 2001) study the properties of discretely adjusted hedges, including Black-Scholes ones, by employing the minimum quadratic criterion. Cochrane and Saa-Requejo (2000) investigate the replication of derivative payoffs with infrequent trading and provide “good deal” bounds on their prices. Discrete hedging has also been employed to address various important issues in finance. In particular, Bakshi, Cao and Chen (1997) employ discrete hedges to empirically evaluate the relative performance of competing option pricing models, while Bakshi and Kapadia (2003) employ discrete hedges to develop a test for the existence and the sign of the volatility risk premium. Jameson and Wilhelm (1992) demonstrate that the risks due to discreteness of portfolio rebalancing explain a significant proportion of option bid-ask spreads.

To investigate the discrete hedging problem, we modify our economic setting of Section 2 as follows. The hedger can now rebalance the tradable wealth only at fixed dates $t = 0, \Delta t, 2\Delta t, \dots, T$, where $\Delta t = T/N$ is the time increment in the model, T the investment horizon and N the number of trading periods. The hedger trades in a riskless bond with return r and a stock with price S_t that follows a stochastic (continuous or discrete) process with one-period stock returns $S_{t+\Delta t}/S_t$ drawn from a known distribution with finite mean $E_t[S_{\tau+\Delta t}/S_\tau]$ and variance $\text{var}_t[S_{\tau+\Delta t}/S_\tau]$, where $\tau \geq t$. In general, we allow the moment parameters to depend on multiple state variables. We consider a hedger who is committed to hold a European-style derivative security with terminal payoff X_T , which depends on the horizon stock price S_T and in general on the past stock price realizations. The hedger’s objective is to minimize the variance criterion (4) subject to the discrete-time budget constraint:

$$W_{t+\Delta t} = e^{r\Delta t}W_t + \theta_t \left(\frac{S_{t+\Delta t}}{S_t} - e^{r\Delta t} \right), \quad (27)$$

where θ_t is the dollar amount invested in the stock. The financial market is in general incomplete since the derivative payoff cannot be perfectly replicated by the available securities in this discrete-time setting even if the stock price follows a simple continuous-time GBM process. The variance

criterion provides a natural metric for evaluating how closely the derivative payoff can be replicated with discrete trading. Consequently, the variance minimization (4) subject to the budget constraint (27) extends the concept of derivative replication to discrete incomplete-market settings.

We proceed by applying the methodology of Section 2 to derive the optimal hedge via dynamic programming in terms of the expected discounted derivative payoff net of replication costs, G , defined analogously to equation (10) (after Proposition 1) as:

$$G_t = E_t \left[X_T e^{-r(T-t)} \right] - E_t \left[W_T^* e^{-r(T-t)} - W_t \right]. \quad (28)$$

Similarly to the discounted net asset payoff (10) in Section 2, the discounted net payoff G in (28) can be interpreted as the value of the self-financing strategy that replicates the derivative payoff as closely as possible, and for which the hedging error is unbiased in the sense that conditionally on time- t information the hedging error is zero on average, so that $E_t[X_T - W_T^*] = 0$. Consequently, we label the discounted net payoff G as the hedger's *derivative replication value*.¹⁰ Proposition 3 reports the optimal hedge, the replication value, and the ensuing hedging error variance.

Proposition 3. *The optimal discrete-time hedge and the derivative replication value are given by:*

$$\theta_t^* = S_t \frac{E_t[e^{-r\Delta t} G_{t+\Delta t}] - G_t}{E_t[e^{-r\Delta t} S_{t+\Delta t}] - S_t}, \quad (29)$$

$$G_t = \frac{1}{\xi_t} E_t \left[\xi_T X_T e^{-r(T-t)} \right], \quad (30)$$

where the process ξ_t is a martingale, given by:

$$\xi_t = \prod_{\tau=0}^{t-\Delta t} \left\{ 1 + \left(\frac{E_\tau[S_{\tau+\Delta t}/S_\tau] - S_{\tau+\Delta t}/S_\tau}{\sqrt{\text{var}_\tau[S_{\tau+\Delta t}/S_\tau]}} \right) \left(\frac{E_\tau[S_{\tau+\Delta t}/S_\tau] - e^{r\Delta t}}{\sqrt{\text{var}_\tau[S_{\tau+\Delta t}/S_\tau]}} \right) \right\} \text{ if } t > 0, \text{ and } \xi_0 = 1. \quad (31)$$

The hedging error variance is given by:

$$\text{var}_t[X_T - W_T^*] = E_t \left[\sum_{\tau=t}^{T-\Delta t} e^{r(T-\tau-\Delta t)} \text{var}_\tau \left(G_{\tau+\Delta t} - \theta_\tau^* \frac{S_{\tau+\Delta t}}{S_\tau} \right) \right]. \quad (32)$$

Proposition 3 provides an analytical characterization of the discrete hedge and replication value in terms of the exogenous model parameters. We observe that the optimal hedge (29) per unit of stock price S_t can be interpreted as the discrete hedging analogue of the standard option delta, measuring the sensitivity of the expected change in the (discounted) replication value to the expected change in the underlying (discounted) stock price. Indeed, with continuous-time complete

¹⁰For comparison, in the quadratic hedging literature (e.g., Bertsimas, Kogan and Lo, 2000; Schweizer, 2001) the replication value is found as the value of W_t that minimizes the mean squared error $E_t(X_T - W_T^*)^2$. However, in contrast to minimum variance hedging, quadratic hedging in general does not produce mean-zero hedging errors (this can be demonstrated even in a simple one-period example). Likewise, Branger and Schlag (2008) demonstrate that biased hedging errors arise when Black-Scholes hedges are applied to discrete hedging problems, which may bias the tests for the existence and the sign of the volatility risk premium.

markets the standard complete-market hedge (9), by employing Itô's Lemma and the PDE for the no-arbitrage value for G , can be rewritten as:

$$\theta_t^{RN} = S_t \frac{E_t[d(e^{-rt}G_t)]}{E_t[d(e^{-rt}S_t)]}. \quad (33)$$

Comparing the expressions for the hedges (29) and (33) we see that both have a similar structure that naturally arises in the context of discrete hedging. This similarity also reinforces our earlier observation that the minimization of the hedging error variance (4) provides a natural extension of the no-arbitrage replication to the case of incomplete markets.

The hedger's replication value (30) is given by the expected deflated terminal payoff of the derivative. So, the structure of the replication value is also analogous to that of the no-arbitrage value in complete markets. The martingale deflator ξ_t is determined analytically in (31) and depends on the unexpected changes in the stock returns per unit of standard deviation weighted by the stock Sharpe ratio. Hence, the stock mean-return μ generally directly affects the derivative replication value through the discrete-time Sharpe ratio. Moreover, comparing the two representations for the replication value in (28) and (30) we see that the deflator ξ_t absorbs the hedging costs (second term in (28)), and so we label ξ as the *hedge-neutral martingale deflator*. Finally, analogously to the complete-market case, the process $e^{-rt}\xi_t S_t$ is also a martingale, as demonstrated in the Appendix. Using this property it can easily be demonstrated, for example, that the replication values for European call and put options satisfy the put-call parity.

We further note that the optimal hedge (29) and the replication value (30) can easily be evaluated numerically by standard Monte Carlo simulation methods, and in some cases explicitly computed in closed form, as demonstrated below. The results of Proposition 3 allow the evaluation of the optimal hedges for stock prices generally being driven by multiple state variables and non-Markovian path-dependent processes, such as GARCH processes for stock return volatilities. These hedges can also be computed for derivative securities with path-dependent payoffs determined by discretely sampled past stock prices (e.g., daily or weekly closing prices), for which closed-form prices are not available even in complete markets. Finally, we note that our discrete-time hedges are more tractable than the quadratic hedges provided in Bertsimas, Kogan and Lo (2001). In particular, the latter are derived in terms of coefficients satisfying recursive equations with expectation operators at each step, which complicates the computations, and do not permit explicit hedges.

The quality of the optimal hedge is measured by the variance of the hedging error (32). We obtain the quality of the hedge as the sum of one-period hedging error variances by applying the law of total variance. Using the expression for $W_{t+\Delta t}$ in (27) we further observe that one-period hedging error variances on the right-hand side of (32) can be rewritten as $\text{var}_\tau[G_{\tau+\Delta t} - W_{\tau+\Delta t}]$. Consequently, by keeping the hedging error variance (32) small, our optimal hedges also keep the whole path of the replication value G_t close to the path of the tradable wealth W_t .

3.2. Application to Option Hedging with GBM Stock Price

We now study the important commonly-studied special case of the derivative being a call option with payoff $X_T = (S_T - K)^+$, $K > 0$, and the stock price following a continuous-time GBM process:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dw_t. \quad (34)$$

When the hedges can be rebalanced continuously the financial market is dynamically complete here, and the closed-form Black-Scholes option price obtains via no-arbitrage valuation. With our incomplete market when the hedging takes place over discrete time intervals, we apply the results of Proposition 3 to derive the optimal hedge and option replication value, as reported in Proposition 4.

Proposition 4. *When the stock price follows GBM the optimal hedge and the replication value of the call option are given by:*

$$\theta_t^* = S_t \frac{E_t[e^{-r\Delta t} G_{t+\Delta t}] - G_t}{E_t[e^{-r\Delta t} S_{t+\Delta t}] - S_t} = \frac{e^{(\mu-r)\Delta t} \sum_{k=0}^{n-1} \omega_{n-1,k} C_k^{BS}(S_t, t) - \sum_{k=0}^n \omega_{n,k} C_k^{BS}(S_t, t)}{e^{(\mu-r)\Delta t} - 1}, \quad (35)$$

$$G_t = \sum_{k=0}^n \omega_{n,k} C_k^{BS}(S_t, t), \quad (36)$$

where $n = (T - t)/\Delta t$, and $C_k^{BS}(S_t, t)$ is the Black-Scholes price with the risk-adjusted discount rate $r_k = r + (\mu - r + k\sigma^2/n)$,

$$C_k^{BS}(S_t, t) = S_t \Phi\left(\frac{\ln(S_t/K) + (r_k + 0.5\sigma^2)(T - t)}{\sigma\sqrt{T - t}}\right) - K e^{-r_k(T-t)} \Phi\left(\frac{\ln(S_t/K) + (r_k - 0.5\sigma^2)(T - t)}{\sigma\sqrt{T - t}}\right),$$

and the weight $\omega_{n,k}$ is given by

$$\omega_{n,k} = \binom{n}{k} \frac{(e^{(\mu-r+\sigma^2)\Delta t} - 1)^{n-k} (e^{\sigma^2\Delta t} - e^{(\mu-r+\sigma^2)\Delta t})^k}{(e^{\sigma^2\Delta t} - 1)^n}. \quad (37)$$

Moreover, as Δt shrinks to zero, the replication value (36) converges to the Black-Scholes price.

Proposition 4 provides a closed-form expression for the optimal hedge (35) and the option replication value (36) in terms of weighted Black-Scholes prices C_k^{BS} with the risk-adjusted discount rate r_k (replacing the interest rate r in the standard Black-Scholes price). Despite numerous applications of discrete-time hedging in finance, to our best knowledge, above is the first to provide closed-form expressions for the replication value and hedging policy. Our closed-form hedges can be applied as an alternative to widely used sub-optimal discretely evaluated Black-Scholes hedges (e.g., Boyle and Emanuel, 1980; Jameson and Wilhelm, 1992; Bakshi, Cao, and Chen, 1997; Bertsimas, Kogan and Lo, 2000; Bakshi and Kapadia, 2003), approximate hedges for small hedging intervals (e.g., Leland 1985; Wilmott, 2006, and literature therein), as well as hedges obtained from local parametric estimation (e.g., Bossaerts and Hillion, 1997, 2003).

As demonstrated in the Appendix, each Black-Scholes price C_k^{BS} in (36) corresponds to a k -period optimal static hedge when the hedger keeps θ_k^* dollars in stock between times t and $t + k\Delta t$ without rebalancing, and invests only in bonds thereafter.¹¹ Consequently, the risk adjustments in r_k arise due to the inability to hedge perfectly, and increase in the length of the static period k , and generally depend on the risk premium $\mu - r$. Therefore, in contrast to Black-Scholes hedges, our hedges depend on the stock mean return μ . As pointed out in Bossaerts and Hillion (1997), the fact that Black-Scholes hedges do not depend on the stock mean return significantly impairs their performance in discrete hedging applications. Furthermore, the replication value formula (36) demonstrates that the option payoff can be hedged by a portfolio of static hedging policies, each accounting for the uncertainty in k periods. The increase in the number of trading periods n improves the diversification of risks in this portfolio, and hence its value converges to Black-Scholes price in the continuous-time limit.

Remark 4 (Alternative representation for the replication value). When the number of hedging intervals n increases, the weights $\omega_{n,k}$ given by (37) may become very large due to the powers of n and the number of choices of k out of n possibilities. Therefore, small computational errors in replication value formula (36) are multiplied by large coefficients, decreasing the precision of calculations.¹² We here present an alternative representation for the replication value G_t which is computationally more efficient, though less intuitive economically than representation (36). As demonstrated in the Appendix, the call option replication value is alternatively given by:

$$G_t = \frac{1}{2\pi\sigma\sqrt{(T-t)}} \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} R(x, \varepsilon)^n \cos(n\varphi(x, \varepsilon)) e^{-\frac{\varepsilon^2}{2}} d\varepsilon \right] (S_t e^x - K)^+ e^{-\frac{(x - (\mu - 0.5\sigma^2)(T-t))^2}{2\sigma^2(T-t)}} dx, \quad (38)$$

where the closed-form expressions for $R(x, \varepsilon)$ and $\varphi(x, \varepsilon)$ are given in the Appendix, and the integrals in (38) can easily be evaluated numerically. This alternative formula avoids large coefficients and is computationally stable.

3.3. Option Hedging under Stock Return Predictability

We next apply our methodology to study optimal hedges and replication values of European call options under stock return predictability. While return predictability is a well-documented and important feature of financial markets (e.g., Lo and MacKinlay, 1988, 1990; Fama and French,

¹¹In particular, the replication value under this static hedge is given by $aC_0^{BS} + bC_k^{BS}$, where a and b are constants, giving rise to Black-Scholes C_k^{BS} prices in (36). The weight $\omega_{n,k}$ is also driven by k and depends on the number of possible choices of k out of n periods remaining in which the static hedge is applied.

¹²The loss of precision can be demonstrated in the following example:

$$(q + 1 - q)^n = \sum_{k=0}^n \binom{n}{k} q^{n-k} (1 - q)^k, \quad \text{where } q = \sqrt{2}.$$

Even though the equation above holds as identity, the numerical calculations of its left-hand and right-hand sides are not equivalent. In computer-based calculations (e.g., MATLAB) $q = \sqrt{2}$ is replaced by an approximation with finite number of digits. The approximation error is then multiplied by a large coefficient, leading to calculation errors when n is large.

1988; Lo and Wang, 1995), it is well known that in complete markets option hedges and prices do not depend on stock mean-returns through which predictability most commonly affects stock prices. However, in a more realistic incomplete-market setting, the ability to predict the moneyness of an option in general affects its value since investors may consider the Black-Scholes price to be too high or too low depending on whether an option is very likely to be out of the money or in the money, respectively. Consequently, the stock mean-return becomes an important parameter affecting the optimal hedge and the replication value in an incomplete-market setting.

The implications of return predictability for derivative pricing have been studied in some related works. In particular, Lo and Wang (1995) consider a continuous-time complete-market setting with mean-reverting detrended stock log-prices. They demonstrate that even though Black-Scholes option prices are functionally independent of mean-returns, the stock return volatility estimated from discretely sampled data depends on the speed of mean-reversion, and hence the Black-Scholes formula should be adjusted to account for predictability in the data. Bertsimas, Kogan and Lo (2000) in a similar setting derive the asymptotic properties of hedging errors under discretely applied Black-Scholes hedges, while Schwartz (1997) studies the prices of commodity futures in a setting with mean-reverting commodity log-prices. To our best knowledge our analysis below is the first to evaluate the impact of predictability on optimal discrete-time hedges and replication values.

We consider a setting where the derivative is again a European call option with payoff $(S_T - K)^+$, $K > 0$, and detrended log-price S is mean-reverting (as in Lo and Wang (1995) and Bertsimas, Kogan, and Lo (2000)):

$$d(\ln(S_t) - \bar{\mu}t) = -\kappa(\ln(S_t) - \bar{\mu}t)dt + \sigma dw_t, \quad (39)$$

where $\bar{\mu} > 0$ is the trend parameter and $\kappa > 0$ is the speed of mean-reversion. We note that the GBM process (34) is a special case of (39) for $\kappa = 0$ and $\bar{\mu} = \mu - 0.5\sigma^2$. When hedges are rebalanced continuously the market is complete, and hence the replication value is given by the standard Black-Scholes formula that does not depend on the stock mean-return. To solve the hedging problem with infrequent trading we observe that the replication value in (30) can recursively be rewritten as:

$$G_t = \frac{1}{\xi_t} E_t[\xi_{t+\Delta t} G_{t+\Delta t} e^{-r\Delta t}], \quad (40)$$

where ξ_t is as given by (31). Since the stock returns $S_{\tau+\Delta t}/S_\tau$ are log-normal, their conditional means and variances are available in closed form (e.g., Lo and Wang, 1995). Substituting these means and variances into the expression for the process ξ_t yields an analytic expression for $\xi_{t+\Delta t}/\xi_t$. The replication value is then computed via backwards induction starting with $G_T = (S_T - K)^+$.

Figure 1 presents the results of numerical computations for a call option with strike price $K = 0.5$ and maturity $T = 1$ for calibrated model parameters.¹³ Panels (a) and (b) of Figure 1

¹³In our numerical computations, we set $\bar{\mu} = 0.064$, $\sigma = 0.166$, $r = 0.01$ so that the (long-run) expected Sharpe ratio for process (39), given by $(\bar{\mu} + 0.5\sigma^2 - r)/\sigma$, matches the Sharpe ratio of 0.41 and stock return volatility of 0.166 for US stock market over the period of 1934-2003 years (e.g., Cogley and Sargent, 2008). Following the literature, we




Figure 1: Option Deltas, Replication Values, and Hedging Quality.

Panels (a) and (b) plot the deltas and replication values as functions of the stock price S for varying levels of speed of mean-reversion, κ , for the following calibrated parameter values: $\bar{\mu} = 0.066$, $\sigma = 0.166$, $r = 0.01$, $K = 0.5$, $n = 24$, and $T = 1$ (see footnote 13). Panel (c) plots the hedging error standard deviations for different values of the number of trading periods n for $\kappa = 0.4$, where the remaining parameters are as in panels (a) and (b).

depict the discrete-time deltas θ^*/S , and the replication values G under bimonthly rebalancing, respectively, as functions of the stock price S for different plausible speeds of mean-reversion κ . For comparison, the Black-Scholes option deltas and prices (solid lines) are also presented. The plots on panels (a) and (b) demonstrate that in the absence of predictability ($\kappa = 0$) the deltas and replication values are closely approximated by their Black-Scholes counterparts, consistently with the literature (e.g., Leland, 1995; Bertsimas, Kogan and Lo, 2000).

The difference between discrete-time and Black-Scholes deltas and replication values, however, becomes significantly larger as the speed of mean-reversion increases. In particular, panel (b)

let the mean-reversion parameter to be $\kappa = 0.4$ or $\kappa = 0.7$. In particular $\kappa = 0.4$ is calibrated in Lo and Wang (1995) from historical daily returns on CRSP value-weighted market index from 1962 to 1990, while $\kappa = 0.7$ is an upper bound for plausible mean-reversion parameters for commodity price processes, as estimated in Schwartz (1997).

reveals that higher predictability decreases the replication values around the strike price, leading to a larger deviation from the Black-Scholes benchmark. To see this, we first note that detrended stock log-prices are negative around the strike price. Therefore, the stock log-prices will revert back to the trend line, and towards the option expiration date the stock price will fluctuate around $e^{\mu T} > 1$. Consequently, higher predictability increases the likelihood that the option with strike price $K = 0.5$ will be in the money. In turn, the increase in the probability of moneyness moves the replication value closer to the lower no-arbitrage bound for option prices, $S_t - Ke^{-r(T-t)}$. This is because with higher probability of moneyness the option payoff $(S_T - K)^+$ is better approximated by the exercise value $S_T - K$, and hence the option replication value should be closer to the replication value of $S_T - K$, which is given by $S_t - Ke^{-r(T-t)}$. Hence, the replicating value is decreased with higher predictability and stock price close to the strike price, which also translates into larger deviations of discrete deltas from Black-Scholes ones.¹⁴ Therefore, the inability to trade continuously considerably affects the optimal hedging strategy in the presence of return predictability, and hence impairs the ability of the Black-Scholes model to predict market prices of options. Accordingly, one economic implication that emerges from our analysis is that the deviations of market prices of options from the Black-Scholes prices should be larger with stronger return predictability.

Panel (c) plots the quality of the discrete hedging, the standard deviation of the hedging error, as a function of the stock price S for varying levels of rebalancing frequency, bimonthly, weekly and daily. We see that the standard deviations decrease with higher number of rebalancing periods and are hump-shaped functions of the stock price S . The quality of the hedging is better for stock prices which are much lower or higher than the strike price since in these cases there is less uncertainty about whether the option will be out of the money or in the money. Furthermore, since below the strike the detrended log-prices are negative and hence revert back to zero, the probability of stock prices going up exceeds the probability of going down. Consequently, the standard deviations are maximized at a stock price below the strike, where the uncertainty about exceeding the strike is higher.

4. Generalizations

We now generalize our main results on minimum-variance hedging in Section 2 to richer settings along three dimensions. First, we consider a more general model in which the hedger is guided by a linear mean-variance criterion over the hedging error. Then, we demonstrate that the minimum-variance hedging model can easily be extended to a richer environment with multiple non-tradable assets and stocks and study dynamic hedging with time-varying correlations. Finally, we extend the analysis to the case of assets following processes with jumps.

¹⁴In the limiting case of perfect predictability with the speed of mean-reversion κ going to infinity, the discrete-time replication value converges to $S_t - Ke^{-r(T-t)}$. Intuitively, since the option is predicted to be in the money with certainty, its payoff can be replicated by a portfolio that buys one stock and borrows $Ke^{-r(T-t)}$ at time t . In contrast to the discrete-time model, the Black-Scholes price while being functionally independent of κ fails to account for the effect of predictability.

4.1. Mean-Variance Hedging and Benchmark Tracking

We here consider a hedger who also accounts for the mean hedging error, and trades it off against the hedging error variance. Such a mean-variance hedging criterion is commonly employed in a variety of, primarily static, settings (e.g., Anderson and Danthine, 1980, 1981, 1983; Hirshleifer, 1988; Duffie, 1989; Duffie and Jackson, 1989). Our analysis in this Section is also related to the literature on portfolio management with benchmarking. In this literature, money managers are evaluated relative to a benchmark portfolio and are concerned about their tracking error, defined as the deviation of a manager's performance from that of the benchmark. The mean-variance tracking error model amounts to mean-variance hedging if we relabel the non-tradable asset X as the benchmark portfolio and observe that tracking error is the negative of hedging error. Roll (1992), Jorion (2003), Gomez and Zapatero (2003), Cornell and Roll (2005) discuss the implications of such benchmarking on portfolio efficiency and asset pricing. Chan, Karceski and Lakonishok (1999) and Costa and Paiva (2002) discuss the implications of estimation risk and robust portfolio selection with benchmarking. These works all employ a static mean-variance framework by either minimizing the tracking error variance for a given mean, or maximizing the tracking error mean for a given variance.

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

$$\max_{\theta_t} E_t[W_T - X_T] - \frac{\gamma}{2} \text{var}_t[W_T - X_T], \quad (41)$$

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 (41) evaluated at the optimal policy. As in Section 2 we consider the time-consistent solution to problem (41) obtained by dynamic programming. Proposition 5 reports the dynamically optimal hedging policy along with the value function.

Proposition 5. *The dynamically optimal mean-variance hedge, θ_t^* , and the corresponding value function, J_t , are given by*

$$\begin{aligned} \theta_t^* &= \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial X_t} + S_t \frac{\partial E_t^*[X_T e^{-r(T-t)}]}{\partial S_t} + \frac{\mu_t - r}{\gamma\sigma_t^2} e^{-r(T-t)} \\ &\quad - \frac{\rho\nu_t}{\sigma_t} X_t \frac{\partial E_t^*\left[\left(\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds\right) e^{-r(T-t)}\right]}{\partial X_t} - S_t \frac{\partial E_t^*\left[\left(\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds\right) e^{-r(T-t)}\right]}{\partial S_t}, \end{aligned} \quad (42)$$

$$\begin{aligned} J_t &= -\frac{\gamma}{2}(1 - \rho^2) E_t \left[\int_t^T \nu_s^2 X_s^2 \left(\frac{\partial E_s^*[X_T - \int_s^T \frac{1}{\gamma} \left(\frac{\mu_\tau - r}{\sigma_\tau}\right)^2 d\tau]}{\partial X_s} \right)^2 ds \right] \\ &\quad + W_t e^{r(T-t)} - E_t^*[X_T] + \frac{1}{2} E_t^* \left[\int_t^T \frac{1}{\gamma} \left(\frac{\mu_s - r}{\sigma_s}\right)^2 ds \right]. \end{aligned} \quad (43)$$

Proposition 5 reveals that the dynamically optimal mean-variance hedge is comprised of three types of terms. The first two terms in (42) 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 (42) 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 value function (43) 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.

The optimal hedge (42) 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 . We next explicitly provide the optimal hedge in a stochastic environment as in one of our examples in Section 2.3, where the non-tradable asset follows the mean-reverting process (15), and the stock price follows the stochastic volatility process (16), which has been employed in portfolio choice context by Chacko and Viceira (2003), and is a special case of the quadratic asset returns model in Liu (2007).¹⁵ Corollary 4 reports the optimal hedge.

Corollary 4. *The optimal hedging policy for the mean-reverting stochastic-volatility (15)–(16) is given by*

$$\begin{aligned} \theta_t^* = & \rho\bar{v}X_t e^{-(r+\kappa+\rho\bar{v}(\mu-r))(T-t)} \left(1 + \frac{(\mu-r)^2}{\gamma(\kappa+\rho\bar{v}(\mu-r))} \right) \\ & + \frac{\mu-r}{\gamma} X_t e^{-r(T-t)} - \frac{\rho\bar{v}X_t(\mu-r)^2}{\gamma(\kappa+\rho\bar{v}(\mu-r))} e^{-r(T-t)}. \end{aligned} \quad (44)$$

The optimal hedge in (44) is comprised of three terms. The first term is given by the pure minimum-variance hedge scaled by a factor, which arises due to the fact that the agent now hedges not only the non-tradable asset but also the fluctuations in the squared market price of risk $((\mu_t - r)/\sigma_t)^2 = (\mu - r)^2 X_t$. Specifically, since the squared market price of risk is proportional to the asset price X , hedging its fluctuations also involves a minimum-variance hedge multiplied by a constant accounting for the compounding effect of integration in (42). This interacts with the pure minimum-variance hedge for the non-tradable asset terminal payoff, giving rise to the scaling factor in the first term of (44). The second term in (44) is given by the pure speculative demand, which is adjusted by the third term in (44), accounting for the residual hedging demand due to fluctuations in the market price of risk.

¹⁵Similarly, it is possible to obtain explicit hedges when the non-tradable asset follows the mean-reverting process (11) and the stock follows a process with mean-reverting Gaussian mean-return, as in Kim and Omberg (1996).

4.2. Dynamic Hedging with Multiple Assets and Time-Varying Correlations

We now demonstrate that the results of Section 2 can be generalized to the case with multiple non-tradable assets and stocks, and provide an application investigating the effects of correlation risk. We consider an economy in which uncertainty is generated by two multi-dimensional Brownian motions $w_X = (w_{X1}, \dots, w_{XN})^\top$ and $w = (w_1, \dots, w_K)^\top$. By ρ we denote the $N \times K$ correlation matrix with elements $\rho = \{\rho_{nk}\}$ representing correlations between the Brownian motions w_{Xn} and w_k .

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

$$\frac{dX_{it}}{X_{it}} = m_i(X_t, t)dt + \nu_i(X_t, t)^\top dw_{Xt}, \quad i = 1, \dots, N, \quad (45)$$

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

$$\frac{dS_{jt}}{S_{jt}} = \mu_j(X_t, S_t, t)dt + \sigma_j(X_t, S_t, t)^\top dw_t, \quad j = 1, \dots, K, \quad (46)$$

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

$$dW_t = [rW_t + \theta_t^\top (\mu_t - r)]dt + \theta_t^\top \sigma_t dw_t. \quad (47)$$

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 (47). The optimal policy is then derived by dynamic programming as in Section 2. Proposition 6 reports the dynamically optimal hedge and its associated quality.

Proposition 6. *The optimal hedging policy and the corresponding variance of hedging error are given by*

$$\theta_t^* = (\nu_t \rho^\top \sigma_t^{-1})^\top I_{X_t} \frac{\partial E_t^*[\phi^\top X_T e^{-r(T-t)}]}{\partial X_t^\top} + I_{S_t} \frac{\partial E_t^*[\phi^\top X_T e^{-r(T-t)}]}{\partial S_t^\top}, \quad (48)$$

$$\text{var}_t[\phi^\top X_T - W_T^*] = E_t \left[\int_t^T \frac{\partial E_s^*[\phi^\top X_T]}{\partial X_s} \nu_s I_{X_s} (I - \rho^\top \rho) I_{X_s} \nu_s^\top \frac{\partial E_s^*[\phi^\top X_T]}{\partial X_s^\top} ds \right], \quad (49)$$

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

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

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 (48) 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 static – we provide one such application below. The expression (49) 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.

We next apply the general results of Proposition 6 to study dynamic hedging in the presence of *correlation risk*, namely the risk associated with the changes in asset values when the correlations between the underlying assets are time-varying and hard to predict. We here provide closed-form expressions for optimal hedges when correlations are time-varying and identify the economic channels through which correlation risk affects dynamic hedging. The time-variation of correlations is well documented for stock indices and other assets (e.g., Bekaert and Harvey, 1995; Moskowitz, 2003; Driessen, Maenhout, and Vilkov, 2009) and has been attributed to market clearing effects, contagion, as well as financial crises (e.g., Cochrane, Longstaff, Santa-Clara, 2008; Pavlova and Rigobon, 2008; Buraschi, Porchia and Trojani, 2010). This, in turn, has led to a growing recent literature studying the economic effects of correlation risk. In particular, Liu (2007) studies portfolio choice with quadratic asset returns and provides closed-form solutions for settings when asset returns correlations and volatilities are time-varying. Driessen, Maenhout, and Vilkov (2009) demonstrate that stock correlation risk is priced and identify the correlation risk premium. Buraschi, Porchia and Trojani (2010) demonstrate that correlation risk significantly affects portfolio choice of CRRA investors and induces large hedging demands.

In general, correlation risk and volatility risk, arising due to time-varying asset return volatilities, are both driven by the time-variation in the components of the volatility matrix σ and hence may be difficult to disentangle. Motivated by the widely employed stock return models with stochastic volatilities but constant correlations, we consider a model with time-varying correlations which includes as special cases dynamics with constant correlations and time-varying square-root (as in

Heston, 1993; Liu, 2007) and inverse square-root (as in Chacko and Viceira, 2003) volatilities – this will subsequently enable us to disentangle the hedging effects of correlation and volatility risks. Specifically, the non-tradable asset and stock prices are assumed to evolve as:

$$dX_t = \kappa(\bar{X} - X_t) + \bar{\nu}\sqrt{X_t}dw_{xt}, \quad (50)$$

$$\frac{dS_{jt}}{S_{jt}} = \mu_j dt + \bar{\sigma}_{j1} \frac{\sqrt{X_t}}{a_1 + b_1 X_t} dw_{1t} + \bar{\sigma}_{j2} \frac{\sqrt{X_t}}{a_2 + b_2 X_t} dw_{2t}, \quad j = 1, 2, \quad (51)$$

where $\kappa > 0$, $\bar{\sigma}_{jk} \neq 0$, the correlation between Brownian motions w_x and w_k is ρ_k , $a_k \geq 0$, $b_k \geq 0$, and $a_k + b_k = 1$, $k = 1, 2$. The time-varying correlations between asset returns are driven by the *volatility shape parameters* a_k and b_k (as elaborated on below), which determine the *volatility weights* $\sigma_{jk}(X_t) \equiv \bar{\sigma}_{jk}\sqrt{X}/(a_k + b_k X)$. The shape parameters allow for considerable flexibility in modeling the volatility weights, which can be increasing, decreasing, or hump-shaped functions of asset price X , depending on the parameters a_k and b_k .

The correlations between the instantaneous changes in asset and stock prices are, in general, time-varying and given by:

$$\text{corr}\left(dX_t, \frac{dS_{jt}}{S_{jt}}\right) = \frac{\rho_1 \bar{\sigma}_{j1}/(a_1 + b_1 X_t) + \rho_2 \bar{\sigma}_{j2}/(a_2 + b_2 X_t)}{\sqrt{\bar{\sigma}_{j1}^2/(a_1 + b_1 X_t)^2 + \bar{\sigma}_{j2}^2/(a_2 + b_2 X_t)^2}}, \quad (52)$$

$$\begin{aligned} \text{corr}\left(\frac{dS_{1t}}{S_{1t}}, \frac{dS_{2t}}{S_{2t}}\right) &= \frac{\bar{\sigma}_{11}\bar{\sigma}_{21}/(a_1 + b_1 X_t) + \bar{\sigma}_{12}\bar{\sigma}_{22}/(a_2 + b_2 X_t)}{\sqrt{\bar{\sigma}_{11}^2/(a_1 + b_1 X_t)^2 + \bar{\sigma}_{12}^2/(a_2 + b_2 X_t)^2}} \\ &\times \frac{1}{\sqrt{\bar{\sigma}_{21}^2/(a_1 + b_1 X_t)^2 + \bar{\sigma}_{22}^2/(a_2 + b_2 X_t)^2}}. \end{aligned} \quad (53)$$

The asset return dynamics with time-varying volatilities but constant correlations are nested as special cases of the dynamics (50)–(51) when the differences in shape parameters, $a_1 - a_2$ and $b_1 - b_2$, are equal to zero. Indeed, for $a_k = 0$ and $b_k = 1$, $k = 1, 2$, we obtain the square-root stochastic volatility model employed in Heston (1993) and Liu (2007), while for $a_k = 1$ and $b_k = 0$, $k = 1, 2$, we obtain the inverse square-root stochastic volatility model employed in Chacko and Viceira (2003). We note that the source of correlation risk is the time-variation in the ratios of volatility weights $\sigma_{j1}(X_t)$ and $\sigma_{j2}(X_t)$, $j = 1, 2$, given by $\bar{\sigma}_{j1}(a_2 + b_2 X_t)/(\bar{\sigma}_{j2}(a_1 + b_1 X_t))$. In particular, when the ratios $\sigma_{j1}(X_t)/\sigma_{j2}(X_t)$ are stochastic, the fraction of stock j instantaneous return variance due to Brownian motion w_k , given by $\sigma_{jk}(X_t)^2/(\sigma_{j1}(X_t)^2 + \sigma_{j2}(X_t)^2)$, is stochastic. Consequently, the correlations (52)–(53) become time-varying due to the changing contributions of Brownian motions w_k to the overall stock return uncertainty. We also observe that the ratios of volatility weights being stochastic is equivalent to the differences in shape parameters, $a_1 - a_2$ and $b_1 - b_2$, being non-zero, and hence these differences quantify the departure of asset price dynamics (50)–(51) from the case of constant correlations.¹⁶ Corollary 5 summarizes our results.

¹⁶The correlations (52) also become constant in the special case of $\bar{\sigma}_{12}$ and $\bar{\sigma}_{21}$ being zero. However, this case implies zero correlation (53) between the stock returns and is assumed away by the condition that $\bar{\sigma}_{jk} \neq 0$ since stock returns are likely to be correlated due to exposure to common market risk factors.

Corollary 5. *The optimal hedge for the time-varying correlation model (50)–(51) is given by*

$$\theta_t^* = \left\{ \begin{pmatrix} \rho_1 \bar{\sigma}_{22} - \rho_2 \bar{\sigma}_{21} \\ \rho_2 \bar{\sigma}_{11} - \rho_1 \bar{\sigma}_{12} \end{pmatrix} (a_1 + b_1 X_t) + \begin{pmatrix} \rho_2 \bar{\sigma}_{21} \\ -\rho_1 \bar{\sigma}_{11} \end{pmatrix} \left((a_1 - a_2) + (b_1 - b_2) X_t \right) \right\} \frac{\bar{\nu} e^{-\tilde{\kappa}(T-t)}}{\det \bar{\sigma}}, \quad (54)$$

where $\tilde{\kappa} = \kappa + \bar{\nu}(\rho_1 b_1 \tilde{\mu}_1 + \rho_2 b_2 \tilde{\mu}_2)$, $(\tilde{\mu}_1, \tilde{\mu}_2)^\top = \bar{\sigma}^{-1}(\mu_1 - r, \mu_2 - r)^\top$, and $\bar{\sigma}$ is a 2×2 matrix with elements $\{\bar{\sigma}_{jk}\}_{j,k=1,2}$.

Corollary 5 provides the optimal hedge in closed-form and demonstrates that it can be decomposed into two components, a time-varying volatility hedge (first term in (54)) and a time-varying correlation hedge (second term in (54)). The volatility hedge is driven by volatility risk alone and is not affected by correlation risk, which is driven by the differences in volatility shape parameters $a_1 - a_2$ and $b_1 - b_2$, and coincides with the optimal hedge with correlation risk absent and only volatility risk present. The correlation hedge only arises in settings with time-varying correlations, and its magnitude is driven by the differences in shape parameters $a_1 - a_2$ and $b_1 - b_2$, which quantify the departure from the constant correlation model, as discussed above.

The expression for the optimal hedge (54) also reveals that correlation risk complicates the dynamic hedging by inducing additional hedging demands. Specifically, the hedging implications are now confounded by the fact that non-zero differences in shape parameters break the perfect co-movement (perfect positive or negative correlation) between the volatility weights $\sigma_{j1}(X_t)$ and $\sigma_{j2}(X_t)$. Hence, assuming for simplicity that $\rho_1 = \rho_2$, while the stock-return source of uncertainty w_k with the volatility weight better co-moving with the non-tradable asset volatility can be used to cancel some of the non-tradable asset volatility, the other source of uncertainty tends to amplify it, complicating the dynamic hedging. Finally, we observe that in general, both the volatility and correlation hedges are time-varying and linearly depend on the asset price X . However, in the special case of square-root volatility weights (when $b_1 = b_2 = 0$) both hedges are constant. Intuitively, in this case the hedging is facilitated by the fact that the instantaneous volatilities of the non-tradable and tradable assets are proportional to each other, and hence the non-tradable asset fluctuations can optimally be reduced by a constant hedge.

4.3. Dynamic Hedging with Jumps

In this Section we extend the analysis of Section 2 and provide tractable closed-form expressions for dynamically optimal minimum-variance hedges when non-tradable asset and stock prices follow processes with jumps. Such processes are commonly employed in the portfolio choice and dynamic hedging literature, though closed-form expressions are rarely obtained. In particular, Liu, Longstaff and Pan (2003), Liu and Pan (2003), Ait-Sahalia, Cacho-Diaz and Hurd (2009) derive semi-explicit solutions for portfolio choice problems with jumps, while Schweizer (1994), Pham, Rheinlander and Schweizer (1998), Bertsimas, Kogan and Lo (2001) study dynamic hedging with quadratic criterion in the presence of jumps but do not provide closed-form solutions.

We here consider an economic setting with uncertainty generated by two Brownian motions w_X and w with correlation parameter ρ as before, and now also by a *Poisson jump process* q with

intensity λ (average rate of jump arrival), representing the number of jumps up to a point in time. The non-tradable asset with price X is assumed to follow the dynamics

$$dX_t = m(X_t, t)X_t dt + \nu(X_t, t)X_t dw_{Xt} + \alpha(X_{t-}, t)X_{t-} dq_t, \quad (55)$$

where m , ν , the jump size α and intensity λ are deterministic functions of X , and X_{t-} denotes the left-side limit of X_t . The dynamics of the stock is similarly modeled as

$$dS_t = \mu(X_t, S_t, t)S_t dt + \sigma(X_t, S_t, t)S_t dw_t + \beta(X_{t-}, S_{t-}, t)S_{t-} dq_t, \quad (56)$$

where μ , σ and the jump size β are deterministic functions of X and S , and S_{t-} denotes the left-side limit of process S_t .¹⁷ We will denote m_t , ν_t , μ_t , σ_t , α_t , β_t and λ_t as shorthand for the coefficients in equations (55)–(56) and the jump intensity evaluated at S_t and X_t , and α_{t-} and β_{t-} as shorthand for the left-side limits of jumps sizes.

Given the hedging policy θ , the hedger's tradable wealth follows the process

$$dW_t = [rW_t + \theta_t(\mu_t - r)]dt + \theta_t \sigma_t dw_t + \theta_{t-} \beta_{t-} dq_t. \quad (57)$$

At each time t the hedger chooses the optimal hedging policy θ_t that minimizes the variance of hedging error, $X_T - W_T$, subject to the budget constraint (57). In contrast to Section 2, the financial market here has a higher degree of incompleteness due to the additional source of risk driven by the jump process q , which cannot be hedged simultaneously with Brownian motion components when only one tradable security is available. The optimal hedging policy is then obtained by adopting the dynamic programming approach developed in Section 2. Proposition 7 reports the dynamically optimal hedge along with the hedging error variance.

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

$$\begin{aligned} \theta_t^* &= \left(\frac{\rho \nu_t}{\sigma_t} X_t \frac{\partial G(X_t, S_t, t)}{\partial X_t} + S_t \frac{\partial G(X_t, S_t, t)}{\partial S_t} \right. \\ &\quad \left. + \frac{\lambda_t \beta_t}{\sigma_t^2} \left[G(X_t(1 + \alpha_t), S_t(1 + \beta_t), t) - G(X_t, S_t, t) \right] \right) \frac{\sigma_t^2}{\sigma_t^2 + \lambda_t \beta_t^2}, \end{aligned} \quad (58)$$

$$\begin{aligned} \text{var}_t[X_T - W_T^*] &= (1 - \rho^2) E_t \left[\int_t^T \nu_s^2 X_s^2 \left(\frac{\partial G(X_s, S_s, s)}{\partial X_s} \right)^2 e^{2r(T-s)} ds \right] \\ &\quad + E_t \left[\int_t^T \lambda_s \left(G(X_s(1 + \alpha_s), S_s(1 + \beta_s), s) - G(X_s, S_s, s) \right)^2 e^{2r(T-s)} ds \right] \\ &\quad + E_t \left[\int_t^T \left(\frac{\rho \nu_s}{\sigma_s} X_s \frac{\partial G(X_s, S_s, s)}{\partial X_s} + S_s \frac{\partial G(X_s, S_s, s)}{\partial S_s} \right)^2 \sigma_s^2 e^{2r(T-s)} \right. \\ &\quad \left. - \left(\frac{\rho \nu_s}{\sigma_s} X_s \frac{\partial G(X_s, S_s, s)}{\partial X_s} + S_s \frac{\partial G(X_s, S_s, s)}{\partial S_s} \right) \right. \\ &\quad \left. + \frac{\lambda_s \beta_s}{\sigma_s^2} \left(G(X_s(1 + \alpha_s), S_s(1 + \beta_s), s) - G(X_s, S_s, s) \right)^2 \frac{\sigma_s^4 e^{2r(T-s)}}{\sigma_s^2 + \lambda_s \beta_s^2} ds \right], \end{aligned} \quad (59)$$

¹⁷If there is a jump at time t , then X_t and S_t denote the values of X and S immediately after the jump while X_{t-} and S_{t-} denote the values immediately before the jump, respectively. Following the literature, parameters m , ν , μ and σ are evaluated at X_t and S_t rather than at X_{t-} and S_{t-} , which does not affect the results since the processes under these two alternatives differ only on a countable number of points (e.g., Duffie, 2001).

where $G(X_t, S_t, t) \equiv E^*[X_T e^{-r(T-t)} | X_t, S_t]$ is the conditional expectation under the unique hedge-neutral measure P^* on which are defined two Brownian motions w^* and w_x^* with correlation ρ and the jump process q^* with intensity λ^* such that the processes for the non-tradable asset, X , and stock price, S , are given by

$$dX_t = \left(m_t - \rho \nu_t \sigma_t \frac{\mu_t - r + \lambda_t \beta_t}{\sigma_t^2 + \lambda_t \beta_t^2} \right) X_t dt + \nu_t X_t dw_{x_t}^* + \alpha_t - X_t - dq_t^*, \quad (60)$$

$$dS_t = \left(\mu_t - \sigma_t^2 \frac{\mu_t - r + \lambda_t \beta_t}{\sigma_t^2 + \lambda_t \beta_t^2} \right) S_t dt + \sigma_t S_t dw_t^* + \beta_t - S_t - dq_t^*, \quad (61)$$

and the intensity parameter λ_t^* is given by

$$\lambda_t^* = \lambda_t \left(1 - \beta_t \frac{\mu_t - r + \lambda_t \beta_t}{\sigma_t^2 + \lambda_t \beta_t^2} \right), \quad (62)$$

assuming that the model parameters are such that $\lambda_t^* > 0$. Moreover, the probability measure P^* is defined by the Radon-Nykodym derivative

$$\frac{dP^*}{dP} = e^{\int_0^T \left[\lambda_s^* - \lambda_s - \frac{1}{2} \left(\sigma_s \frac{\mu_s - r + \lambda_s \beta_s}{\sigma_s^2 + \lambda_s \beta_s^2} \right)^2 \right] ds - \int_0^T \sigma_s \frac{\mu_s - r + \lambda_s \beta_s}{\sigma_s^2 + \lambda_s \beta_s^2} dw_s + \int_0^T \ln \frac{\lambda_s^*}{\lambda_s} dq_s}.$$

The dynamically optimal hedge (58) generalizes the previous optimal hedge (5) by incorporating the effect of jumps. Specifically, the new, third term in (58) captures the sensitivity of the hedge-neutral asset value G to potential jumps in the non-tradable asset or stock prices. The hedge quality (59) is now comprised of three terms. The new second and third terms quantify the additional source of incompleteness due to Poisson jumps. Moreover, the third term in (59) can be rewritten as:

$$E_t \left[\int_0^T \left((\theta_s^{*,no \ jumps})^2 \sigma_s^2 - (\theta_s^*)^2 (\sigma_s^2 + \lambda_s \beta_s^2) \right) e^{2r(T-s)} ds \right],$$

where θ^* and $\theta^{*,no \ jumps}$ denote the optimal hedges with and without jumps, respectively, and hence directly captures the effect of the difference between optimal hedges with and without jumps on the quality of the hedge.

We next demonstrate the tractability of our analysis by computing the optimal hedge in a jump-diffusion setting employed by Liu and Pan (2003) in the context of portfolio choice. The non-tradable asset, X , follows a square-root mean-reverting process

$$dX_t = \kappa(\bar{X} - X_t)dt + \bar{\nu} \sqrt{X_t} dw_{xt}, \quad (63)$$

with $\kappa > 0$, and the stock price, S , follows a jump process

$$dS_t = (r + \bar{\mu} X_t) S_t dt + \bar{\sigma} \sqrt{X_t} S_t dw_t + \beta S_t - dq_t, \quad (64)$$

with jump intensity $\lambda_t = \bar{\lambda} X_t$, $\bar{\lambda} > 0$, $\beta > -1$, and $\bar{\sigma}^2 > \beta \bar{\mu}$ so that λ_t^* in (62) is positive. Recovering the optimal hedge amounts to finding the conditional expectation of the discounted terminal payoff X_T under the hedge-neutral measure, similar to the examples in Section 2.3. Corollary 6 reports the optimal hedge and the hedging error variance.

Corollary 6. *The optimal hedging policy and the corresponding variance of the hedging error for the mean-reverting stochastic-volatility model with jumps (63)–(64) are given by*

$$\theta_t^* = \frac{\rho\bar{\nu}\bar{\sigma}}{\bar{\sigma}^2 + \bar{\lambda}\beta^2} e^{-\left(\kappa+r+\rho\bar{\nu}\bar{\sigma}(\bar{\mu}+\bar{\lambda}\beta)/(\bar{\sigma}^2+\bar{\lambda}\beta^2)\right)(T-t)}, \quad (65)$$

$$\begin{aligned} \text{var}_t[X_T - W_T^*] &= (1 - \rho^2)\bar{\nu}^2 \int_t^T e^{-2\left(\kappa+\rho\bar{\nu}\bar{\sigma}(\bar{\mu}+\bar{\lambda}\beta)/(\bar{\sigma}^2+\bar{\lambda}\beta^2)\right)(T-s)} \left(\bar{X} + (X_t - \bar{X})e^{-\kappa(s-t)}\right) ds \\ &\quad + \int_0^T \left((\rho\bar{\nu})^2 e^{-2\left(\kappa+\rho\bar{\nu}\bar{\mu}/\bar{\sigma}\right)(T-s)} - \frac{(\rho\bar{\nu}\bar{\sigma})^2}{\bar{\sigma}^2 + \bar{\lambda}\beta^2} e^{-2\left(\kappa+\rho\bar{\nu}\bar{\sigma}(\bar{\mu}+\bar{\lambda}\beta)/(\bar{\sigma}^2+\bar{\lambda}\beta^2)\right)(T-s)} \right) \\ &\quad \times \left(\bar{X} + (X_t - \bar{X})e^{-\kappa(s-t)}\right) ds. \end{aligned} \quad (66)$$

Corollary 6 explicitly reveals the effect of the jump intensity parameter $\bar{\lambda}$ and the jump size β on the optimal hedge (65) and the hedging error variance (66). In particular, it can easily be demonstrated that the optimal hedge is a decreasing function of the intensity parameter $\bar{\lambda}$. This is because the presence of jumps increases the extent of market incompleteness, and hence reduces the attractiveness of stocks for hedging. The comparative statics with respect to other model parameters can be analyzed along the lines of Section 2.3.

5. Conclusion

This work tackles the problem of dynamic hedging in incomplete markets and provides tractable optimal hedges according to the traditional minimum-variance criterion over the hedging error. The optimal hedges are shown to retain both the simple structure of complete-market hedges and the intuitive features of static hedges, and are in terms of the familiar Greeks, widely employed in risk management applications. Moreover, in contrast to the existing literature, the hedges are derived via dynamic programming and hence are time-consistent. The baseline analysis is then applied to study the replication and hedging of derivatives when trading occurs at discrete periods of time, rendering the market incomplete. In this discrete-time setting, a generalized Black-Scholes formula for the value of the replicating portfolio is obtained in closed form. Moreover, it is demonstrated that the optimal discrete time hedges significantly deviate from the classic Black-Scholes hedges in the presence of return predictability. 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 (2010) 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

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

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

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

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 , $G_t e^{r(T-t)}$ and $W_t e^{r(T-t)}$, substituting them into (A.4) and computing the variance term, we obtain the equation

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

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

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

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

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

with the terminal condition $J_T = 0$. The Feynman-Kac solution (Karatzas and Shreve, 1991) to equation (A.48) is then given by

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

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

$$\frac{\partial G_t}{\partial t} + \left(m_t - \rho \nu_t \frac{\mu_t - r}{\sigma_t} \right) X_t \frac{\partial G_t}{\partial X_t} + r S_t \frac{\partial G_t}{\partial S_t} + \frac{1}{2} \left(\nu_t^2 X_t^2 \frac{\partial^2 G_t}{\partial X_t^2} + 2 \rho \nu_t \sigma_t X_t S_t \frac{\partial^2 G_t}{\partial X_t \partial S_t} + \sigma_t^2 S_t^2 \frac{\partial^2 G_t}{\partial S_t^2} \right) - r G_t = 0, \quad (\text{A.10})$$

with the terminal condition $G_T = X_T$. Its Feynman-Kac solution is then given by

$$G_t = E_t^* [X_T e^{-r(T-t)}], \quad (\text{A.11})$$

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.11) 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 $dw_{X_t} = \rho dw_t + \sqrt{(1 - \rho^2)} dw_t^\perp$, where $w_t^\perp \equiv (w_{X_t} - \rho w_t) / \sqrt{(1 - \rho^2)}$ 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)

$$E_t \left[W_T^* - W_t e^{r(T-t)} \right] = E_t \left[\int_t^T \theta_s^* (\mu_s - r) e^{r(T-s)} ds \right], \quad (\text{A.12})$$

and then substituting (A.11) and (A.12) into (A.5).

Q.E.D.

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

$$dX_t = (\kappa + \rho \bar{\nu}) \left(\frac{\kappa \bar{X}}{\kappa + \rho \bar{\nu}} - X_t \right) dt + \bar{\nu} dw_{X_t}^*, \quad (\text{A.13})$$

for which the conditional moments are well-known (e.g., Vasicek, 1977), yielding

$$E_t^* [X_T] = \frac{\kappa \bar{X}}{\kappa + \rho \bar{\nu}} + \left(X_t - \frac{\kappa \bar{X}}{\kappa + \rho \bar{\nu}} \right) e^{-(\kappa + \rho \bar{\nu})(T-t)}.$$

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

$$dX_t = \left(\kappa + \rho\bar{\nu}(\mu - r) \right) \left(\frac{\kappa\bar{X}}{\kappa + \rho\bar{\nu}(\mu - r)} - X_t \right) dt + \bar{\nu}\sqrt{X_t}dw_{Xt}^*. \quad (\text{A.14})$$

The conditional expectation of X_T is well-known (e.g., Cox, Ingersoll, and Ross, 1985) to be

$$E_t^*[X_T] = \frac{\kappa\bar{X}}{\kappa + \rho\bar{\nu}(\mu - r)} + \left(X_t - \frac{\kappa\bar{X}}{\kappa + \rho\bar{\nu}(\mu - r)} \right) e^{-(\kappa + \rho\bar{\nu}(\mu - r))(T-t)}.$$

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:

$$0 = E_t \left[d \text{var}_s(X_T - W_T) + \text{var}_s(dE_s[X_T - W_T]) \right]. \quad (\text{A.15})$$

Integrating (A.15) from t to T yields

$$\text{var}_t[X_T - W_T] = E_t \left[\int_t^T \frac{\text{var}_s(dE_s[X_T - W_T])}{ds} ds \right]. \quad (\text{A.16})$$

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

$$dE_t[X_T - W_T] = (\dots)dt + \nu_t X_t \frac{\partial E_t[X_T]}{\partial X_t} dw_{Xt} + \sigma_t S_t \frac{\partial E_t[X_T]}{\partial S_t} dw_t - \theta_t \sigma_t e^{r(T-t)} dw_t. \quad (\text{A.17})$$

Substituting (A.17) into (A.16) and computing $\text{var}_s(dE_t[X_T - W_T])$, we obtain:

$$\text{var}_t[X_T - W_T] = E_t \left[\int_t^T \left(\theta_s \sigma_s e^{r(T-s)} - \rho \nu_s X_s \frac{\partial E_s[X_T]}{\partial X_s} - \sigma_s S_s \frac{\partial E_s[X_T]}{\partial S_s} \right)^2 + (1 - \rho^2) \left(\frac{\partial E_s[X_T]}{\partial X_s} \right)^2 ds \right]. \quad (\text{A.18})$$

Minimizing the expression under the integral in (A.18) 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.16) 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 Remark 3. The optimal pre-commitment hedge (26) for the case of $r = 0$ and $W_0 = 0$ has been obtained by Duffie and Richardson (1991) in the context of futures hedging.¹⁸ 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

$$d\tilde{W}_t = \theta_t \tilde{\mu}_t dt + \theta_t \tilde{\sigma}_t dw_t, \quad (\text{A.19})$$

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.19) reduces to the case with $r = 0$ and $\tilde{W}_0 = 0$, and hence the pre-commitment hedge (26) is easily obtained from the solution in Duffie and Richardson.

Next, 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

$$E_t^*[X_T] = X_t e^{(m - \rho\nu(\mu - r)/\sigma)(T-t)}. \quad (\text{A.20})$$

Consequently, by rearranging terms in (10) and substituting (A.20) we obtain:

$$E_t[(X_T - W_T^*)e^{-r(T-t)}] = X_t e^{(m - r - \rho\nu(\mu - r)/\sigma)(T-t)} - W_t. \quad (\text{A.21})$$

Hence, the second term in (26) hedges the deviations $E_t[(X_T - W_T^*)e^{-r(T-t)}]$ from its time-zero value (compounded by 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)). *Q.E.D.*

Proof of Proposition 3. Analogously to the proof of Proposition 1, we derive the HJB equation for our discrete-time problem as:

$$J_t = \min_{\theta_t} E_t[J_{t+\Delta t}] + e^{(T-t-\Delta t)r} \text{var}_t \left[G_{t+\Delta t} - \frac{S_{t+\Delta t}}{S_t} \theta_t \right]. \quad (\text{A.22})$$

Similarly to Proposition 2, the value function J_t does not depend on W_t . Therefore, solving the minimization problem in (A.22) we obtain the optimal hedge as:

$$\theta_t^* = \frac{\text{cov}_t(G_{t+\Delta t}, S_{t+\Delta t}/S_t)}{\text{var}_t[S_{t+\Delta t}/S_t]}. \quad (\text{A.23})$$

Furthermore, solving the HJB equation (A.22) recursively we obtain the hedging error variance $J_t = \text{var}_t[X_T - W_T^*]$ as reported in (32).

Next, from the definition of G_t in (28) we obtain the following recursive representation:

$$G_t = e^{-r\Delta t} E_t[G_{t+\Delta t}] - E_t \left[W_{t+\Delta t} e^{-r\Delta t} - W_t \right]. \quad (\text{A.24})$$

¹⁸For 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.

Substituting $W_{t+\Delta t}$ from (27) into (A.24) we determine the optimal hedge (29) as reported in Proposition 3. Moreover, substituting $W_{t+\Delta t}$ and then θ_t^* from (A.23) into (A.24), after some simple algebra we obtain the following discrete-time analogue of the PDE in (A.10):

$$G_t = E_t \left[\left\{ 1 + \left(E_t \left[\frac{S_{t+\Delta t}}{S_t} \right] - \frac{S_{t+\Delta t}}{S_t} \right) \frac{E_t[S_{t+\Delta t}/S_t] - e^{r\Delta t}}{\text{var}_t[S_{t+\Delta t}/S_t]} \right\} G_{t+\Delta t} e^{-r\Delta t} \right]. \quad (\text{A.25})$$

Iterating the recursive equation (A.25) and noting that $G_T = X_T$ we obtain an analytical expression for G_t as reported in (30)–(31), which is a discrete-time analogue of the Feynman-Kac formula.

Finally, we demonstrate that the process $e^{-rt}\xi_t S_t$ is a martingale – the proof of ξ_t being a martingale follows analogously. Using (31) for $e^{-rt}\xi_t S_t$ we obtain:

$$\begin{aligned} E_t[e^{-r(t+\Delta t)}\xi_{t+\Delta t}S_{t+\Delta t}] &= e^{-r(t+\Delta t)}\xi_t S_t \\ &\quad \times E_t \left\{ \frac{S_{t+\Delta t}}{S_t} + \frac{S_{t+\Delta t}}{S_t} \left(\frac{E_t[S_{t+\Delta t}/S_t] - S_{t+\Delta t}/S_t}{\sqrt{\text{var}_t[S_{t+\Delta t}/S_t]}} \right) \left(\frac{E_t[S_{t+\Delta t}/S_t] - e^{r\Delta t}}{\sqrt{\text{var}_t[S_{t+\Delta t}/S_t]}} \right) \right\} \\ &= e^{-r(t+\Delta t)}\xi_t S_t \left\{ E_t \left(\frac{S_{t+\Delta t}}{S_t} \right) - \left(\frac{\text{var}_t[S_{t+\Delta t}/S_t]}{\sqrt{\text{var}_t[S_{t+\Delta t}/S_t]}} \right) \left(\frac{E_t[S_{t+\Delta t}/S_t] - e^{r\Delta t}}{\sqrt{\text{var}_t[S_{t+\Delta t}/S_t]}} \right) \right\} \\ &= e^{-rt}\xi_t S_t. \end{aligned}$$

Q.E.D.

Proof of Proposition 4. Since the stock price S follows a GBM process, the stock log-returns $\varepsilon_k = \ln(S_{t+k\Delta t}/S_{t+(k-1)\Delta t})$, $k \geq 1$, are i.i.d. and $\varepsilon_k \sim N(\hat{\mu}, \hat{\sigma}^2)$, where $\hat{\sigma} = \sigma\sqrt{\Delta t}$ and $\hat{\mu} = (\mu - 0.5\sigma^2)\Delta t$. Therefore, the horizon stock price is given by:

$$S_T = S_t e^{\varepsilon_1 + \dots + \varepsilon_n}, \quad (\text{A.26})$$

where $\varepsilon_i \sim N(\hat{\mu}, \hat{\sigma}^2)$ and $n = (T-t)/\Delta t$. Using the law of iterated expectations we rewrite formula (30) for G_t in terms of the conditional expectation $F_n(x) = E[\xi_T/\xi_t | \varepsilon_1 + \dots + \varepsilon_n = x]$:

$$G_t = E_t \left[F_n(x) (S_t e^x - K)^+ \right], \quad (\text{A.27})$$

where $x \sim N(n\hat{\mu}, n\hat{\sigma}^2)$. Substituting ξ_t from (31) into the definition of $F_n(x)$, noting that $S_{t+k\Delta t}/S_{t+(k-1)\Delta t} = e^{\varepsilon_k}$, and opening the brackets we obtain:

$$F_n(x) = \sum_{k=0}^n \binom{n}{k} a^{n-k} b^k F_{nk}(x), \quad \text{where } F_{nk}(x) = E \left[e^{\varepsilon_1 + \dots + \varepsilon_k} | \varepsilon_1 + \dots + \varepsilon_n = x \right], \quad (\text{A.28})$$

where the coefficients a and b are given by:

$$a = \left(1 + \frac{e^{\mu\Delta t}(e^{\mu\Delta t} - e^{r\Delta t})}{e^{2\mu\Delta t}(e^{\sigma^2\Delta t} - 1)} \right) e^{-r\Delta t}, \quad b = -\frac{e^{\mu\Delta t} - e^{r\Delta t}}{e^{2\mu\Delta t}(e^{\sigma^2\Delta t} - 1)} e^{-r\Delta t}. \quad (\text{A.29})$$

To compute $F_{nk}(x)$ explicitly, we first rewrite it as $F_{nk}(x) = E[e^y | y+z = x]$, where $y \sim N(k\hat{\mu}, k\hat{\sigma}^2)$, $z \sim N((n-k)\hat{\mu}, (n-k)\hat{\sigma}^2)$, $x \sim N(n\hat{\mu}, n\hat{\sigma}^2)$, and y and z are independent. Since both y and x are

normally distributed, the conditional distribution of y conditional on x is given by (e.g., Shiryaev, 1995):

$$y|x \sim N\left(\mu_y + (x - \mu_x) \text{cov}(x, y)/\sigma_x^2, \sigma_y^2 - \text{cov}(x, y)^2/\sigma_x^2\right), \quad (\text{A.30})$$

where μ_y and μ_x denote the means, σ_y and σ_x the variances of x and y , respectively. Since $x = y + z$ and y and z are independent, $\text{cov}(x, y) = \sigma_y^2$. Substituting the expressions for μ_x , μ_y , σ_y , σ_x and $\text{cov}(x, y)$ into (A.30) we find that $y|x \sim N\left(xk/n, k(n-k)\hat{\sigma}^2/n\right)$, and hence:

$$F_{nk}(x) = E\left[e^y | y + z = x\right] = e^{kx/n + 0.5k(n-k)\hat{\sigma}^2/n}. \quad (\text{A.31})$$

Next, we determine $E[F_{nk}(x)(S_t e^x - K)^+]$ in closed form as follows:

$$\begin{aligned} E[F_{nk}(x)(S_t e^x - K)^+] &= \frac{e^{0.5k(n-k)\hat{\sigma}^2/n}}{\sqrt{2\pi n\hat{\sigma}}} \int_{-\infty}^{\infty} e^{xk/n} (S_t e^x - K)^+ e^{-\frac{(x-n\hat{\mu})^2}{2n\hat{\sigma}^2}} dx \\ &= \frac{e^{k\hat{\mu} + 0.5k\hat{\sigma}^2}}{\sqrt{2\pi n\hat{\sigma}}} \int_{-\infty}^{\infty} (S_t e^x - K)^+ e^{-\frac{(x-n\hat{\mu}-k\hat{\sigma}^2)^2}{2n\hat{\sigma}^2}} dx \\ &= e^{k\hat{\mu} + 0.5k\hat{\sigma}^2 + n\hat{\mu} + 0.5\hat{\sigma}^2 + k\hat{\sigma}^2} C_k^{BS}(S_t, t), \end{aligned} \quad (\text{A.32})$$

where C_k^{BS} is as given in Proposition 4, and the last equality in (A.32) is obtained in the same way as in the derivation of the Black-Scholes formula. Substituting (A.32) into (A.28) after some algebra we obtain the formula for G_t as reported in (36).

To prove the convergence to the Black-Scholes price, we first note that $F_n(x)$ in (A.28) can be rewritten as:

$$F_n(x) = E\left[\left(a + b e^{x/n + 0.5\hat{\sigma}^2 + i\hat{\sigma}\varepsilon/\sqrt{n}}\right)^n | x\right], \quad (\text{A.33})$$

where $i = \sqrt{-1}$, $\varepsilon \sim N(0, 1)$, ε and x are independent. Indeed, opening the brackets in (A.33) and using the property of characteristic functions for normally distributed random variables that $E[\exp(ik\hat{\sigma}\varepsilon/\sqrt{n})] = \exp(-0.5k^2\hat{\sigma}^2/n)$ we obtain (A.28) with $F_{nk}(x)$ given by (A.31). Substituting $\hat{\sigma} = \sigma\sqrt{\Delta t}$, $\hat{\mu} = (\mu - 0.5\sigma^2)\Delta t$, $\Delta t = (T - t)/n$, a and b from (A.29) into (A.33), and using Taylor's expansions we obtain:

$$\left(a + b e^{x/n + 0.5\hat{\sigma}^2 + i\hat{\sigma}\varepsilon/\sqrt{n}}\right)^n = \left(1 - \frac{\mu-r}{\sigma} \frac{\tilde{\varepsilon}\sqrt{T-t}}{n} - \frac{\mu-r}{\sigma} \frac{i\varepsilon\sqrt{T-t}}{n} + O\left(\frac{1}{n^2}\right)\right)^n, \quad (\text{A.34})$$

where $\tilde{\varepsilon} = (x - (\mu - 0.5\sigma^2)(T - t))/(\sigma\sqrt{T-t})$, and hence $\tilde{\varepsilon} \sim N(0, 1)$. From (A.34) we obtain the following upper bound:

$$\begin{aligned} \left|\left(a + b e^{x/n + 0.5\hat{\sigma}^2 + i\hat{\sigma}\varepsilon/\sqrt{n}}\right)^n\right| &< \left(1 + \frac{\mu-r}{\sigma} \frac{|\tilde{\varepsilon}|\sqrt{T-t}}{n} + \frac{\mu-r}{\sigma} \frac{|\varepsilon|\sqrt{T-t}}{n} + O\left(\frac{1}{n^2}\right)\right)^n \\ &< e^{C(|\tilde{\varepsilon}| + |\varepsilon|)}, \end{aligned} \quad (\text{A.35})$$

for a sufficiently large constant C . Moreover, the equality (A.34) implies point-wise convergence:

$$\lim_{n \rightarrow \infty} \left(a + b e^{x/n + 0.5\hat{\sigma}^2 + i\hat{\sigma}\varepsilon/\sqrt{n}}\right)^n = e^{-\frac{\mu-r}{\sigma} \tilde{\varepsilon}\sqrt{T-t} - \frac{\mu-r}{\sigma} i\varepsilon\sqrt{T-t}}. \quad (\text{A.36})$$

Since the expression on the left-hand side of (A.35) is bounded by an integrable function, by Lebesgue's Convergence Theorem (see Shiryaev, 1995) the point-wise convergence in (A.36) implies the convergence of expectations:

$$\begin{aligned}\lim_{n \rightarrow \infty} F_n(x) &= E \left[e^{-\frac{\mu-r}{\sigma} \tilde{\varepsilon} \sqrt{T-t} - \frac{\mu-r}{\sigma} i \varepsilon \sqrt{T-t}} | \tilde{\varepsilon} \right] \\ &= e^{-\frac{\mu-r}{\sigma} \tilde{\varepsilon} \sqrt{T-t} - 0.5 \left(\frac{\mu-r}{\sigma} \right)^2 (T-t)}.\end{aligned}\tag{A.37}$$

The right-hand side of the second equality in (A.37) is the state price density in the continuous-time Black-Scholes economy. Therefore, applying Lebesgue's Theorem one more time we obtain the convergence to the Black-Scholes price:

$$\begin{aligned}\lim_{n \rightarrow \infty} E \left[F_n(x) (S_t e^x - K)^+ \right] &= E \left[e^{-\frac{\mu-r}{\sigma} \tilde{\varepsilon} \sqrt{T-t} - 0.5 \left(\frac{\mu-r}{\sigma} \right)^2 (T-t)} (S_t e^{\mu(T-t) + \tilde{\varepsilon} \sigma \sqrt{T-t}} - K)^+ \right] \\ &= C^{BS}(S_t, t).\end{aligned}$$

Finally, we demonstrate that the modified Black-Scholes price C_k^{BS} corresponds to a k -period static hedge that keeps θ_k^* dollars in the stock between times t and $t + k\Delta t$ without rebalancing, and invests only in bonds thereafter. Consider the investor minimizing the variance of the hedging error with terminal wealth $W_T = e^{r(n-k)\Delta t} W_{t+k\Delta t}$, where the wealth at time $t + k\Delta t$ is given by

$$W_{t+k\Delta t} = e^{rk\Delta t} W_t + \theta_k \left(\frac{S_{t+k\Delta t}}{S_t} - e^{rk\Delta t} \right).\tag{A.38}$$

Solving this static optimization problem and following all the steps as in the multiperiod case we find that the replication value of the static hedge is given by $G_{kt} = a_k C_0^{BS} + b_k C_k^{BS}$, where a_k and b_k are some constants. Therefore, the formula for the replication value G_t in (36) can equivalently be rewritten in terms of G_{kt} , and hence the call option replication value equals the value of a portfolio of static hedges.

Q.E.D.

Proof of Remark 4 (Alternative representation for the replication value). Applying Euler's formula for the exponents of complex numbers to (A.33) we obtain:

$$\begin{aligned}F_n(x) &= E \left[\left(a + b e^{x/n + 0.5\hat{\sigma}^2} \cos\left(\frac{\hat{\sigma}\varepsilon}{\sqrt{n}}\right) + i b e^{x/n + 0.5\hat{\sigma}^2} \sin\left(\frac{\hat{\sigma}\varepsilon}{\sqrt{n}}\right) \right)^n \right] \\ &= E \left[R(x, \varepsilon)^n \left(\cos \varphi(x, \varepsilon) + i \sin \varphi(x, \varepsilon) \right)^n \right],\end{aligned}\tag{A.39}$$

where

$$R(x, \varepsilon) = \sqrt{\left(a + b e^{x/n + 0.5\hat{\sigma}^2} \cos\left(\frac{\hat{\sigma}\varepsilon}{\sqrt{n}}\right) \right)^2 + \left(b e^{x/n + 0.5\hat{\sigma}^2} \sin\left(\frac{\hat{\sigma}\varepsilon}{\sqrt{n}}\right) \right)^2},\tag{A.40}$$

$$\cos \varphi(x, \varepsilon) = \frac{a + b e^{x/n + 0.5\hat{\sigma}^2} \cos\left(\frac{\hat{\sigma}\varepsilon}{\sqrt{n}}\right)}{R(x, \varepsilon)}, \quad \sin \varphi(x, \varepsilon) = \frac{b e^{x/n + 0.5\hat{\sigma}^2} \sin\left(\frac{\hat{\sigma}\varepsilon}{\sqrt{n}}\right)}{R(x, \varepsilon)},\tag{A.41}$$

and $\hat{\sigma} = \sigma\sqrt{\Delta t}$. Now applying Moivre's formula for the powers of complex numbers to (A.39) and noting that $F_n(x)$ is a real-value function, we get:

$$\begin{aligned} F_n(x) &= E\left[R(x, \varepsilon)^n \cos(n\varphi(x, \varepsilon))\right], \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} R(x, \varepsilon)^n \cos(n\varphi(x, \varepsilon)) e^{-\frac{\varepsilon^2}{2}} d\varepsilon. \end{aligned} \quad (\text{A.42})$$

Substituting (A.42) into (A.27) we obtain the alternative representation (38). Angle $\varphi(x, \varepsilon)$ is found as the argument of the complex number $z = \cos \varphi(x, \varepsilon) + i \sin \varphi(x, \varepsilon)$, i.e. $\arg(z)$.¹⁹ *Q.E.D.*

Proof of Proposition 5. 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:

$$J(X_t, S_t, W_t, t) = E_t[W_T^* - X_T] - \frac{\gamma}{2} \text{var}_t[W_T^* - X_T]. \quad (\text{A.43})$$

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 (41), 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 (43). *Q.E.D.*

Proof of Corollary 4. We first observe that for the stock return dynamics (16), the squared market price of risk is given by $((\mu_s - r)/\sigma_s)^2 = (\mu - r)^2 X_s$. Consequently, given the structure of the optimal hedge (42), finding it explicitly reduces to computing the expectations $E_t^*[X_s]$ and $E_t^*[X_T]$ under the probability measure P^* and their derivatives with respect to X_t . Under P^* the process (15) for X follows the dynamics (A.14). The conditional expectation of X_s is then given by (e.g., Cox, Ingersoll, and Ross, 1985):

$$E_t^*[X_s] = \frac{\kappa \bar{X}}{\kappa + \rho \bar{\nu}(\mu - r)} + \left(X_t - \frac{\kappa \bar{X}}{\kappa + \rho \bar{\nu}(\mu - r)} \right) e^{-(\kappa + \rho \bar{\nu}(\mu - r))(s-t)}. \quad (\text{A.44})$$

The expression (A.44) allows us to obtain $E_t^*[(\mu_s - r)/\sigma_s]^2$ and $E_t^*[X_T]$ explicitly. Substituting the last two expressions in (42) yields the optimal hedge (44). *Q.E.D.*

¹⁹In MATLAB the argument of z is computed by function $\text{angle}(z)$, and can further be explicitly derived in terms of inverse trigonometric functions.

Proof of Proposition 6. Proposition 6 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.*

Proof of Corollary 5. We first note that the volatility matrix σ_t for the stock return dynamics (53) is given by $\sigma_t = \bar{\sigma} \text{diag}\{\sqrt{X_t}/(a_1 + b_1 X_t), \sqrt{X_t}/(a_2 + b_2 X_t)\}$, where $\text{diag}\{c_1, c_2\}$ denotes a 2×2 diagonal matrix with c_k , $k = 1, 2$, on main diagonal. Therefore, we find that $\sigma_t^{-1} = \text{diag}\{(a_1 + b_1 X_t)/\sqrt{X_t}, (a_2 + b_2 X_t)/\sqrt{X_t}\} \bar{\sigma}^{-1}$. Using this expression for σ_t^{-1} , from Proposition 6 we obtain that under the hedge-neutral measure the asset price X evolves as:

$$dX_t = \left(\kappa(\bar{X} - X_t) - \bar{\nu} \sqrt{X_t} (\rho_1, \rho_2) \text{diag}\left\{ \frac{a_1 + b_1 X_t}{\sqrt{X_t}}, \frac{a_2 + b_2 X_t}{\sqrt{X_t}} \right\} \bar{\sigma}^{-1} (\mu_1 - r, \mu_2 - r)^\top \right) dt + \bar{\nu} \sqrt{X_t} dw_{Xt}^*.$$

After some algebra involving the calculation of products of vectors and matrices in the dynamics for X we obtain:

$$dX_t = \tilde{\kappa}(\tilde{X} - X_t)dt + \bar{\nu} \sqrt{X_t} dw_{Xt}^*,$$

where $\tilde{\kappa}$ is as in Corollary 5 and $\tilde{X} = (\kappa \bar{X} - \bar{\nu}(\rho_1 \tilde{\mu}_1 a_1 + \rho_2 \tilde{\mu}_2 a_2))/\tilde{\kappa}$. The conditional expectation of X_T is then given by: $E_t^*[X_T] = \tilde{X} + (X_t - \tilde{X}) \exp\{-\tilde{\kappa}(T - t)\}$ (e.g., Vasicek, 1977). Substituting $E_t^*[X_T]$ into the expression for the optimal hedge in (48) and after some algebra, we obtain expression (54). *Q.E.D.*

Proof of Proposition 7. The proof follows the proof of Proposition 1. The integral representation of the terminal tradable wealth is given by:

$$W_T = W_t e^{r(T-t)} + \int_t^T \theta_s (\mu_s - r) e^{r(T-s)} ds + \int_t^T \theta_s \sigma_s e^{r(T-s)} dw_s + \int_t^T \theta_s - \beta_s e^{r(T-s)} dq_s. \quad (\text{A.45})$$

It can then be shown that the value function J_t satisfies the recursive equation (A.4) with the terminal condition $J_T = 0$, where G_t is defined by

$$G(X_t, S_t, W_t, t) \equiv E_t[X_T e^{-r(T-t)} - (W_T^* e^{-r(T-t)} - W_t)]. \quad (\text{A.46})$$

Substituting (A.45) for W_T^* under the optimal hedge θ_t^* into (A.46), and noting that $E_t[dq_t] = \lambda_t dt$, we obtain:

$$G(X_t, S_t, W_t, t) = E_t[X_T e^{-r(T-t)} - \int_t^T \theta_s^* (\mu_s - r + \lambda_s \beta_s) e^{r(t-s)} ds]. \quad (\text{A.47})$$

As in the proof of Proposition 1, θ_t^* , J_t and G_t do not depend on wealth W_t . Applying Itô's Lemma (e.g., Duffie, 2001) to J_t , $G_t e^{r(T-t)}$ and $W_t e^{r(T-t)}$, substituting into the recursive representation

(A.4), noting that $\text{var}_t[dq_t] = \lambda_t dt$, and computing the variance term in (A.4) we obtain the equation

$$\begin{aligned}
0 = \mathcal{D}J_t &+ \left(\nu_t^2 X_t^2 \left(\frac{\partial G_t}{\partial X_t} \right)^2 + 2\rho\nu_t\sigma_t X_t S_t \frac{\partial G_t}{\partial X_t} \frac{\partial G_t}{\partial S_t} + \sigma_t^2 S_t^2 \left(\frac{\partial G_t}{\partial S_t} \right)^2 \right) e^{2r(T-t)} \\
&+ \lambda_t \left(G(X_t(1 + \alpha_t), S_t(1 + \beta_t), t) - G(X_t, S_t, t) \right)^2 e^{2r(T-t)} \\
&+ \min_{\theta_t} \left[\theta_t^2 (\sigma_t^2 + \lambda_t \beta_t^2) - 2\theta_t \sigma_t \left(\rho\nu_t X_t \frac{\partial G_t}{\partial X_t} + \sigma_t S_t \frac{\partial G_t}{\partial S_t} \right) \right. \\
&\quad \left. - 2\theta_t \lambda_t \beta_t \left(G(X_t(1 + \alpha_t), S_t(1 + \beta_t), t) - G(X_t, S_t, t) \right) \right] e^{2r(T-t)}, \quad (\text{A.48})
\end{aligned}$$

subject to $J_T = 0$. The minimization in (A.48) yields the optimal hedge θ^* given by (58). Substituting θ^* back into the HJB equation (A.48) and invoking the Feynman-Kac theorem we obtain the optimal hedging error variance (59).

Applying the Feynman-Kac theorem to (A.47), we obtain a PDE for G_t . Substituting θ_t^* from (59) into this PDE we obtain the equation for G_t in terms of the exogenous model parameters. It can then be demonstrated that this new PDE has the Feynman-Kac solution

$$G(X_t, S_t, t) = E^*[X_T e^{-r(T-t)} | X_t, S_t],$$

where the expectation is under the unique probability measure P^* on which are defined the Brownian motions w_x^* and w and the Poisson jump process q such that under measure P^* the asset X and stock S follow the processes (60)–(61) and the jump intensity is given by (62). The expression for the Radon-Nikodym derivative follows from the analogue of Girsanov's theorem for processes with jumps (e.g., Proposition 3.1 in Jarrow and Madan, 1995). *Q.E.D.*

Proof of Corollary 6. Under the probability measure P^* , the process (63) becomes

$$dX_t = \left(\kappa + \rho\bar{\nu}\bar{\sigma} \frac{\bar{\mu} + \bar{\lambda}\beta}{\bar{\sigma}^2 + \bar{\lambda}\beta^2} \right) \left(\frac{\kappa\bar{X}}{\kappa + \rho\bar{\nu}\bar{\sigma}(\bar{\mu} + \bar{\lambda}\beta)/(\bar{\sigma}^2 + \bar{\lambda}\beta^2)} - X_t \right) dt + \bar{\nu}\sqrt{\bar{X}_t} dw_{xt},$$

for which the conditional expectation is well-known (e.g., Cox, Ingersoll and Ross, 1985):

$$\begin{aligned}
E_t^*[X_T] &= \frac{\kappa\bar{X}}{\kappa + \rho\bar{\nu}\bar{\sigma}(\bar{\mu} + \bar{\lambda}\beta)/(\bar{\sigma}^2 + \bar{\lambda}\beta^2)} \\
&+ \left(X_t - \frac{\kappa\bar{X}}{\kappa + \rho\bar{\nu}\bar{\sigma}(\bar{\mu} + \bar{\lambda}\beta)/(\bar{\sigma}^2 + \bar{\lambda}\beta^2)} \right) e^{-\left(\kappa + \rho\bar{\nu}\bar{\sigma}(\bar{\mu} + \bar{\lambda}\beta)/(\bar{\sigma}^2 + \bar{\lambda}\beta^2) \right) (T-t)}.
\end{aligned}$$

Substituting this into the expressions in Proposition 7 yields (65)–(66). *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.
- Ait-Sahalia, Y., J. Cacho-Diaz, and T.R. Hurd, 2009, "Portfolio Choice with Jumps: A Closed-Form Solution," *Annals of Applied Probability*, 19, 556-584.
- 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.
- Bakshi, G., C. Cao, and Z. Chen, 1997, "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52, 2003-2049.
- Bakshi, G., and N. Kapadia, 2003, "Delta-Hedged Gains and the Negative Volatility Market Risk Premium," *Review of Financial Studies*, 16, 527-566.
- Basak, S., and G. Chabakauri, 2010, "Dynamic Mean-Variance Asset Allocation," *Review of Financial Studies*, 23, 2970-3016.
- Bekaert, G., and C. Harvey, 1995, "Time-Varying World Market Integration," *Journal of Finance*, 55, 565-613.
- Bernanke, B.S, and M. Woodford, 1997, "Inflation Forecasts and Monetary Policy," *Journal of Money, Credit and Banking*, 29, 653-684.
- Bertsimas, D., L. Kogan, and A.W. Lo, 2000, "When is Time Continuous?" *Journal of Financial Economics*, 55, 173-204.
- 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.
- Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-654.
- Bossaerts and Hillion, 1997, "Local Parametric Analysis of Hedging in Discrete Time," *Journal of Econometrics*, 81, 243-272.
- Bossaerts and Hillion, 2003, "Local Parametric Analysis of Derivatives Pricing and Hedging," *Journal of Financial Markets*, 6, 573-605.
- Boyle and Emanuel, 1980, "Discretely Adjusted Option Hedges," *Journal of Financial Economics*, 8, 259-282.
- Brandt, M.W., 2003, "Hedging Demands in Hedging Contingent Claims," *Review of Economics and Statistics*, 85, 119-140.
- Branger, N., and C. Schlag, 2008, "Can Tests Based on Option Hedging Errors Correctly Identify Volatility Risk Premia?" *Journal of Financial and Quantitative Analysis*, 43, 1055-1090.
- 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.
- Buraschi, A., P. Porchia, and F. Trojani, 2010, "Correlation Risk and optimal Portfolio Choice," *Journal of Finance*, 65, 393-420.
- Campbell, J.Y., A.W. Lo, and A.C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press.
- Chacko, G., and L.M. Viceira, 2005, "Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets," *Review of Financial Studies*, 18, 1369-1402.
- 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.
- Cochrane, J.H., 2001, "Long-Term Debt and optimal Policy in the Fiscal Theory of the Price Level," *Econometrica*, 69, 69-116.
- Cochrane, J.H., F.A. Longstaff, and P. Santa-Clara, 2008, "Two Trees," *Review of Financial Studies*, 21, 347-385.
- Cochrane, J., and Saa-Requejo, 2000, "Good Deal Asset Price Bounds in Incomplete Markets," *Journal of Political Economy*, 108, 79-119.
- 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.
- Cvitanić, J., 2000, "Minimizing Expected Loss of Hedging in Incomplete and Constrained Markets," *SIAM Journal of Control and Optimization*, 38, 1050-1066.
- Cvitanić, J., and F. Zapatero, 2004, *Introduction to the Economics and Mathematics of Financial Markets*, First Edition, The MIT Press.
- Detemple, J., R. Garcia, and M. Rindisbacher, 2003, "A Monte Carlo Method for Optimal Portfolios," *Journal of Finance*, 58, 401-446.
- Detemple, J., and M. Rindisbacher, 2005, "Closed-Form Solutions for Optimal Portfolio Selection with Stochastic Interest Rates," *Mathematical Finance*, 15, 539-568.
- Detemple, J., and S. Sundaresan, 1999, "Nontraded Asset Valuation with Portfolio Constraints: A Binomial Approach," *Review of Financial Studies*, 12, 835-872.
- Dixit, R.K., and R.S. Pindyck, 1994, *Investment under Uncertainty*, Princeton University Press.
- Driessen, J., P.J. Maenhout, and G. Vilkov, 2009, "The Price of Correlation Risk: Evidence from Equity Options," *Journal of Finance*, 64, 1377-1405.
- Duffie, D., 1989, *Futures Markets*, Prentice Hall.
- Duffie, D., 2001, *Dynamic Asset Pricing Theory*, Third Edition, Princeton University Press.
- 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.
- Evans, G.W., and B. McGough, 2007, "Optimal Constrained Interest-Rate Rules," *Journal of Money, Credit and Banking*, 39, 1335-1356.
- Fama, E., and k. French, 1988, "Permanent and Temporary Components of Stock Prices," *Journal of Political Economy*, 96, 246-273.
- Figlewski, S., 1984, "Hedging Performance and Basis Risk in Stock Index Futures," *Journal of Finance*, 39, 657-669.
- Figlewski, S., 1989, "Options Arbitrage in Imperfect Markets," *Journal of Finance*, 44, 1289-1311.
- 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.
- Gomez, J.-P., and F. Zapatero, 2003, "Asset Pricing Implications of Benchmarking: a Two-Factor CAPM," *European Journal of Finance*, 9, 343-357.
- Gourieroux C., L.P. Laurent, and H. Pham, 1998, "Mean-Variance Hedging and Numeraire," *Mathematical Finance*, 8, 179-200.
- Grenadier, S.R., and N. Wang, 2007, "Investment under Uncertainty and Time-Inconsistent Preferences," *Journal of Financial Economics*, 84, 2-39.
- Harris, C., and D. Laibson, 2001, "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 69, 935-957.
- Henderson, V., 2005, "Explicit Solutions to an Optimal Portfolio Choice Problem with Stochastic Income," *Journal of Economic Dynamics and Control*, 29, 1237-1266.
- Heston, S.L., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327-343.
- Hirshleifer, D., 1988, "Residual Risk, trading Costs, and Commodity Futures Risk Premia," *Review of Financial Studies*, 1, 173-193.
- 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.

- Jameson M., and W. Wilhelm, 1992, "Market Making in the options Markets and the Costs of Discrete hedge Rebalancing," *Journal of Finance*, 47, 765-779.
- Jarrow, R., and D. Madan, 1995, "Option Pricing Under the Term Structure of Interest Rates to Hedge Systematic Discontinuities in Asset Returns," *Mathematical Finance*, 5, 311-336.
- 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.
- Leland, H.E., 1985, "Option pricing and Replication with Transaction Costs," *Journal of Finance*, 40, 1283-1301.
- Lioui, A., and P. Poncet, 2000, "The Minimum Variance Hedge Ratio Under Stochastic Interest Rates," *Management Science*, 46, 658-668.
- Liu, J., 2007, "Portfolio Selection in Stochastic Environments," *Review of Financial Studies*, 20, 1-39.
- Liu, J., F.A. Longstaff, and J. Pan, 2003, "Dynamic Asset Allocation with Event Risk," *Journal of Finance*, 58, 231-259.
- Liu, J., and J. Pan, "Dynamic Derivative Strategies," *Journal of Financial Economics*, 69, 401-430.
- Lo, A.W., and C. MacKinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies*, 1, 41-66.
- Lo, A.W., and C. MacKinlay, 1990, "When Are Contrarian Profits Due to Stock Market Overreaction?" *Review of Financial Studies*, 3, 175-206.
- Lo, A.W., and J. Wang, 1995, "Implementing Option Pricing Models when Asset Returns are Predictable," *Journal of Finance*, 50, 87-129.
- McDonald, R., 2006, *Derivatives Markets*, Second Edition, Addison Wesley.
- Moskowitz, T., 2003, "An Analysis of Covariance Risk and Pricing Anomalies," *Review of Financial Studies*, 16, 417-457.
- Musiela, M., and M. Rutkowski, 1998, *Martingale Methods in Financial Modelling*, Springer-Verlag.
- Pavlova, A., and R. Rigobon, 2008, "The Role of Portfolio Constraints in the International Propagation of Shocks," *Review of Economic Studies*, 75, 1215-1256.
- Peleg, B., and M.E. Yaari, 1973, "On the Existence of Consistent Course of Action when Tastes are Changing," *Review of Economic Studies*, 40, 391-401.

- Pham, H., T. Rheinlander, and M. Schweizer, 1998, "Mean-Variance Hedging for Continuous Processes: New Proofs and Examples," *Finance and Stochastics*, 2, 173-198.
- Pindyck, R.S., 2004, "Volatility and Commodity Price Dynamics," *Journal of Futures Markets*, 24, 1029-1047.
- Rolfo, J., 1980, "Optimal Hedging under Price and Quantity Uncertainty: The Case of a Cocoa Producer," *Journal of Political Economy*, 88, 100-116.
- Roll, R., 1992, "A Mean-Variance Analysis of Tracking Error," *Journal of Portfolio Management*, 18, 13-23.
- Sargent, T.J., and N. Wallace, 1975, "Rational Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule," *Journal of Political Economy*, 83, 241-254.
- 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.
- Shiryaev, A., 1995, *Probability*, Second Edition, Springer.
- 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., 2000, "Open-Economy Inflation Targeting," *Journal of International Economics*, 50, 155-183.
- 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.
- Vasicek, O., 1977, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177-188.
- Viceira, L.M., 2001, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income," *Journal of Finance*, 56, 433-470.
- Wilmott, P., 2006, *Paul Wilmott on Quantitative Finance*, Second Edition, John Wiley & Sons.