Utility-based valuation and hedging of basis risk with partial information

Michael Monoyios

Mathematical Institute, University of Oxford

May 20, 2010

Abstract

We analyse the valuation and hedging of a claim on a non-traded asset using a correlated traded asset under a partial information scenario, when the asset drifts are unknown constants. Using a Kalman filter and a Gaussian prior distribution for the unknown parameters, a full information model with random drifts is obtained. This is subjected to exponential indifference valuation. An expression for the optimal hedging strategy is derived. An asymptotic expansion for small values of risk aversion is obtained via PDE methods, following on from payoff decompositions and a price representation equation. Analytic and semi-analytic formulae for the terms in the expansion are obtained when the minimal entropy measure coincides with the minimal martingale measure. Simulation experiments are carried out which indicate that the filtering procedure can be beneficial in hedging, but sometimes needs to be augmented with the increased option premium, that takes into account parameter uncertainty, in order to be effective. Empirical examples are presented which conform to these conclusions.

1 Introduction

This article analyses the optimal valuation and hedging of a contingent claim in an incomplete market under a partial information scenario. The hedger does not know the values of the assets' expected returns, which are filtered from price observations, and trading strategies are required to be adapted to the asset price filtration. We assume volatilities and correlation are known constants, so we assume approximately continuous price monitoring. We make this approximation as drift parameter uncertainty is much more severe than covariance uncertainty, as discussed by Rogers [27] in an optimal investment context, and by Monoyios [22] in an optimal hedging problem.

Partial information problems have usually been studied in the context of optimal investment (Rogers [27], Lakner [19], Brendle [8], and Björk, Davis and Landén [7]). Nagai and Peng [24] treat risk-sensitive control, while Pham [25] and Xiong and Zhou [28] study mean-variance portfolio problems. Utility-based hedging of claims under partial information has received little attention, though some mention of partial information pricing was made in Dufresne and Hugonnier [9].

The incomplete market is a basis risk model in which a claim on a non-tradeable asset Y is hedged using a correlated traded stock S. This has been studied by many authors in the completely observable case, often when asset prices follow log-Brownian motions with constant parameters (Davis [10], Henderson [15], Monoyios [21, 22]) or when some parameters depend on the non-traded asset only (Musiela and Zariphopoulou [23] and, in a multi-dimensional case, Imkeller and co-authors [1, 2, 3]). In the scalar versions of these models, an explicit nonlinear expectation representation for the exponential utility-based price is available.

In our approach, significant differences arise. We begin with lognormal processes for S, Y, but with unknown constant drifts, hence considered as random variables with some prior distribution at time zero. After filtering the drifts from price observations, the resulting full information model written in the observation filtration has stochastic drifts which depend on

both asset prices. This is a departure from the models in the literature and the explicit results of the full information case are no longer available.

Preliminary ideas on using filtering methods to deal with drift parameter uncertainty in the basis risk model were outlined by Monoyios [22]. That analysis used two separate one-dimensional filters for each asset drift, and the indifference valuation and hedging program was not carried out. Here, we use a two-dimensional filter, thus introducing coupling between the optimal filters connected with the drifts of the assets. The filtered drift of S can depend on Y and vice versa. The rather specialised case in [22] is not, in general, obtained in the full two-dimensional analysis.

We carry out the utility-based valuation and hedging program in the derived completely observable model with random drifts. We derive an optimal hedging formula (Theorem 1), giving the hedge ratio in terms of derivatives of the indifference price with respect to both asset prices. This is a modification of the full information result (which contains derivatives with respect to the non-traded factors only). The extra term reflects additional risk induced by drift parameter uncertainty.

Using the PDE satisfied by the indifference price we derive some payoff decompositions (Lemma 1, Corollary 1) and a price representation equation (Corollary 2) written under the minimal entropy martingale measure (MEMM) Q^E. Mania and Schweizer [20] (see also Becerer [5] and Kallsen and Rheinländer [17]) obtain similar (but less explicit) results using backward stochastic differential equation methods in a continuous semimartingale model. Our results are explicit, in that the terms appearing in the relations are identified in terms of the indifference price and its derivatives, given the more concrete setting of our model. These results are used to derive an asymptotic expansion (Theorem 2) of the indifference price, to linear order in the risk aversion, with the zeroth order term being the marginal price (the zero risk aversion limit of the indifference price). Small risk aversion results (of a slightly different form) for indifference valuation have been obtained by Kramkov and Sirbu [18], but for a utility function defined on the positive real line, so not directly applicable here.

The asymptotic expansion relies on identifying the MEMM Q^E. In general, this involves the solution of a stochastic control problem. To obtain more explicit formulae and carry out numerical experiments, we specialise in Section 5 to the case where $Q^E = Q^M$, the minimal martingale measure. This corresponds to the case where the prior variance of S is less than or equal to that of Y. We show (Theorem 3) that, even with stochastic risk premia generated by the Kalman filter, the distribution of the terminal value of the non-traded asset is lognormal, with a time-dependent variance and mean that is a function of both asset prices and time. From this we derive new analytic formulae for most of the terms in the price expansion, involving BS-style formulae with a random dividend yield that is a function of time and the asset prices, and a time-dependent volatility. One term in the first order correction is not obtainable in closed form. This is the Q^M-expectation of the quadratic variation of the gains process of the marginal hedging strategy (the optimal strategy in the zero risk aversion limit), and so is computed via simulation.

We investigate numerically whether utility-based valuation and hedging, when coupled with learning based on filtering the asset drifts, can outperform other techniques. We test against (i) the optimal price and hedge in the absence of filtering, implemented using the full information results with the initial values of the asset drifts (so not updated by filtering), (ii) the perfect correlation Black-Scholes (BS)-style price and hedge, and (iii) its correlation-weighted modification, proposed by Hulley and McWalter [16], in which the latter two methods do not require knowledge of the drift parameters. As a calibration, we also use the genuine full information hedge, using the true asset drifts, once again fixed through the hedging interval. We simulate many asset price histories and compute the distribution of terminal hedging errors, with the hedging programs starting at the same initial wealth. We also compute the effect of using the appropriate option premium (allied to the particular hedging method) as the initial wealth. The latter computation is designed to show the combined effect of *valuation* as well as hedging on the final profit and loss.

The results indicate that filtering the asset price drifts can often improve the terminal profit and loss distribution, but there are instances when the prior is well suited to producing

good results without filtering, or when the correlation-weighted BS hedge is also a good hedge, because its approximation of zero drift for the underlying asset under the minimal measure is a fortuitous one. The partial information method almost always gives an improvement when one combines the filtered hedge with the appropriate valuation at time zero for the claim. This can be traced to the fact that the parameter uncertainty induces a higher effective volatility, so the agent charges a higher premium for the claim than without parameter uncertainty.

Finally, we show some examples with real data, hedging a basket of stocks with index futures, and these conform to our earlier conclusions.

The rest of the paper is organised as follows. In Section 2 we set up the model and use the Kalman-Bucy filter to convert the partial information model to a full information model with random drifts, given in Proposition 1. In Section 3 we carry out exponential indifference valuation in the derived full information model, present the optimal hedging formula (Theorem 1), and discuss the required smoothness of the dual value function, required for the validity of our results. In Section 4 we derive the payoff decompositions and price representation results leading to the asymptotic expansion (Theorem 2) for the indifference price. Section 5 derives analytic formulae for all but one term in the price expansion in the case when $Q^E = Q^M$. Numerical experiments are carried out in Section 6, and some empirical examples in Section 7. Section 8 concludes.

2   The model

The setting is a probability space $(\Omega, \mathcal{F}, P)$, equipped with a filtration $\mathbb{F} := (\mathcal{F}_t)_{0 \leq t \leq T}$ carrying a two-dimensional Brownian motion (B^S, Z^S). Define a Brownian motion B^Y correlated with B^S according to

$$B^Y := \rho B^S + \sqrt{1 - \rho^2} Z^S,$$

with $\rho \in [-1, 1]$ a known constant.

A traded stock price $S := (S_t)_{0 \leq t \leq T}$ and a non-traded asset price $Y := (Y_t)_{0 \leq t \leq T}$ follow

$$dS_t = \sigma^S S_t (\lambda^S dt + dB_t^S), \quad dY_t = \sigma^Y Y_t (\lambda^Y dt + dB_t^Y),$$

where $\sigma^S > 0$ and $\sigma^Y > 0$ are known constants.

The risk premia λ^S, λ^Y are $\mathcal{F}_0$-measurable random variables, so would be known constants if a financial agent had access to the filtration $\mathbb{F}$. The new feature in this article is that an agent will be required to use strategies adapted to the observation filtration generated by the asset prices, so λ^S, λ^Y will be unknown constants whose values will be filtered from price observations. For simplicity, the interest rate is taken to be zero.

In taking σ^S, σ^Y, ρ as known, we imply that they could be inferred from quadratic and cross-variations, so this is an approximation of continuous asset price monitoring. We make this assumption to focus on the more severe problem of drift uncertainty, though it would be interesting to study similar problems with discrete monitoring of asset prices, and unknown volatilities and correlation as well as uncertain drifts.

A European contingent claim pays the non-negative random variable $C(Y_T)$ at time T, where $C(\cdot)$ is a bounded continuous function.

As is well-known, if the correlation is perfect then the claim can be perfectly replicated by a BS-style hedge that does not require knowledge of the asset price drifts, and so is robust with respect to drift parameter uncertainty.

In the completely observable incomplete case exponential utility-based valuation has been studied by a number of authors. In [21, 22] it is shown that optimal valuation combined with hedging is beneficial compared with the BS approach, in terms of the terminal hedging error distribution produced by selling the claim at the appropriate price (the indifference price or the BS price) and investing the proceeds in the corresponding hedging portfolio. If one focuses exclusively on the hedging strategy, so begins with the same initial wealth for both strategies, Hulley and McWalter [16] provide evidence that the the improved performance is not always guaranteed, signifying that utility indifference methods rely to some extent on their initial valuation of the derivative to be effective. We shall see examples of this in Section 6.

We focus here on the issue of drift parameter uncertainty, and the fact that indifference pricing requires perfect knowledge of λ^S, λ^Y, which are virtually impossible to estimate accurately. This can ruin the effectiveness of indifference methods, as shown in [22]. We relax the assumption that the agent knows the values of λ^S, λ^Y, and investigate if optimal valuation and hedging is made feasible in this case by filtering the asset price drifts.

The return parameters λ^S, λ^Y are hence modelled as random variables with some prior distribution. Define the processes ξ^S, ξ^Y by

$$\xi_t^S := \frac{1}{\sigma^S} \int_0^t \frac{dS_u}{S_u} = \lambda^S t + B_t^S, \quad \xi_t^Y := \frac{1}{\sigma^Y} \int_0^t \frac{dY_u}{Y_u} = \lambda^Y t + B_t^Y, \quad 0 \le t \le T.$$

Using the Itô formula these may be expressed as deterministic functions of the asset prices, given by

$$\xi_t^S \equiv \xi^S(t, S_t) = \frac{1}{\sigma^S} \log\left(\frac{S_t}{S_0}\right) + \frac{1}{2}\sigma^S t, \quad \xi_t^Y \equiv \xi^Y(t, Y_t) = \frac{1}{\sigma^Y} \log\left(\frac{Y_t}{Y_0}\right) + \frac{1}{2}\sigma^Y t. \quad (1)$$

For brevity of notation we shall often write $\zeta_t \equiv \zeta(t, S_t, Y_t)$ for any process ζ that is a function of time and current asset prices, whenever no confusion arises.

We regard the two-dimensional process

$$\Xi_t := \begin{pmatrix} \xi_t^S \\ \xi_t^Y \end{pmatrix}, \quad 0 \le t \le T,$$

as an "observation process" in a Kalman-Bucy filtering framework, corresponding to noisy observations of the "signal process" (in our case, an unknown constant) Λ, defined by

$$\Lambda := \begin{pmatrix} \lambda^S \\ \lambda^Y \end{pmatrix}.$$

Define the observation filtration $\widehat{\mathbb{F}} := (\widehat{\mathcal{F}}_t)_{0 \le t \le T}$ by

$$\widehat{\mathcal{F}}_t = \sigma(\xi_u^S, \xi_u^Y; 0 \le u \le t), \quad 0 \le t \le T.$$

We assume a Gaussian prior distribution, given by

$$\text{Law}(\Lambda | \widehat{\mathcal{F}}_0) = \text{N}(\Lambda_0, V_0), \quad (2)$$

with

$$\Lambda_0 = \begin{pmatrix} \lambda_0^S \\ \lambda_0^Y \end{pmatrix}, \quad V_0 = \begin{pmatrix} v_0^S & c_0 \\ c_0 & v_0^Y \end{pmatrix}, \quad c_0 = \rho \min(v_0^S, v_0^Y), \quad (3)$$

for given constants $\lambda_0^S, \lambda_0^Y, v_0^S, v_0^Y$, which may be determined as described below.

A motivation for this prior is the idea that an agent might use data before time zero to make a point estimate of Λ, and could then use the distribution of the estimator as the prior. With historical data for ξ^S (respectively, ξ^Y) over a time interval interval t_S (respectively, t_Y), then an unbiased estimator of Λ is Gaussian according to (2) and (3) with λ_0^i equal to the point estimate of λ^i, and $v_0^i = 1/t_i$, for $i = S, Y$. Hence, we shall suppose that Λ, considered as a random variable, is bivariate normal according to (2) and (3). This distribution will be updated via subsequent observations of ξ^S, ξ^Y over the hedging interval $[0, T]$.

2.1 Two-dimensional Kalman-Bucy filter

We are firmly within the realm of a two-dimensional Kalman filtering problem. In [22], Monoyios outlined preliminary ideas connecting the partial information basis risk model to a Kalman filtering problem, and two one-dimensional Kalman-Bucy filters were used to obtain the filtered random drifts. Here, we properly treat the filtering problem with a two-dimensional Kalman filter.

Remark 1. In principle, one could model the unknown risk premia λ^S, λ^Y as processes following linear SDEs, without leaving the Kalman filtering framework, as Brendle [8] does in the context of an optimal portfolio problem. We do not pursue this here, as we are seeking maximally explicit formulae for the indifference price and optimal hedge. It is not certain at this point that this is feasible in the linear SDE framework. This is left for future research.

The observation and signal SDEs are

$$d\Xi_t = \Lambda dt + D d\mathbf{B}_t, \quad d\Lambda = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

where

$$D = \begin{pmatrix} 1 & 0 \\ \rho & \sqrt{1-\rho^2} \end{pmatrix}, \quad \mathbf{B}_t = \begin{pmatrix} B_t^S \\ Z_t^S \end{pmatrix}.$$

The optimal filter $\widehat{\Lambda}_t := E[\Lambda | \widehat{\mathcal{F}}_t], 0 \le t \le T$, is therefore the two conditional expectations

$$\widehat{\lambda}_t^i := E[\lambda^i | \widehat{\mathcal{F}}_t], \quad 0 \le t \le T, \quad i = S, Y.$$

The conditional variances and covariance are defined by

$$
\begin{aligned}
\mathbf{v}_t^i &:= E\left[(\lambda^i - \widehat{\lambda}_t^i)^2 \,\middle|\, \widehat{\mathcal{F}}_t \right], \quad 0 \le t \le T, \quad i = S, Y, \\
\mathbf{c}_t &:= E\left[(\lambda^S - \widehat{\lambda}_t^S)(\lambda^Y - \widehat{\lambda}_t^Y) \,\middle|\, \widehat{\mathcal{F}}_t \right], \quad 0 \le t \le T,
\end{aligned}
$$

and the covariance matrix will be denoted by

$$V_t := \begin{pmatrix} \mathbf{v}_t^S & \mathbf{c}_t \\ \mathbf{c}_t & \mathbf{v}_t^Y \end{pmatrix}, \quad 0 \le t \le T. \tag{4}$$

As usual with a Kalman filter, this will be a deterministic function of time. For $\rho^2 \neq 1$, define the functions

$$m_t := \min(\mathbf{v}_t^S, \mathbf{v}_t^Y), \quad M_t := \max(\mathbf{v}_t^S, \mathbf{v}_t^Y), \quad b_t := \frac{M_t - \rho^2 m_t}{1 - \rho^2}, \quad 0 \le t \le T, \tag{5}$$

and note that $b_t = m_t = M_t$ when the asset variances $\mathbf{v}_t^S = \mathbf{v}_t^Y$ are equal.

The Kalman-Bucy filter converts the partial information model to a completely observable model as given below.

Proposition 1. *On the filtered probability space $(\Omega, \widehat{\mathcal{F}}_T, \widehat{\mathbb{F}}, P)$ we have a full information model with asset price dynamics given by*

$$dS_t = \sigma^S S_t (\widehat{\lambda}_t^S dt + d\widehat{B}_t^S), \quad dY_t = \sigma^Y Y_t (\widehat{\lambda}_t^Y dt + d\widehat{B}_t^Y), \tag{6}$$

where $\widehat{B}^S, \widehat{B}^Y$ are correlated $(P, \widehat{\mathbb{F}})$-Brownian motions with correlation ρ, and $\widehat{\lambda}^S, \widehat{\lambda}^Y$ are $\widehat{\mathbb{F}}$-adapted processes, given in terms of the asset prices and time as follows.

For $i, j \in \{S, Y\}$, if $m_0 = \mathbf{v}_0^i < \mathbf{v}_0^j = M_0$, then

$$\widehat{\lambda}_t^i = \frac{\lambda_0^i + m_0 \xi_t^i}{1 + m_0 t}, \quad \widehat{\lambda}_t^j = \frac{\lambda_0^j + b_0 \xi_t^j}{1 + b_0 t} - \rho \left(\frac{\lambda_0^i + b_0 \xi_t^i}{1 + b_0 t} - \frac{\lambda_0^i + m_0 \xi_t^i}{1 + m_0 t} \right), \quad 0 \le t \le T, \quad i, j \in \{S, Y\}. \tag{7}$$

The covariance matrix V in (4) is given by

$$\mathbf{v}_t^i = m_t, \quad \mathbf{v}_t^j = M_t = \rho^2 m_t + (1 - \rho^2) b_t, \quad \mathbf{c}_t = \rho m_t, \quad 0 \le t \le T, \quad i, j \in \{S, Y\}, \tag{8}$$

with m, M, b defined in (5), and m, b given by

$$m_t = \frac{m_0}{1 + m_0 t}, \quad b_t = \frac{b_0}{1 + b_0 t}, \quad 0 \le t \le T. \tag{9}$$

In the case that $m_0 = \mathbf{v}_0^S = \mathbf{v}_0^Y = M_0$, then (7)–(9) still hold, with $b_0 = m_0$, and hence $b_t = m_t$ for all $t \in [0, T]$.

Proof. By the Kalman-Bucy filter (for example, Theorem V.9.2 in Fleming and Rishel [12]), the process $\widehat{\Lambda}$ satisfies the SDE

$$d\widehat{\Lambda}_t = V_t \left(DD^T\right)^{-1} (d\Xi_t - \widehat{\Lambda}_t dt) =: V_t \left(DD^T\right)^{-1} dN_t, \quad \widehat{\Lambda}_0 = \Lambda_0, \tag{10}$$

where N is the innovations process, defined by

$$N_t := \Xi_t - \int_0^t \widehat{\Lambda}_u du, \quad 0 \le t \le T, \tag{11}$$

and is a two-dimensional correlated $\widehat{\mathbb{F}}$-Brownian motion, given by

$$N_t = \begin{pmatrix} \widehat{B}_t^S \\ \widehat{B}_t^Y \end{pmatrix}, \quad \langle \widehat{B}^S, \widehat{B}^Y \rangle_t = \rho t, \quad 0 \le t \le T. \tag{12}$$

Using (11), (12) and the price dynamics in the form

$$d\begin{pmatrix} S_t \\ Y_t \end{pmatrix} = \begin{pmatrix} \sigma^S S_t \\ \sigma^Y Y_t \end{pmatrix} d\Xi_t,$$

gives the dynamics (6) of S, Y in the observation filtration $\widehat{\mathbb{F}}$.

The covariance matrix V_t satisfies the Riccati equation

$$\frac{dV_t}{dt} = -V_t \left(DD^T\right)^{-1} V_t,$$

with V_0 given in (3). Then $F_t := V_t^{-1}$ satisfies the Lyapunov equation

$$\frac{dF_t}{dt} = \left(DD^T\right)^{-1}.$$

Lengthy (but straightforward) calculations confirm that the Lyapunov equation is solved by (8) and (9). Using these formulae in the filtering equation (10) we find that for $i, j \in \{S, Y\}$, if $m_0 = v_0^i < v_0^j = M_0$, then

$$\begin{aligned} d\widehat{\lambda}_t^i &= m_t d\widehat{B}_t^i = m_t(d\xi_t^i - \widehat{\lambda}_t^i dt), \quad \widehat{\lambda}_0^i = \lambda_0^i, \\ d(\widehat{\lambda}_t^j - \rho\widehat{\lambda}_t^i) &= b_t(d\widehat{B}_t^j - \rho d\widehat{B}_t^i) = b_t[d(\xi_t^j - \rho\xi_t^i) - (\widehat{\lambda}_t^j - \rho\widehat{\lambda}_t^i)dt], \quad \widehat{\lambda}_0^j = \lambda_0^j. \end{aligned}$$

and in the case that $m_0 = v_0^S = v_0^Y = M_0$, these SDEs are valid with $b = m$. Solving these SDEs we obtain (7).

$\square$

Remark 2. Written explicitly, the dependence of the random risk premia on the asset prices is according to

$$\begin{aligned} &\widehat{\lambda}_t^S \equiv \widehat{\lambda}^S(t, S_t), &&\widehat{\lambda}_t^Y \equiv \widehat{\lambda}^Y(t, S_t, Y_t), &&\text{if } v_0^S < v_0^Y, \\ &\widehat{\lambda}_t^S \equiv \widehat{\lambda}^S(t, S_t), &&\widehat{\lambda}_t^Y \equiv \widehat{\lambda}^Y(t, Y_t), &&\text{if } v_0^S = v_0^Y, \\ &\widehat{\lambda}_t^S \equiv \widehat{\lambda}^S(t, S_t, Y_t), &&\widehat{\lambda}_t^Y \equiv \widehat{\lambda}^Y(t, Y_t), &&\text{if } v_0^S > v_0^Y, \end{aligned} \tag{13}$$

satisfying the SDEs

$$\begin{aligned} &d\widehat{\lambda}_t^S = m_t d\widehat{B}_t^S, &&d\widehat{\lambda}_t^Y - \rho d\widehat{\lambda}_t^S = b_t(d\widehat{B}_t^Y - \rho d\widehat{B}_t^S), &&\text{if } v_0^S < v_0^Y, \\ &d\widehat{\lambda}_t^S = m_t d\widehat{B}_t^S, &&d\widehat{\lambda}_t^Y = m_t d\widehat{B}_t^Y, &&\text{if } v_0^S = v_0^Y, \\ &d\widehat{\lambda}_t^Y = m_t d\widehat{B}_t^Y, &&d\widehat{\lambda}_t^S - \rho d\widehat{\lambda}_t^Y = b_t(d\widehat{B}_t^S - \rho d\widehat{B}_t^Y), &&\text{if } v_0^S > v_0^Y. \end{aligned} \tag{14}$$

When both asset drifts have equal prior variance, the optimal filters decouple, somewhat similarly to the situation that arises when one uses two one-dimensional Kalman filters on each asset, as done in Monoyios [22] (though one can have different prior variances for each asset with two one-dimensional filters).

In general, the asset price with smaller prior variance enters the formulae for both random risk premia. With the interpretation that the prior distribution is set using past data for the asset prices, then the longer historical dataset influences the optimal filters for both assets, and the shorter dataset only gives information on its respective asset price drift. The intuition behind this result is that estimation of the drift of a geometric Brownian motion depends only on the length of the time interval for which it is observed.

3 Exponential valuation and hedging with random drifts

On the stochastic basis $(\Omega, \widehat{\mathcal{F}}, \widehat{\mathbb{F}}, P)$, we consider exponential indifference valuation and hedging of the claim.

An agent trades the stock with $\widehat{\mathbb{F}}$-adapted strategy $\theta = (\theta_t)_{0 \le t \le T}$, an S-integrable process representing the number of shares held in the portfolio. Denote the portfolio wealth process by $X = (X_t)_{0 \le t \le T}$. For $t \in [0, T]$, given $X_t = x$, the wealth evolution is given by

$$X_u = x + \int_t^T \theta_u dS_u = x + \sigma^S \int_t^T \pi_u(\widehat{\lambda}_u^S du + d\widehat{B}_u^S), \quad t \le u \le T,$$

where $\pi := \theta S$. Denote by Θ (respectively, Π) the set of admissible θ (respectively, π), defined shortly.

The set of equivalent martingale measures is denoted by

$$\mathbb{P}_e := \{Q \sim P | S \text{ is a local } (Q, \widehat{\mathbb{F}})\text{-martingale}\}.$$

Denote by $H(Q, P)$ the relative entropy between $Q \in \mathbb{P}_e$ and P:

$$H(Q, P) := E\left[\frac{dQ}{dP} \log \frac{dQ}{dP}\right], \quad (\text{if finite, else } H(Q, P) := \infty).$$

The set of measures with finite relative entropy is denoted by

$$\mathbb{P}_{e,f} := \{Q \in \mathbb{P}_e | H(Q, P) < \infty\},$$

and we assume this set is nonempty. The set of admissible strategies is defined in a similar manner to Becherer [4] and Mania and Schweizer [20], as

$$\Theta := \{\theta | (\theta \cdot S) \text{ is a } (Q, \widehat{\mathbb{F}})\text{-martingale for all } Q \in \mathbb{P}_{e,f}\}.$$

For measures $Q \in \mathbb{P}_{e,f}$ denote the likelihood ratio process by a $(P, \widehat{\mathbb{F}})$-martingale Γ^Q:

$$\Gamma_t^Q := \left.\frac{dQ}{dP}\right|_{\widehat{\mathcal{F}}_t}, \quad 0 \le t \le T. \tag{15}$$

We assume the agent has an exponential utility function

$$U(x) := \exp(-\alpha x), \quad x \in \mathbb{R}, \quad \alpha > 0,$$

with risk aversion parameter α.

The primal value function is the maximal expected utility from terminal wealth from trading, with the additional random terminal endowment of a short position in the claim:

$$u^C(t, x, s, y) := \sup_{\pi \in \Pi} E_{t,x,s,y}[U(X_T - C(Y_T))], \tag{16}$$

where $E_{t,x,s,y}$ denotes expectation given $(X_t, S_t, Y_t) = (x, s, y)$, for $t \in [0, T]$. Denote by u^0 the value function when no claim is sold.

The indifference selling price at time $t \le T$ is $p(t, S_t, Y_t)$, where the function $p : [0, T] \times \mathbb{R}_+^2$ is defined as usual by

$$u^C(t, x + p(t, s, y), s, y) = u^0(t, x, s, y).$$

As always with exponential utility, we anticipate that the indifference price is independent of initial wealth.

Denote the optimal strategy for (16) by π^C, and the optimal strategy with no claim by π^0. The optimal hedging strategy $\pi^{(H)}$ is defined by

$$\pi^{(H)} := \pi^C - \pi^0. \tag{17}$$

Well-known duality theory for the problem (16) (see Delbaen et al [11], for example) implies that the primal value function has the representation

$$u^C(t, x, s, y) = -\exp\left(-\alpha x - H^C(t, s, y)\right), \tag{18}$$

where the function H^C originates from the dual problem to (16), and is defined by

$$H^C(t, s, y) := \inf_{Q \in \mathbb{P}_{e,f}} E^Q_{t,s,y}\left[\log\left(\frac{\Gamma^Q_T}{\Gamma^Q_t}\right) - \alpha C(Y_T)\right], \tag{19}$$

where $E^Q_{t,s,y}$ denotes Q-expectation given $(S_t, Y_t) = (s, y)$.

Denote by H^0 the function in (19) when no claim is present, or equivalently when $C \equiv 0$. This is the value function corresponding to the problem of minimising the relative entropy between $Q \in \mathbb{P}_{e,f}$ and P, so that $H^0(0, \cdot, \cdot) = H(Q^E, P)$. To emphasise the link with Q^E, we shall sometimes write $H^0 \equiv H^E$.

Applying the definition of the the indifference price and using the separable form (18) of the value function leads to the well-known entropic representation

$$p(t, s, y) = -\frac{1}{\alpha}(H^C(t, s, y) - H^0(t, s, y)). \tag{20}$$

3.1 Optimal hedging theorem

The result below is a representation for the optimal hedging strategy in terms of derivatives of the indifference price.

Theorem 1. *Suppose the indifference pricing function* $p : [0, T] \times \mathbb{R}^+$ *is of class* $C^{1,2}([0, T] \times \mathbb{R}^+)$. *Then the optimal hedge for a short position in the claim is to hold* $\theta^{(H)}_t$ *shares of* S_t *at* $t \in [0, T]$, *given by*

$$\theta^{(H)}_t = \left(\frac{\partial p}{\partial s}(t, S_t, Y_t) + \rho\frac{\sigma^Y}{\sigma^S}\frac{Y_t}{S_t}\frac{\partial p}{\partial y}(t, S_t, Y_t)\right), \quad 0 \leq t \leq T.$$

Remark 3. The required regularity of the indifference price for the validity of the theorem is established in Section 3.3. The additional term $p_s(t, S_t, Y_t)$, compared with other papers [1, 15, 21, 23] in which the drift parameters do not depend on the traded stock price, reflects the additional risk induced by parameter uncertainty. It is easy to see that the formula still holds if the volatilities and correlation are also functions of time and current asset prices.

Proof. The HJB equation associated with the primal the value function (16) is

$$\frac{\partial u^C}{\partial t} + \max_\pi \mathcal{A}_{X,S,Y} u^C = 0,$$

where $\mathcal{A}_{X,S,Y}$ denotes the generator of (X, S, Y) under P. Performing the maximisation in the HJB equation gives the optimal feedback control $\pi^C(t, s, y)$ in terms of derivatives of the value function. Then using the separable form (18) of the value function, we obtain the optimal strategy as $\pi^C_t = \pi^C(t, S_t, Y_t)$, where

$$\pi^C(t, s, y) = \frac{\widehat{\lambda}^S}{\sigma^S \alpha} - \frac{1}{\alpha}\left(sH^C_s + \rho\frac{\sigma^Y}{\sigma^S}yH^C_y\right).$$

A similar formula holds for the optimal strategy π^0 in the case when no claim is present, with H^C replaced by H^0. Applying the definition (17) of the optimal hedging strategy, we obtain

$$\pi^{(H)}(t, s, y) = -\frac{1}{\alpha}\left[s(H^C_s - H^0_s) + \rho\frac{\sigma^Y}{\sigma^S}y(H^C_y - H^0_y)\right].$$

The result now follows from the entropic representation (20) of the indifference price.

$\square$

3.2 The dual stochastic control problem

We consider the dual problem in (19) from a stochastic control perspective, to derive a PDE for $H^C(t, s, y)$. We discuss the existence of a sufficiently smooth solution to this PDE, so that the value function and indifference price are smooth enough for the hedging theorem to be valid. From the resulting PDE satisfied by the indifference price we shall derive, in the next section, payoff decompositions and an indifference price representation equation, leading to an asymptotic expansion for the indifference price, for small values of risk aversion.

Parametrise the measures $Q \in \mathbb{P}_{e,f}$ via $\widehat{\mathbb{F}}$-adapted processes ψ in the P-martingale Γ^Q of (15), according to

$$\Gamma_t^Q = \mathcal{E}(-\widehat{\lambda}^S \cdot \widehat{B}^S - \psi \cdot \widehat{Z}^S)_t, \quad 0 \leq t \leq T,$$

where $\widehat{Z}^S$ is a $(P, \widehat{\mathbb{F}})$-Brownian motion orthogonal to $\widehat{B}^S$. By the Girsanov Theorem we have a two-dimensional $(Q, \widehat{\mathbb{F}})$-Brownian motion $(\widehat{B}^{S,Q}, \widehat{Z}^{S,Q})$ defined by

$$\widehat{B}_t^{S,Q} := \widehat{B}_t^S + \int_0^t \widehat{\lambda}_u^S du, \quad \widehat{Z}_t^{S,Q} := \widehat{Z}_t^S + \int_0^t \psi_u du, \quad 0 \leq t \leq T,$$

with $\psi = 0$ corresponding to the minimal martingale measure Q^M, so that $\widehat{Z}^S$ is also a $(Q^M, \widehat{\mathbb{F}})$-Brownian motion. Then, for $Q \in \mathbb{P}_{e,f}$, we have

$$E_{t,s,y}^Q \left[\log \frac{\Gamma_T^Q}{\Gamma_t^Q} \right] = E_{t,s,y}^Q \left[\frac{1}{2} \int_t^T \left[(\widehat{\lambda}_u^S)^2 + \psi_u^2 \right] du \right] < \infty, \tag{21}$$

where the integrability condition on the right hand side is associated with the finite entropy condition and ensures that the stochastic integrals on the left hand side have zero expectation.

Let Ψ denote the set of integrands ψ such that (21) is satisfied. Then H^C in (19) is the value function of the stochastic control problem

$$H^C(t, s, y) := \inf_{\psi \in \Psi} E_{t,s,y}^Q \left[\frac{1}{2} \int_t^T \left[(\widehat{\lambda}_u^S)^2 + \psi_u^2 \right] du - \alpha C(Y_T) \right], \tag{22}$$

where, under $Q \in \mathbb{P}_{e,f}$, the state variables S, Y follow

$$\begin{array}{rcl} dS_t & = & \sigma^S S_t d\widehat{B}_t^{S,Q}, \\ dY_t & = & \sigma^Y Y_t \left[(\widehat{\lambda}_t^Y - \rho \widehat{\lambda}_t^S - \sqrt{1-\rho^2} \psi_t) dt + d\widehat{B}_t^{Y,Q} \right], \end{array} \tag{23}$$

and $\widehat{B}^{Y,Q}$ is a $(Q, \widehat{\mathbb{F}})$-Brownian motion given by $\widehat{B}^{Y,Q} = \rho \widehat{B}^{S,Q} + \sqrt{1-\rho^2} \widehat{Z}^{S,Q}$. The risk premia $\widehat{\lambda}^S, \widehat{\lambda}^Y$ are functions of the asset prices according to (13).

The HJB equation for H^C is

$$H_t^C + \mathcal{A}_{S,Y}^{Q^M} H^C + \frac{1}{2} (\widehat{\lambda}^S)^2 + \min_\psi \left[\frac{1}{2} \psi^2 - \sqrt{1-\rho^2} \sigma^Y y \psi H_y^C \right] = 0, \quad H^C(T, s, y) = -\alpha C(y), \tag{24}$$

where $\mathcal{A}_{S,Y}^{Q^M}$ is the generator of (S, Y) under Q^M. The function H^0 corresponding to finding the minimal entropy measure Q^E satisfies (24) with terminal condition $H^0(T, s, y) = 0$.

The optimal feedback control in (24) is ψ^C, given by

$$\psi^C(t, s, y) = \sqrt{1-\rho^2} \sigma^Y y H_y^C(t, s, y). \tag{25}$$

and in particular, the problem with no claim gives the integrand $\psi^E \equiv \psi^0$ in the density of the minimal entropy measure:

$$\frac{dQ^E}{dP} = \mathcal{E}(-\widehat{\lambda}^S \cdot \widehat{B}^S - \psi^E \cdot \widehat{Z}^S)_T,$$

where the process ψ^E is given by $\psi_t^E = \psi^E(t, S_t, Y_t)$, $0 \leq t \leq T$, with

$$\psi^E(t, s, y) = \sqrt{1-\rho^2} \sigma^Y y H_y^0(t, s, y). \tag{26}$$

Inserting the optimal feedback control (25) back into the Bellman equation yields the PDE

$$H_t^C + \mathcal{A}_{S,Y}^{Q^M} H^C + \frac{1}{2}(\widehat{\lambda}^S)^2 - \frac{1}{2}(1-\rho^2)(\sigma^Y y H_y^C)^2 = 0, \quad H^C(T,s,y) = -\alpha C(y).$$

Using the fact that H^0 satisfies the same PDE with zero terminal condition, along with the entropic representation (20) of the indifference price, we subtract the PDE for H^0 from that for H^C to yield that the indifference price satisfies

$$p_t + \mathcal{A}_{S,Y}^{Q^M} p + \frac{1}{2}\alpha(1-\rho^2)(\sigma^Y y p_y)^2 - \frac{1}{2}\sqrt{1-\rho^2}\sigma^Y y p_y \psi^E(t,s,y) = 0, \quad p(T,s,y) = C(y),$$

where we have used (26). Using the fact that the drift of Y under Q^E is given by (23) with $\psi = \psi^E$, we recast the above PDE into the form

$$p_t + \mathcal{A}_{S,Y}^{Q^E} p + \frac{1}{2}\alpha(1-\rho^2)(\sigma^Y y p_y)^2 = 0, \quad p(T,s,y) = C(y). \tag{27}$$

For $\alpha = 0$ this PDE becomes linear and the indifference price becomes the the so-called marginal utility-based price p^E, given as an expectation of the payoff under Q^E, as is well-known:

$$p^E(t,s,y) := \lim_{\alpha \to 0} p(t,s,y) = E_{t,s,y}^{Q^E} C(Y_T).$$

Remark 4 (The case $Q^E = Q^M$). In the case when $\mathrm{v}_0^S \le \mathrm{v}_0^Y$, then by (13), $\widehat{\lambda}^S$ loses all dependence on the non-traded asset price, the infimum in the dual problem (22) for $C = 0$ is achieved by $\psi^E = 0$, and the MEMM coincides with the minimal martingale measure, $Q^E = Q^M$. Then H^0 loses dependence on the non-traded asset price and (26) gives $\psi^E = 0$, as it should. The indifference pricing PDE (27) then becomes tractable, as the unknown function $\psi^E(t,s,y)$ in the generator $\mathcal{A}_{S,Y}^{Q^E}$ is set to zero. We shall therefore focus in Section 5 on the case $Q^E = Q^M$ to obtain more explicit results.

3.3 Regularity of the value function

In this section we discuss the required smoothness of the dual value function H^C, and hence of the indifference price, required for the validity of Theorem 1. We outline how smoothness can be established in the case when $\mathrm{v}_0^S \le \mathrm{v}_0^Y$. The arguments are identical (but with modified formulae in places) when $\mathrm{v}^S > \mathrm{v}_0^Y$.

The existence of sufficiently smooth solutions to semi-linear PDEs of the type (24) has been considered by Pham [26] and Benth and Karlsen [6], and similar techniques could in principle be used to establish that H^C is indeed a classical solution to (24). We do not pursue this here, but instead follow Davis [10] and make the transformations

$$\mathcal{S}_t := \frac{1}{\sigma^S} \log S_t, \quad \mathcal{Y}_t := \frac{1}{\sigma^Y} \log Y_t, \quad f(y) := C(\exp(\sigma^Y y)).$$

The function H^C expressed in the new variables is J, defined by the stochastic control problem

$$J(t,s,y) := \inf_{\psi} E^Q \left[\frac{1}{2} \int_t^T \left[\mathcal{L}^2(t,\mathcal{S}_u) + \psi_u^2 \right] du + f(\mathcal{Y}_T) \,\middle|\, \mathcal{S}_t = s, \mathcal{Y}_t = y \right], \tag{28}$$

where the function $\mathcal{L}(t,s)$ is given by transforming the formula for $\widehat{\lambda}^S$ from Proposition 1 into the new variables:

$$\mathcal{L}(t,s) = m_t \left(\frac{\lambda_0^S}{m_0} - \mathcal{S}_0 + \frac{1}{2}\sigma^S t + s \right).$$

The dynamics of the state variables in (28) are

$$
\begin{aligned}
d\mathcal{S}_t &= -\frac{1}{2}\sigma^S dt + d\widehat{B}_t^{S,Q}, \\
d\mathcal{Y}_t &= [a(t,\mathcal{S}_t,\mathcal{Y}_t) - \sqrt{1-\rho^2}\psi_t]dt + d\widehat{B}_t^{Y,Q},
\end{aligned}
$$

where $a(t, s, y)$ is equal to $\widehat{\lambda}_t^Y - \rho \widehat{\lambda}_t^S - \frac{1}{2}\sigma^Y$, transformed to the new variables:

$$a(t, s, y) = b_t \left[\frac{\lambda_0^Y - \rho\lambda_0^S}{b_0} - (\mathcal{Y}_0 - \rho\mathcal{S}_0) + \frac{1}{2}(\sigma^Y - \rho\sigma^S)t + y - \rho s \right] - \frac{1}{2}\sigma^Y,$$

with $b_0 = m_0$ (and hence $b_t = m_t$) in the case where $\mathrm{v}_0^S = \mathrm{v}_0^Y$. Then a is Lipschitz in s, y for all $t \in [0, T]$. The point of making this transformation is that (28) is a standard form of stochastic control problem, whose solution can be shown to be given by a classical solution of the HJB equation

$$J_t + a(t, s, y)J_y - \frac{1}{2}\sigma^S J_s + \frac{1}{2}J_{ss} + \rho J_{sy} + \frac{1}{2}J_{yy} + \frac{1}{2}\mathcal{L}^2(t, s) + \min_\psi \left[\frac{1}{2}\psi^2 - \sqrt{1 - \rho^2}\psi J_y \right] = 0, \quad (29)$$

with $J(T, s, y) = f(y)$. The salient feature of this equation is that it is uniformly elliptic or parabolic[1] for $\rho^2 \leq 1$, and the proof of Lemma 1 and Theorem 2 in Davis [10] can be adapted to show that the value function J is the unique classical solution of (29), and hence that the function $H^C(t, s, y)$ is a classical solution of (24), and thus the primal value function u^C is smooth enough for the proof of Theorem 1 to be valid.

4 Payoff decompositions and asymptotic expansions

4.1 Payoff decompositions and price representation

We shall obtain an asymptotic representation of the indifference price valid for small values of risk aversion, following from payoff decompositions and a price representation equation.

We work under the MEMM Q^E, with asset price dynamics given by (23) with $\psi = \psi^E$. Define the local $(Q^E, \widehat{\mathbb{F}})$-martingale L by

$$L_t := \sqrt{1 - \rho^2}\sigma^Y \int_0^t Y_u p_y(u, S_u, Y_u) d\widehat{Z}_u^{S, Q^E}, \quad 0 \leq t \leq T, \quad (30)$$

where $\widehat{Z}^{S, Q^E}$ is a $(Q^E, \widehat{\mathbb{F}})$-Brownian motion orthogonal to the Brownian motion $\widehat{B}^{S, Q^E}$ driving the stock under Q^E. Note that from (25), (26), the entropic representation (20) of the indifference price and the integrability condition (21), that L is a $(Q^E, \widehat{\mathbb{F}})$-martingale. We have the following decomposition of the claim payoff.

Lemma 1 (Payoff decomposition). *The claim payoff admits the decomposition*

$$C(Y_T) = p(t, S_t, Y_t) + \int_t^T \theta_u^{(H)} dS_u + L_T - L_t - \frac{1}{2}\alpha(\langle L\rangle_T - \langle L\rangle_t), \quad 0 \leq t \leq T, \quad (31)$$

where $\theta^{(H)}$ is the optimal hedging strategy for the claim, given in Theorem 1.

Remark 5. Mania and Schweizer [20] have obtained a similar result in a backward stochastic differential equation representation (see equation (4.5) in Theorem 13 of [20]), in a more abstract model. The new feature of Lemma 1 is that the integrand $\theta^{(H)}$ and the Q^E-martingale L are given in terms of derivatives of the indifference price. Naturally, we have been able to obtain this because we are considering a more explicit model, allowing the Itô formula and PDE methods to be exploited.

Remark 6. Lemma 1 is similar in spirit to Theorem 7 in Musiela and Zariphopoulou [23], but our hedging strategy depends, of course, on derivatives of the indifference price with respect to both state variables (S, Y), rather than the single variable Y of [23].

---

[1]In other words, writing the second order terms as $a_{11}J_{ss} + 2a_{12}J_{sy} + a_{22}J_{yy}$, we have $a_{12}^2 \leq a_{11}a_{22}$.

Proof of Lemma 1. This is established by computing the differential $dp(t, S_t, Y_t)$ under Q^E and using the PDE satisfied by p, to give

$$dp(t, S_t, Y_t) = -\frac{1}{2}\alpha(1-\rho^2)[\sigma^Y Y_t p_y(t, S_t, Y_t)]^2 dt + p_s(t, S_t, Y_t)dS_t + \sigma^Y Y_t p_y(t, S_t, Y_t)d\widehat{B}_t^{Y,Q^E},$$
(32)

where $\widehat{B}^{Y,Q^E}$ is a $(Q^E, \widehat{\mathbb{F}})$-Brownian motion driving the non-traded asset price. We have

$$\widehat{B}^{Y,Q^E} = \rho\widehat{B}^{S,Q^E} + \sqrt{1-\rho^2}\widehat{Z}^{S,Q^E},$$

as well as $dS_t = \sigma^S S_t d\widehat{B}_t^{S,Q^E}$. Using these relations, along with the optimal hedging formula of Theorem 1 and the definition (30) of L, converts (32) to

$$dp(t, S_t, Y_t) = -\frac{1}{2}\alpha d\langle L\rangle_t + \theta_t^{(H)}dS_t + dL_t,$$

and the result follows by integrating from t to T.

$\square$

Define the process L^E as the $\alpha \to 0$ limit of L, given by replacing the indifference price p with the marginal price p^E in the definition (30) of L:

$$L_t^E := \sqrt{1-\rho^2}\sigma^Y \int_0^t Y_u p_y^E(u, S_u, Y_u)d\widehat{Z}_u^{S,Q^E}, \quad 0 \le t \le T.$$
(33)

We have the immediate corollary below, from setting $\alpha = 0$ in Lemma 1, the Föllmer-Schweizer-Sondermann [13, 14] decomposition of the payoff under Q^E in our specific model.

Corollary 1 (Föllmer-Schweizer-Sondermann decomposition). *The claim payoff admits the decomposition*

$$C(Y_T) = p^E(t, S_t, Y_t) + \int_t^T \theta_u^E dS_u + L_T^E - L_t^E, \quad 0 \le t \le T,$$
(34)

where p^E is the marginal utility-based price of the claim, θ^E is the marginal hedging strategy for the claim, given by Theorem 1 with p^E in place of the indifference price, and L^E is the process in (33).

The following corollary of Lemma 1 follows by taking conditional expectations of (31) under Q^E, given $(S_t, Y_t) = (s, y)$.

Corollary 2 (Indifference price representation). *The indifference pricing function $p : [0,T] \times \mathbb{R}_+^2$ has the representation*

$$p(t, s, y) = p^E(t, s, y) + \frac{1}{2}\alpha E_{t,s,y}^{Q^E}[\langle L\rangle_T - \langle L\rangle_t].$$
(35)

Remark 7. A more abstract form of this result appears in Mania and Schweizer [20] (see the last equation before Theorem 18 in their paper). Once again, the new feature here is the explicit identification of the Q^E-martingale L with the derivative of the indifference price according to (30).

4.2 Asymptotic expansions

Denote by

$$v(t, s, y) := \text{var}_{t,s,y}^{Q^E}[C(Y_T)],$$

the conditional variance of the claim payoff under Q^E. Define the gains process G^E associated with the marginal hedging strategy by

$$G_t^E := \int_0^t \theta_u^E dS_u, \quad 0 \le t \le T.$$

The asymptotic expansion for the indifference price to first order in α is then given by the following theorem.

Theorem 2. *The indifference pricing function $p(t, s, y)$ has the asymptotic expansion*

$$p(t, s, y) = p^E(t, s, y) + \frac{1}{2}\alpha \left(v(t, s, y) - E_{t,s,y}^{Q^E}[\langle G^E \rangle_T - \langle G^E \rangle_t] \right) + O(\alpha^2).$$

Proof. Write

$$p(t, s, y) = p^E(t, s, y) + \alpha p^{(1)}(t, s, y) + O(\alpha^2).$$

Insert this expansion into the price representation equation (35), and use the definition (30) of L, to obtain

$$\alpha p^{(1)}(t, s, y) + O(\alpha^2) = \frac{1}{2}\alpha(1-\rho^2)(\sigma^Y)^2 E_{t,s,y}^{Q^E} \int_t^T Y_u^2 \left(p_y^E(u, S_u, Y_u) + \alpha p_y^{(1)}(u, S_u, Y_u) + O(\alpha^2) \right)^2 du.$$

Equating terms of order α and using the definition (33) of L^E, we obtain

$$p^{(1)}(t, s, y) = \frac{1}{2} E_{t,s,y}^{Q^E}[\langle L^E \rangle_T - \langle L^E \rangle_t]. \tag{36}$$

Now use the Föllmer-Schweizer-Sondermann decomposition (34) to compute

$$
\begin{aligned}
v(t, s, y) = \text{var}_{t,s,y}^{Q^E}[C(Y_T)] &= E_{t,s,y}^{Q^E}[(C(Y_T) - p^E(t, s, y))^2] \\
&= E_{t,s,y}^{Q^E} \left[(G_T^E - G_t^E + L_t^E - L_t^E)^2 \right].
\end{aligned}
$$

The processes G^E, L^E are orthogonal Q^E-martingales, so this becomes

$$v(t, s, y) = E_{t,s,y}^{Q^E}[\langle G^E \rangle_T - \langle G^E \rangle_t + \langle L^E \rangle_T - \langle L^E \rangle_t],$$

and inserting this into (36) gives the result. $\qquad\square$

Remark 8. Contrast the expansion in Theorem 2 with the corresponding result in the full information case, which is obtained from a Taylor expansion of a non-linear expectation representation of the indifference price, and which is of the form (see [21], for instance)

$$p^{\text{FI}}(t, y) = p^{E,\text{FI}}(t, y) + \frac{1}{2}\alpha(1 - \rho^2)v^{\text{FI}}(t, y) + O(\alpha^2), \quad \text{(full information expansion)} \tag{37}$$

(where the superscript FI denotes full information). In this case, $Q^E = Q^M$, and the indifference price has no dependence on s. As a result, $\langle L^{E,\text{FI}} \rangle = ((1 - \rho^2)/\rho^2)\langle G^{E,\text{FI}} \rangle$, and the variance of the payoff in the full information case is given by

$$
\begin{aligned}
v^{\text{FI}}(t, y) &:= \text{var}_{t,y}^{Q^E}[C(Y_T)] \\
&= E_{t,y}^{Q^E} \left[\langle L^{E,\text{FI}} \rangle_T - \langle L^{E,\text{FI}} \rangle_t + \langle G^{E,\text{FI}} \rangle_T - \langle G^{E,\text{FI}} \rangle_t \right] \\
&= \frac{1}{\rho^2} E_{t,y}^{Q^E} \left[\langle G^{E,\text{FI}} \rangle_T - \langle G^{E,\text{FI}} \rangle_t \right].
\end{aligned}
$$

The first order term in the price expansion of the theorem reduces to

$$\frac{1}{2}\alpha \left(v^{\text{FI}}(t, y) - E_{t,y}^{Q^E} \left[\langle G^{E,\text{FI}} \rangle_T - \langle G^{E,\text{FI}} \rangle_t \right] \right) = \frac{1}{2}\alpha(1 - \rho^2)v^{\text{FI}}(t, y),$$

in accordance with (37). This expansion can be written in analytic form, so numerical computation for non-zero risk aversion is considerably easier in the completely observable case than in the partial information model.

5 Analytic formulae when $Q^E = Q^M$

From now on we specialise to the case when $\mathrm{v}_0^S \leq \mathrm{v}_0^Y$, so that $\widehat{\lambda}_t^S \equiv \widehat{\lambda}^S(t, S_t)$ does not depend on the non-traded asset price, and $Q^E = Q^M$, as in Remark 4.

For convenience, define the process $\widehat{\nu}$ by

$$\widehat{\nu}_t := \widehat{\lambda}_t^Y - \rho \widehat{\lambda}_t^S, \quad 0 \leq t \leq T,$$

with $\widehat{\nu}_0 = \lambda_0^Y - \rho \lambda_0^S$. The dynamics of Y under Q^M are

$$dY_t = \sigma^Y Y_t (\widehat{\nu}_t dt + d\widehat{B}_t^{Y,Q^M}),$$

with $\widehat{B}^{Y,Q^M}$ a Q^M-Brownian motion.

The dynamics of $\widehat{\nu}_t$ under Q^M follow from transforming (14) from P to Q^M, giving

$$d\widehat{\nu}_t = \sqrt{1 - \rho^2} b_t d\widehat{Z}_t^S,$$

where $\widehat{Z}^S$ is a Q^M-Brownian motion (and also a P-Brownian motion) perpendicular to the Brownian motion $\widehat{B}^{S,Q^M}$ driving the stock, related to $\widehat{B}^{Y,Q^M}$ by

$$\widehat{B}^{Y,Q^M} = \rho \widehat{B}^{S,Q^M} + \sqrt{1 - \rho^2} \widehat{Z}^S. \tag{38}$$

From Proposition 1, when $\mathrm{v}_0^S \leq \mathrm{v}_0^Y$, $\widehat{\nu}$ is a function of the current asset prices, $\widehat{\nu}_t \equiv \widehat{\nu}(t, S_t, Y_t)$, given by

$$\widehat{\nu}(t, s, y) = \frac{\widehat{\nu}_0 + b_0(\xi^Y(t, y) - \rho \xi^S(t, s))}{1 + b_0 t}, \tag{39}$$

with $b_0 = m_0$ in the case that $\mathrm{v}_0^S = \mathrm{v}_0^S$, and where $\xi^S(t, s)$ and $\xi^Y(t, y)$ are the functions in (1) evaluated at $S_t = s$ and $Y_t = y$.

The following result shows that $\log Y_T$ is normal under Q^M, with a mean which depends on the current asset prices, and a with a time-dependent variance.

Theorem 3. *Suppose* $\mathrm{v}_0^S \leq \mathrm{v}_0^Y$. *Under* Q^M, *conditional on* $S_t = s, Y_t = y$,

$$\log Y_T \sim \mathrm{N}\left(\mu(t, s, y), \Sigma^2(t)\right),$$

with

$$
\begin{aligned}
\mu(t, s, y) &= \log y + \sigma^Y \widehat{\nu}(t, s, y)(T - t) - \frac{1}{2}\left(\sigma^Y\right)^2 (T - t), \\
\Sigma^2(t) &= \left[1 + (1 - \rho^2) b_t (T - t)\right] (\sigma^Y)^2 (T - t),
\end{aligned}
\tag{40}
$$

where $\widehat{\nu}(t, s, y)$ *is given in (39) and* b_t *is given in (9). (When* $\mathrm{v}_0^S = \mathrm{v}_0^Y$ *the same formulae hold with* $b_t = m_t$, *for* $0 \leq t \leq T$.)

Proof. We use the SDEs for Y and $\widehat{\nu}$ under Q^M. Applying the Itô formula to $\log Y$ under Q^M, we obtain

$$\log Y_T = \log Y_t + \sigma^Y \int_t^T \widehat{\nu}_u du - \frac{1}{2}(\sigma^Y)^2 (T - t) + \sigma^Y \int_t^T d\widehat{B}_u^{Y,Q^M}, \quad 0 \leq t \leq T. \tag{41}$$

From the dynamics of $\widehat{\nu}_t$ under Q^M we have

$$\widehat{\nu}_u = \widehat{\nu}_t + \sqrt{1 - \rho^2} \int_t^u b_r d\widehat{Z}_r^S, \quad t \leq u \leq T.$$

Hence, after changing the order of integration in a double integral, we have

$$\int_t^T \widehat{\nu}_u du = \widehat{\nu}_t (T - t) + \sqrt{1 - \rho^2} \int_t^T b_u (T - u) d\widehat{Z}_u^S.$$

Inserting this into (41) and using (38) gives

$$
\begin{aligned}
\log Y_T \;=\;&\; \log Y_t + \sigma^Y \left[\widehat{\nu}_t (T-t) + \sqrt{1-\rho^2} \int_t^T (1+b_u(T-u))\, d\widehat{Z}_u^S + \rho \int_t^T d\widehat{B}_u^{S,Q^M} \right] \\
&\; - \frac{1}{2}(\sigma^Y)^2 (T-t), \quad 0 \le t \le T.
\end{aligned}
$$

The stochastic integrals are orthogonal Gaussian processes with zero conditional expectation given $(S_t, Y_t) = (s, y)$, from which the distribution of $\log Y_T$ follows.

$\square$

Using this distribution we obtain BS-style formulae for the marginal price with a dividend rate that depends on the asset prices, and with a time-dependent volatility. We also get an analytic formula for the conditional variance $v(t, s, y)$ in Theorem 2.

Denote by μ_j the j^{th} moment of the payoff under Q^M:

$$
\mu_j(t, s, y) := E_{t,s,y}^{Q^M} \left[C^j(Y_T) \right], \quad t \le T, \quad j \in \mathbb{N}.
$$

The marginal pricing function $p^E(t, s, y) = E_{t,s,y}^{Q^M} C(Y_T)$ is given by

$$
p^E(t, s, y) = \mu_1(t, s, y),
$$

and the variance $v(t, s, y) := \text{var}_{t,s,y}^{Q^M}[C(Y_T)]$ is given by

$$
v(t, s, y) = \mu_2(t, s, y) - \mu_1^2(t, s, y). \tag{42}
$$

For a put option of strike K, the first and second moments are given by the following formulae.

Lemma 2. *With $\mu(t, s, y)$ and $\Sigma(t)$ as in Theorem 3, define $q \equiv q(t, s, y)$ by*

$$
\mu(t, s, y) = \log y - q(t, s, y) - \frac{1}{2}\Sigma^2(t).
$$

The marginal price at time $t \in [0, T]$ of a put option with payoff $(K - Y_T)^+$ is $p^E(t, S_t, Y_t)$, given by

$$
\begin{aligned}
p^E(t, s, y) \;=\;&\; K\Phi(-d_1(t, s, y) + \Sigma(t)) - y \exp(-q(t, s, y))\Phi(-d_1(t, s, y)), \\
d_1(t, s, y) \;:=\;&\; \frac{1}{\Sigma(t)} \left[\log\left(\frac{y}{K}\right) - q(t, s, y) + \frac{1}{2}\Sigma^2(t) \right],
\end{aligned}
$$

where $\Phi(\cdot)$ denotes the standard cumulative normal distribution function.

The second moment of the payoff under Q^M is given by

$$
\begin{aligned}
\mu_2(t, s, y) \;=\;&\; K^2 \Phi(-d_1(t, s, y) + \Sigma(t)) - 2Ky \exp(-q(t, s, y))\Phi(-d_1(t, s, y)) \\
&\; + y^2 \exp(\Sigma^2(t) - 2q(t, s, y))\Phi(-d_1(t, s, y) - \Sigma(t)).
\end{aligned}
$$

To implement the optimal hedging strategy we need the derivatives with respect to s and y of the indifference price. Approximating the indifference price by the asymptotic expansion of Theorem 2, we obtain the derivatives of p^E, and of the variance v. The calculations are similar (but not identical to) those in [21, 22], but some nice cancellations occur to give formulae with a similar flavour to those in [21, 22].

Lemma 3. *The marginal hedging strategy for a short put position is*

$$
\theta^E(t, s, y) = -\rho \frac{\sigma^Y}{\sigma^S} \frac{y}{s} \exp(-q(t, s, y))\Phi(-d_1(t, s, y)).
$$

The derivatives of the variance in (42) satisfy

$$
v_s(t,s,y) + \rho \frac{\sigma^Y}{\sigma^S} \frac{y}{s} v_y(t,s,y)
$$

$$
= -2\rho \frac{\sigma^Y}{\sigma^S} \frac{y}{s} e^{-q(t,s,y)} \left[K\Phi(-d_1(t,s,y)) - y \exp(\Sigma^2(t) - q(t,s,y))\Phi(-d_1(t,s,y) - \Sigma(t)) \right]
$$

$$
- 2p^E(t,s,y)\theta^E(t,s,y).
$$

The final term we need to implement the optimal hedge with the asymptotic expansion of Theorem 2 is the quadratic variation term involving $\langle G^E \rangle$. Using the explicit formula for the marginal price, this term is given as follows.

Lemma 4. *The expectation under Q^M of the quadratic variation $\langle G^E \rangle$ for a put option satisfies*

$$
E_{t,s,y}^{Q^M}[\langle G^E \rangle_T - \langle G^E \rangle_t] = \left(\rho\sigma^Y\right)^2 E_{t,s,y}^{Q^M} \int_t^T \left[Y_u \exp(-q(u, S_u, Y_u))\Phi(-d_1(u, S_u, Y_u))\right]^2 du.
$$

This expression will be evaluated numerically by Monte-Carlo simulation, as will its derivatives with respect to s and y.

6 Numerical results and experiments

We conducted extensive numerical investigations into the partial information valuation and hedging scheme.

An initial gauge of the effect of parameter uncertainty on the valuation and hedging of the claim is given in Table 2, showing put option prices and hedging strategies at time zero from the partial information model and the full information model, for varying correlation and for risk aversions $\alpha = 0$ and $\alpha = 0.01$. We set the means of the prior equal to the true risk premia. These and other parameters are shown in Table 1. The option prices are higher in the partial information case as the agent incorporates the risk from parameter uncertainty into a higher effective volatility, given by (40). This effect does not translate to the hedges, in that the agent generally takes a bigger short position with full information, except for high correlation and non-zero risk aversion. This is a first indication that the biggest effect of parameter uncertainty on utility-based methods is incorporated into the valuation of the claim, as opposed to the hedge. This will be borne out by simulation results which follow, and is not altogether surprising. Utility-based valuation insists on achieving unchanged utility only when selling the option at the indifference price, and not necessarily by trading judiciously with an option premium that is at odds with the agent's objective.

Note the dependence of the results on the Q^M-drift of Y, given by $\sigma^Y \widehat{\nu}$ in the partial information case, and by the corresponding quantity with the true risk premia in the full information case. The negative of this drift acts as a dividend yield in the option pricing formulae, the put option premia increase as this drift decreases, and the hedge ratios become larger in absolute value. This is reflected in the lower half of Table 2.

We then conducted simulation experiments on the hedging of the claim. Using a given prior distribution λ_0^i, v_0^i, $i = S, Y$, either chosen or estimated randomly from simulated data over $[-t_i, 0]$, $i = S, Y$ (and setting λ_0^i to the point estimates of the risk premia, with $v_0^i = 1/t_i$), we hedged a short position in a put option of strike K over $[0, T]$, using some fixed rebalancing interval δt (we used $\delta t = 1/252$, one trading day in all results we report). The prior distribution was updated over $[0, T]$ and the claim was hedged using the optimal strategy of Theorem 1 with the indifference price approximated by the asymptotic expansion of Theorem 2, or by its $\alpha \to 0$ limit, the marginal price. We generated a terminal hedging error over the simulated paths and repeated this over many price paths to produce a hedging error distribution, and compared this distribution with that produced from alternative strategies, namely:

- The utility-based hedge which does not incorporate the learning from filtering. This uses the indifference hedging formulae of the completely observable incomplete model with

16

Table 1: Partial and full information (FI) put option prices and hedge ratios at time zero, for risk aversion $\alpha = 0$ (indexed with superscript E) and $\alpha = 0.01$ from the first order in α expansions (indexed with superscript 1). The parameters in the upper half of the table are as in Table 1, and in the lower half of the Table we have $\lambda^Y = \lambda_0^Y = 0.325$. The BS price is 9.95 and the BS hedge is -0.3752.

ρ	p^E	p^1	$p^{E,\mathrm{FI}}$	$p^{1,\mathrm{FI}}$	θ^E	θ^1	$\theta^{E,\mathrm{FI}}$	$\theta^{1,\mathrm{FI}}$
0.25	11.50	12.88	6.68	7.21	-0.0704	-0.0779	-0.0737	-0.0791
0.5	11.64	12.91	7.60	8.08	-0.1493	-0.1653	-0.1601	-0.1689
0.75	11.19	12.15	8.60	8.91	-0.2426	-0.2701	-0.2588	-0.2667
0.85	10.76	11.53	9.02	9.22	-0.2870	-0.3213	-0.3016	-0.3073
0.95	10.11	10.63	9.45	9.52	-0.3389	-0.3823	-0.3463	-0.3485
0.25	12.58	14.08	7.84	8.45	-0.0743	-0.0817	-0.0816	-0.0872
0.5	13.03	14.42	9.12	9.68	-0.1593	-0.1752	-0.1786	-0.1877
0.75	12.96	14.04	10.52	10.88	-0.2633	-0.2905	-0.2902	-0.2980
0.85	12.72	13.60	11.11	11.35	-0.3147	-0.3483	-0.3385	-0.3441
0.95	12.32	12.91	11.71	11.80	-0.3769	-0.4193	-0.3889	-0.3910

Table 2: Parameters for prices and hedge ratios in Table 2

S_0	Y_0	K	T	λ^S	σ^S	λ^Y	σ^Y	λ_0^S	λ_0^Y	v_0^S	v_0^Y
100	100	100	1 year	0.5	0.3	0.425	0.25	0.5	0.425	2	2

the asset drifts set to their initial values λ_0^S, λ_0^Y, and kept fixed throughout the hedging time-frame. To order α the indifference price at $t \in [0, T]$ is given by (37), where we take the drift of Y under Q^M to be $\sigma^Y(\lambda_0^Y - \rho\lambda_0^S)$. The optimal hedging strategy is given by Theorem 1 (with, of course, no derivative with respect to s).

- The BS-style hedge which assumes that S is a perfect proxy for Y, given by

$$\theta_t^{\mathrm{BS}} = \frac{\sigma^Y}{\sigma^S}\frac{Y_t}{S_t}\frac{\partial}{\partial y}\mathrm{BS}(t, Y_t; \sigma^Y), \quad 0 \le t \le T. \tag{43}$$

  where $\mathrm{BS}(t, y; \sigma^Y)$ denotes the BS formula with underlying asset price y and volatility σ^Y.

- A variant of the BS hedge proposed by Hulley and McWalter [16], which multiplies the hedge in (43) by the correlation ρ. In effect, this approximates the Q^M-drift of Y by zero.

We also carried out the tests using the full information hedge, with the drift of Y under Q^M taken to be its true value $\sigma^Y(\lambda^Y - \rho\lambda^S) =: \sigma^Y\nu$. This strategy is not available in reality to the agent, but we include it as a calibration. It is not guaranteed to produce the best results, as any finite sample of data may not reflect the true drifts of the assets.

The initial wealth was set to the time zero BS price for all the strategies. We also computed the hedging error when starting with the appropriate option premium corresponding to the hedging program, to assess the effect of valuation as well as the benefits or otherwise of the hedging strategy.

First, we present results in which we used the marginal hedging strategies (with or without filtering) for the utility-based methods. In this case analytic formulae were available for all quantities involved, and with the ensuing fast computation we carried out extensive simulations over a range of scenarios, a representative sample of which we report.

Table 3 shows summary statistics for the hedging error distributions when the prior risk premia were set equal to the true risk premia, $\lambda_0^i = \lambda^i$, for $i = S, Y$. Of course, in this case, the

full information results will be identical to those without filtering. We use this as a base case and then vary the prior to illustrate how the benefits or otherwise of the partial information approach vary with the prior.

The hedge with filtering gives a higher mean, median and expected utility than the other strategies. The standard deviation is higher than that of the unfiltered hedge and the correlation-weighted BS hedge. The BS hedge is the worst performer, and is massively improved by weighting it by the correlation. With the parameters used in the upper half of the table, the true drift of Y under the minimal measure is positive (we have $\nu := \lambda^Y - \rho\lambda^S = 0.05$). Under these conditions the utility-based hedges tend to generate a positive terminal wealth by under-hedging (that is, taking a small short position), as the option tends to end up out of the money. The correlation-weighted BS hedge in effect approximates the Q^M-drift of Y drift by zero, so lower than the true value. It then tends to over-hedge (that is, take a larger a short position than the utility-based hedges) and thus under-perform. Conversely, if we change the risk premium of Y to $\lambda^Y = 0.325$, then the Q^M-drift of Y becomes negative, and we obtain the results in the lower half of Table 3 (over a fresh set of simulated paths). Here, the correlation-weighted BS hedge outperforms the utility-based hedge without learning, but once again the filtering has improved the performance so that it turned out to be the best strategy in terms of mean, median and expected utility.

If we incorporate the effect of valuation, and start the hedging programs at the indifference price implied by the hedging method, then the filtering procedure massively outperforms the other methods. We found this to be true in virtually all our simulation results.

These initial results indicate that the filtering procedure can improve the performance of the utility-based method without filtering, regardless of initial wealth, and regardless of the fact the the unfiltered hedge used the true values of the risk premia. The latter point reflects the fact that asset data over any finite time period may well not reflect the true values of the drifts (the notorious difficulty of drift parameter uncertainty) and the updating of filtering has counteracted this to some extent.

The quality of strategies that do not use filtering appear to be related to the perceived Q^M-drift of Y relative to its true value. This is indicated by the results in Table 4. In the upper half of the table, the agent who does not incorporate filtering perceives the Q^M-drift of Y as greater than the true value, though both are positive. The non-filtered strategy under-hedges, and when the option is unlikely to be exercised, this is a successful policy (when all hedges begin with common initial wealth). The improvement in hedging offered by filtering is now less pronounced (the mean hedging error with filtering is larger than without, but the median is not). A caveat to this is that if agents use their respective option premia as the initial wealth, then the agent who does not incorporate learning will suffer. Similarly, the correlation-weighted BS hedge approximates the Q^M-drift of Y by zero, and does not perform so well in this case. Overall, the filtering procedure appears to be of benefit. This is supported by the results in the lower half of the table. Now the agent who does not filter perceives the Q^M-drift of Y as negative, and the correlation-weighted BS hedge outperforms the utility-based hedge without filtering, but if we incorporate filtering, then the utility-based method has an improved performance.

Occasionally, we found that the filtering procedure was not beneficial, usually when the prior gave an extremely poor estimate of $\nu := \lambda^Y - \rho\lambda^S$. The filtering then appears to be of limited use in improving the hedge, and relies on its increased valuation of the claim to be at all effective.

To this end, the upper half of Table 5 shows results in the case that the risk premium of Y is initially badly over-estimated and the risk premium of S is badly under-estimated. In this scenario, $\nu = -0.375$, while the initial value of this quantity in the prior is $\nu_0 = 0.425$. The non-filtering agent perceives the put as much less risky than in reality, and his hedging strategy produces losses. But the prior is so poor that the filtering procedure fails to update it drastically enough to change this perception, and the filtered hedge is also poor. The filtering method performs better than the non-filtered hedge if the agents incorporate their initial valuations of the option into the initial wealth, but in this case both strategies do worse than the correlation-weighted BS hedge. If we improve the prior a little, we get the results shown in the lower half

18

Table 3: Hedging error statistics as a fraction of the initial wealth (the BS price at time zero), over $40,000$ asset price paths, with the risk premia of the prior set to be equal to the true risk premia. The utility-based hedges used the marginal ($\alpha = 0$) hedging strategies. Parameters in the upper half of the table are as in Table 1, $\delta t = 1/252$, and $\rho = 0.75$. The time zero option premia are $p_0^E = 11.93$, $p_0^{E,\mathrm{NF}} = p^{E,\mathrm{FI}} = 9.39$ (NF denoting no filtering), $p_0^{\mathrm{BS}} = 9.95$. Figures in parentheses show results obtained if the option premium corresponding to the hedging method was used as the initial wealth (so standard deviations and the BS-related statistics are unchanged), again as a percentage of the time zero BS price. The last column of the table is the expected utility of the terminal wealth with random endowment. In the lower half of the Table we set $\lambda_0^Y = \lambda^Y = 0.325$, with all other parameters as before. In this case the initial utility-based option premia are $p_0^E = 12.96$, $p_0^{E,\mathrm{NF}} = p^{E,\mathrm{FI}} = 10.52$.

	Mean	SD	Median	$EU(X_T - C(Y_T))$
Filtered hedge	0.0647 (0.2637)	0.7625	0.2347 (0.4337)	-0.9965 (-0.9770)
Non-filtered hedge	0.0551 (-0.0005)	0.7449	0.2211 (0.1654)	-0.9973 (-1.0029)
BS hedge	-0.0753	0.8159	0.0226	-1.0109
ρ.BS hedge	0.0480	0.7444	0.2044	-0.9980
Full information hedge	0.0551 (-0.0005)	0.7449	0.2211 (0.1654)	-0.9973 (-1.0029)
With $\lambda^Y = \lambda_0^Y = 0.325$				
Filtered hedge	-0.0397 (0.2628)	0.8066	0.1291 (0.4316)	-1.0073 (-0.9774)
Non-filtered hedge	-0.0548 (0.0026)	0.7907	0.1129 (0.1703)	-0.9973 (-1.0029)
BS hedge	-0.1731	0.8572	-0.0550	-1.0211
ρ.BS hedge	-0.0480	0.7916	0.1288	-1.0080
Full information hedge	-0.0548 (0.0026)	0.7907	0.1129 (0.1703)	-0.9973 (-1.0029)

Table 4: Hedging error statistics with $\lambda_0^Y = 0.525$ (upper table) and $\lambda_0^Y = 0.325$ (lower table). The utility-based hedges used $\alpha = 0$, and all other parameters as in Tables 1 and 3. The time zero option premia are for the upper table are $p_0^E = 10.94$, $p_0^{E,\mathrm{NF}} = 8.34$, $p^{E,\mathrm{FI}} = 9.39$, with the BS price and the full information price as in Table 3. For the lower table we have $p_0^E = 12.96$, $p_0^{E,\mathrm{NF}} = 10.52$.

With $\lambda_0^Y = 0.525$	Mean	SD	Median	$EU(X_T - C(Y_T))$
Filtered hedge	0.0650 (0.1652)	0.7726	0.2474 (0.3476)	-0.9966 (-0.9867)
Non-filtered hedge	0.0628 (-0.0988)	0.7566	0.2510 (0.0895)	-0.9967 (-1.0128)
BS hedge	-0.0770	0.8253	0.0240	-1.0112
ρ.BS hedge	0.0416	0.7543	0.2029	-0.9988
Full information hedge	0.0486 (-0.0071)	0.7547	0.2189 (0.1632)	-0.9981 (-1.0036)
With $\lambda_0^Y = 0.325$				
Filtered hedge	0.0568 (0.3593)	0.7684	0.2198 (0.5224)	-0.9974 (-0.9678)
Non-filtered hedge	0.0391 (0.0966)	0.7513	0.1882 (0.2456)	-0.9990 (-0.9933)
BS hedge	-0.0784	0.8255	0.0237	-1.0113
ρ.BS hedge	0.0462	0.7513	0.2047	-0.9983
Full information hedge	0.0534 (-0.0023)	0.7516	0.2201 (0.1645)	-0.9976 (-1.0031)

Table 5: Hedging error statistics with $\lambda^Y = 0$ and $\lambda^S_0 = 0$ (upper table). Other parameters are as in Tables 1 and 3. The time zero option premia are $p^E_0 = 8.50$, $p^{E,\text{NF}}_0 = 5.83$, $p^{E,\text{FI}} = 14.63$, with the BS price as in Table 3. In the lower table the prior is improved: $\lambda^S_0 = 0.4$, $\lambda^Y_0 = 0.1$, and the time zero option premia are $p^E_0 = 14.59$, $p^{E,\text{NF}}_0 = 12.34$, with other prices unchanged.

$\lambda^S_0 = 0$, $\lambda^Y_0 = 0.425$	Mean	SD	Median	$EU(X_T - C(Y_T))$
Filtered hedge	-0.3935 (-0.5393)	0.9583	-0.1996 (-0.3454)	-1.0448 (-1.0600)
Non-filtered hedge	-0.3670 (-0.7810)	0.9672	-0.1068 (-0.5208)	-1.0421 (-1.0859)
BS hedge	-0.5645	0.9820	-0.4048	-1.0629
ρ.BS hedge	-0.4251	0.9310	-0.2105	-1.0478
Full information hedge	-0.4710 (0.0001)	0.9157	-0.3069 (0.1641)	-1.0524 (-1.0042)

$\lambda^S_0 = 0.4$, $\lambda^Y_0 = 0.1$	Mean	SD	Median	$EU(X_T - C(Y_T))$
Filtered hedge	-0.4231 (0.0343)	0.9364	-0.2560 (0.2104)	-1.0486 (-1.0010)
Non-filtered hedge	-0.4503 (-0.2102)	0.9211	-0.2585 (-0.0184)	-1.0503 (-1.0255)
BS hedge	-0.5654	0.9815	-0.4036	-1.0630
ρ.BS hedge	-0.4247	0.9312	-0.2068	-1.0477
Full information hedge	-0.4709 (0.0001)	0.9158	-0.3031 (0.1680)	-1.0630 (-1.0477)

of the table. Now the filtered hedge improves the unfiltered hedge, but still not to the point of out-performing the correlation-weighted BS hedge, unless the initial option premium is used as the initial wealth.

The conclusion to be drawn from these tests is that filtering can be of benefit, provided the prior is not extremely poor. If a fortuitous combination of parameters aids other methods, then the partial information approach needs to incorporate its valuation of the claim into the initial wealth to be effective.

To confirm these ideas, we conducted some modified experiments. Table 6 shows results of two simulation experiments in which the prior was randomly generated using data before time zero. These results are in line with our earlier conclusions. In both cases the filtering has improved the hedging performance of the unfiltered hedge.

We conducted a modified simulation experiment in which a different prior was generated randomly on each simulation run. The results are shown in Table 7. The non-filtered hedge then out-performed the filtered hedge. With multiple priors, the quality of the prior averaged out over many estimations to be perfectly adequate, so that the unfiltered hedge performed well. When the initial option premium allied to the hedging method is used as the initial wealth, the filtering method is more effective.

The main insurance provided against parameter uncertainty by utility-based methods appears to be incorporated into the initial valuation of the claim. The intuition here is that the rate of learning provided by the filter on the asset drifts is simply not great enough (especially with an unfortunate prior) to counteract parameter uncertainty without the extra insurance of an increased option premium.

Very similar results were obtained when using non-zero risk aversion. As an example, we show in Table 8 the hedging error statistics obtained using 5,000 paths and one randomly generated prior. In this case the prior gives $\nu_0 = -1.3267$, against a true value of $\nu = 0.05$. The unfiltered hedge is improved by the filtering, but both are out-performed by the correlation-weighted BS hedge. But if the option premia are incorporated into the initial wealth, the partial information hedge is the best performer.

The overall conclusion to be drawn from these results is that the filtering procedure can sometimes help in hedging, but there are instances when it is of limited, if any use, unless accompanied by the safety net of the increased option premium that it would charge, taking into account the uncertainty in parameter estimation. If one considers the combined valuation and hedging program together, then incorporating parameter uncertainty and using a filtering

Table 6: Hedging error statistics as a fraction of the initial wealth (the time zero BS price), over $40,000$ asset price paths. For the utility-based hedges, we used the marginal hedges. The true risk premia and volatilities were as in Table 1 (so $\nu = 0.05$) as was the maturity and option strike, and the correlation was $\rho = 0.75$. A single prior distribution was estimated using data before time zero, with $S_{-t_S} = Y_{-t_Y} = 90$, $t_S = t_Y = 0.5$, so the prior variances are as in Table 1. The risk premia of the prior were estimated from this data. In the upper table this gave $\lambda_0^S = 0.8289$, $\lambda_0^Y = 0.1173$, and the initial asset prices were $S_0 = 99.65$, $Y_0 = 89.91$. The time zero option premia were $p_0^E = 23.72$, $p_0^{E,NF} = 22.84$, $p_0^{E,FI} = 14.64$, $p_0^{BS} = 15.33$. In the lower table we have $\lambda_0^S = -0.9736$, $\lambda_0^Y = -0.1654$, $S_0 = 76.04$, $Y_0 = 86.79$, $p_0^E = 12.45$, $p_0^{E,NF} = 9.68$, $p_0^{E,FI} = 16.61$, $p_0^{BS} = 17.34$.

With $\nu_0 = -0.5044$	Mean	SD	Median	$EU(X_T - C(Y_T))$
Filtered hedge	0.0455 (0.5933)	0.5962	0.1295 (0.6772)	-0.9973 (-0.9170)
Non-filtered hedge	0.0104 (0.5011)	0.5901	0.0718 (0.5624)	-1.0025 (-0.9219)
BS hedge	-0.0621	0.6437	-0.0052	-1.0145
ρ.BS hedge	0.0404	0.5902	0.1382	-0.9980
Full information hedge	0.0440 (-0.0009)	0.5908	0.1471 (0.1022)	-0.9974 (-1.0043)
With $\nu_0 = 0.5648$				
Filtered hedge	0.0767 (-0.2049)	0.5674	0.1621 (-0.1195)	-0.9917 (-1.0413)
Non-filtered hedge	0.0749 (-0.3501)	0.5632	0.1853 (-0.2398)	-0.9919 (-1.0678)
BS hedge	-0.0601	0.5955	-0.0181	-1.0159
ρ.BS hedge	0.0370	0.5472	0.1139	-0.9982
Full information hedge	0.0399 (-0.0018)	0.5479	0.1199 (0.0782)	-0.9977 (-1.0049)

Table 7: Hedging error statistics as a fraction of initial wealth (the BS price at time zero) over $40,000$ simulations. We repeatedly estimated the prior using data before time zero, with $t_S = t_Y = 0.5$, $T = 1$, $\delta t = 1/252$, $S_{-t_S} = Y_{-t_Y} = 90$, $\lambda^S = 0.5$, $\sigma^S = 0.3$, $\lambda^Y = 0.425$, $\sigma^Y = 0.25$, $\rho = 0.75$, $K = 100$, $\alpha = 0$. The average initial wealth was $\bar{X}_0 = 14.57$. The average initial asset prices at time zero were $\bar{S}_0 = 97.15$, $\bar{Y}_0 = 94.93$. The average time zero risk premia were $\bar{\lambda}_0^S = 0.5079$, $\bar{\lambda}_0^Y = 0.4266$. The average initial prices were $\bar{p}_0^E = 18.41$, $\bar{p}_0^{E,NF} = 17.24$, $\bar{p}_0^{E,FI} = 13.97$, $\bar{p}_0^{BS} = \bar{X}_0$.

	Mean	SD	Median	$EU(X_T - C(Y_T))$
Filtered hedge	0.0835 (0.2230)	0.7907 (0.9659)	0.1973 (0.2857)	-0.9943 (-0.9601)
Non-filtered hedge	0.1054 (0.0892)	0.7987 (0.9706)	0.2183 (0.1101)	-0.9945 (-0.9721)
BS hedge	-0.0649	0.8305	0.0033	-1.0120
ρ.BS hedge	0.0480	0.7504	0.1532	-0.9975
Full information hedge	0.0547 (0.0041)	0.7498 (0.7496)	0.1668 (0.1188)	-0.9970 (-1.0030)

Table 8: Hedging error statistics as a fraction of initial wealth (the BS price at time zero) over $5,000$ simulations, using the $O(\alpha)$ utility-based hedges. A single prior was randomly generated using data over $[-t_i, 0]$, $i = S, Y$, with $t_S = t_Y = 0.5$, $T = 1$, $S_{-t_S} = Y_{-t_Y} = 90$, $\lambda^S = 0.5$, $\sigma^S = 0.3$, $\lambda^Y = 0.4$, $\sigma^Y = 0.25$, $\rho = 0.75$, $K = 100$, $\alpha = 0.01$. The initial wealth was the time zero BS price, $X_0 = p_0^{\mathrm{BS}} = 13.58$. The initial asset prices at time zero were $S_0 = 123.67$, $Y_0 = 92.86$. The time zero risk premia were $\lambda_0^S = 0.2.2688$, $\lambda_0^Y = 0.3749$. The initial option prices were $p_0 = 35.05$, $p_0^{\mathrm{NF}} = 34.34$, $p_0^{\mathrm{FI}} = 13.68$.

	Mean	SD	Median	$EU(X_T - C(Y_T))$
Filtered hedge	-0.0069 (1.5745)	0.6504	0.0602 (1.6415)	-1.0049 (-0.8107)
Non-filtered hedge	-0.0473 (1.4814)	0.6636	-0.0225 (1.5062)	-1.0106 (-0.8211)
BS hedge	-0.0922	0.7073	-0.0142	-1.0174
ρ.BS hedge	0.0207	0.6466	0.1420	-1.0011
Full information hedge	0.0190 (0.0267)	0.6464	0.1381 (0.1457)	-1.0014 (-1.0003)

Table 9: Stocks comprising the non-traded basket.

Abbey National	British Airports Authority	BAE Systems	British Gas	
Boots PLC	British Telecom	Shell	Tesco	Vodafone

approach is of benefit.

7 Empirical examples

In this section we illustrate how the hedging approach we have developed can be applied in an empirical situation. We tackle the case of hedging a basket of nine UK stocks using futures contracts on the FTSE100 index. We do not claim to be carrying out an exhaustive empirical testing procedure, which is a good topic for future research.

We obtained daily (closing price) data from January 1, 1990 to 30 August, 2003, on the closest to maturity futures contract on the FTSE100 index, and on nine stocks (listed in Table 9) used to construct an equally weighted basket. All data were obtained from Datastream.

We consider a put option on the basket. We estimate the risk premia and volatilities of the futures price and the basket, and the correlation, from the returns of a selected time period ending at the inception date of the option. We use the estimated risk premia in the prior, and either keep them fixed and use the full information model formulae, or update them via filtering and use the partial information formulae.

The first example is an at the money put option written on October 1, 1996 and maturing March 31, 1998. Data from April 1, 1996 up to the inception date of the option were used to estimate the volatilities, correlation and the prior risk premia, and as usual we set $\mathrm{v}_0^i = 1/t_i$, $i = S, Y$, where t_i is the time between the 1 April and 1 October. The relevant parameters are given in Table 10, along with the risk aversion α. We used a zero interest rate. We valued and hedged the option assuming daily portfolio re-balancing.

Figure 1 shows the futures price (scaled down by a factor of 10) and the basket price paths

Table 10: Empirical parameters used to value and hedge a put option on a basket of stocks from October 1996 to March 1998. The hedge is illustrated in Figure 1.

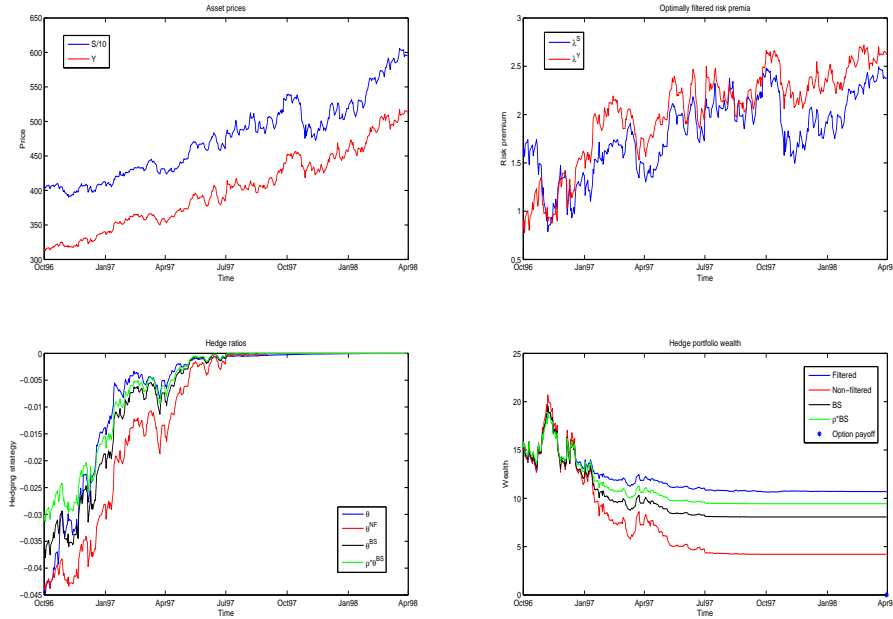| S_0	Y_0	K	λ_0^S	σ^S	λ_0^Y	σ^Y	$\mathrm{v}_0^S = \mathrm{v}_0^Y$	ρ	α
4028	313.1	313.1	1.5667	0.1009	0.7980	0.1039	1.9944	0.8233	0.01

Figure 1: Asset price paths, filtered risk premia and hedge portfolios from October 1996 to March 1998. The parameters are as in Table 10. The estimates of volatilities, correlation and the prior risk premia were obtained from six months of data before the inception of the option. The average values of the filtered risk premia over the hedging time-frame were $\bar{\lambda}^S = 1.8105$, $\bar{\lambda}^Y = 2.0241$. In the figure, all hedges begin with an initial wealth equal to the time zero BS price, $p^{\mathrm{BS}} = 15.86$. The utility-based option premia are $p_0 = 33.42$ (partial information) and $p_0^{\mathrm{NF}} = 30.67$ (without filtering).

over the option life, along with the filtered risk premia, the hedge ratios and hedge portfolios over these paths. One can see the filtered hedge outperforming the others in the graph on the bottom right.

When all the strategies began with an initial wealth equal to the time zero BS price the terminal hedging errors as a fraction of this initial wealth are 0.6747 for the optimal hedge with filtering, 0.2654 without filtering, -0.038 for the BS hedge and -0.0313 for the correlation-modulated BS hedge. So in this case the filtering was of great benefit. This can be traced to the fact that the prior risk premia give $\widehat{\nu}_0 = -0.4919$, while the average filtered risk premia over the option life give $\bar{\widehat{\nu}} = 0.5335$. In other words, the unfiltered hedge and the modulated BS hedge would tend to hedge too aggressively, and this is counteracted by the learning procedure of the filtering. If one uses the appropriate initial option premia as the initial wealth, the partial information outperforms the others even more. So, in this example, the parameters of the prior and the subsequent realisations of the risk premia meant that the filtering method was of benefit without recourse to the extra insurance provided by the increased partial information option premium.

A second example shows how if the parameters of the prior are fortuitously set for the unfiltered hedge, it can apparently do well, but this is also loaded with danger, as the option premium such an agent would use would lead to losses.

In Figure 2 we hedge an at the money put option written on 1 October 1997 and maturing on 31 March 1999. Data from 1 April 1997 up to the option inception were used for estimating the prior and other parameters, shown in Table 11. In this case the terminal hedging errors as a fraction of the (common) initial wealth are 0.5087 for the filtered hedge, 0.6603 for the unfiltered

Table 11: Empirical parameters used to value and hedge a put option on a basket of stocks from October 1997 to March 1999. The hedge is illustrated in Figure 2.

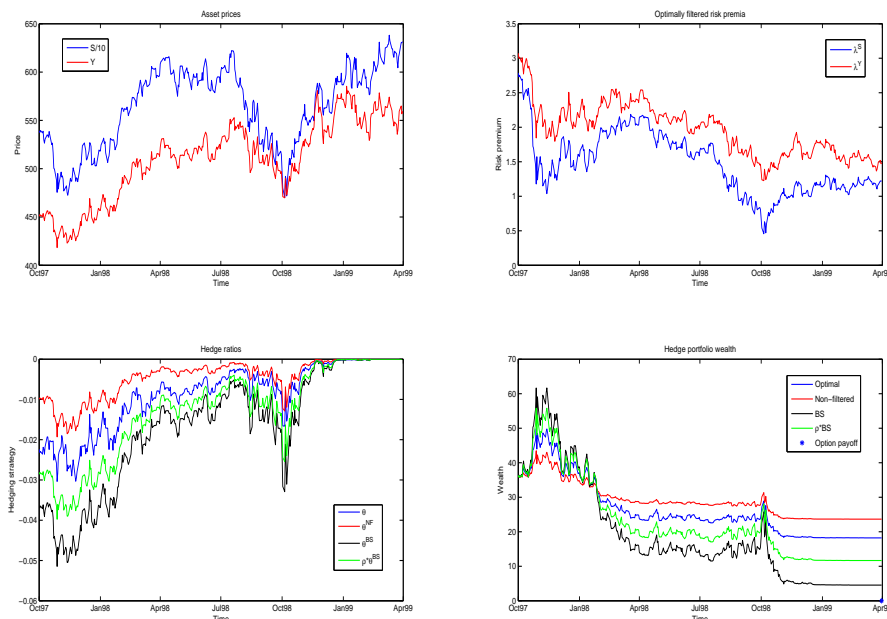| S_0	Y_0	K	λ_0^S	σ^S	λ_0^Y	σ^Y	$\mathrm{v}_0^S = \mathrm{v}_0^Y$	ρ	α
5392	452.7	452.7	2.7878	0.1727	3.0729	0.1624	1.9944	0.7726	0.01

Figure 2: Asset price paths, filtered risk premia and hedge portfolios from October 1997 to March 1999. Parameters are as in Table 11. The average values of the filtered risk premia over the hedging time-frame were $\bar{\lambda}^S = 1.8105$, $\bar{\lambda}^Y = 2.0241$. In the figure, all hedges begin with an initial wealth equal to the time zero BS price, $p^{\mathrm{BS}} = 35.82$. The utility-based option premia are $p_0 = 21.86$ (partial information) and $p_0^{\mathrm{NF}} = 7.39$ (without filtering).

hedge, 0.1267 for the BS hedge and 0.3253 for the correlation-modulated BS hedge. On this basis, the unfiltered hedge is the best performer, and this can be traced to the values of the prior risk premia versus the realised values over the hedging time-frame. We have $\hat{\nu}_0 = 0.9190$, while the average values of the filtered risk premia over $[0, T]$ give $\bar{\hat{\nu}} = 0.8216$. In other words, the option is unlikely to end up in the money, the unfiltered hedge takes an even higher estimate for the Q^M-drift of Y, and hence under-hedges and realises a larger profit than the filtered hedge. The correlation weighted BS hedge approximates ν by zero, tends to hedge unnecessarily aggressively, and generates losses. If the agents use their respective option premia as the initial wealth, then the unfiltered hedge performs much worse than the filtered hedge. The terminal error as a fraction of the BS price is then 0.1189 for the hedge with filtering, -0.1333 without filtering. Although on this basis both hedges then do worse than the modified BS hedge, it is significant that the partial information model has generated a terminal profit even though its prior is not fortuitously set up to do so.

These empirical examples tend to reinforce the conclusions from the simulation experiments, that a partial information approach can have benefits, but these are not categorical. Any improvement comes in part from the updating procedure applied to the dynamic hedge, but a significant benefit of the partial information method is encapsulated in the increased option

24

premium, that takes into account the uncertainty in the drift parameters.

8   Conclusions

Using a Kalman filter and a Gaussian prior distribution for the uncertain risk premia of a two-asset basis risk model, we have derived representations and an asymptotic expansion for utility-based prices and hedging strategies. Analytic results are possible for zero risk aversion, and semi-analytic results for small risk aversion. Simulation results show that the filtering procedure can help the hedging performance of the utility-based hedge, but the improvements are not universal. If we incorporate the agent's subjective valuation of the claim into the initial wealth, the partial information model almost invariably performs better, by incorporating extra insurance against parameter uncertainty.

This highlights a feature of utility indifference methods that some may regard as a drawback. They succeed primarilty by charging an appropriate price for a non-exchange traded option. This is implicit by definition, since they insist on achieving utility indifference by setting an appropriate initial option price. Developing other methods of hedging in incomplete markets, which work well independently of which price is charged for a claim, with an associated numerical and empirical study, would be a good topic for future research.

References

[1] S. Ankirchner and P. Imkeller, *Hedging with residual risk: a BSDE approach.* Preprint, 2009.

[2] S. Ankirchner, P. Imkeller, and A. Popier, *Optimal cross hedging of insurance derivatives*, Stoch. Anal. Appl., 26 (2008), pp. 679–709.

[3] S. Ankirchner, P. Imkeller, and G. Reis, *Pricing and hedging of derivatives based on non-tradable underlyings*, Math. Finance, (2009). In press.

[4] D. Becherer, *Utility-indifference hedging and valuation via reaction-diffusion systems*, Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci., 460 (2004), pp. 27–51.

[5] ——, *Bounded solutions to backward SDEs with jumps for utility optimization and indifference hedging*, Ann. Appl. Probab., 16 (2006), pp. 2027–2054.

[6] F. E. Benth and K. H. Karlsen, *A PDE representation of the density of the minimal entropy martingale measure in stochastic volatility markets*, Stochastics, 77 (2005), pp. 109–137.

[7] T. Björk, M. H. A. Davis, and L. C., *Optimal investment under partial information.* Preprint, 2008.

[8] S. Brendle, *Portfolio selection under incomplete information*, Stochastic Process. Appl., 116 (2006), pp. 701–723.

[9] P. Collin Dufresne and J. Hugonnier, *Pricing and hedging in the presence of extraneous risks*, Stochastic Process. Appl., 117 (2007), pp. 742–765.

[10] M. H. A. Davis, *Optimal hedging with basis risk*, in From stochastic calculus to mathematical finance, Y. Kabanov, R. Liptser, and J. Stoyanov, eds., Springer, Berlin, 2006, pp. 169–187.

[11] F. Delbaen, P. Grandits, T. Rheinländer, D. Samperi, M. Schweizer, and C. Stricker, *Exponential hedging and entropic penalties*, Math. Finance, 12 (2002), pp. 99–123.

[12] W. H. Fleming and R. W. Rishel, *Deterministic and stochastic optimal control*, Springer-Verlag, Berlin, 1975. Applications of Mathematics, No. 1.

[13] H. Föllmer and M. Schweizer, *Hedging of contingent claims under incomplete information*, in Applied stochastic analysis (London, 1989), vol. 5 of Stochastics Monogr., Gordon and Breach, New York, 1991, pp. 389–414.

[14] H. Föllmer and D. Sondermann, *Hedging of nonredundant contingent claims*, in Contributions to mathematical economics, North-Holland, Amsterdam, 1986, pp. 205–223.

[15] V. Henderson, *Valuation of claims on nontraded assets using utility maximization*, Math. Finance, 12 (2002), pp. 351–373.

[16] H. Hulley and T. A. McWalter, *Quadratic hedging of basis risk.* Preprint, 2008.

[17] J. Kallsen and T. Rheinländer, *Asymptotic utility-based pricing and hedging for exponential utility*. Preprint, 2008.

[18] D. Kramkov and M. Sîrbu, *Asymptotic analysis of utility-based hedging strategies for small number of contingent claims*, Stochastic Process. Appl., 117 (2007), pp. 1606–1620.

[19] P. Lakner, *Optimal trading strategy for an investor: the case of partial information*, Stochastic Process. Appl., 76 (1998), pp. 77–97.

[20] M. Mania and M. Schweizer, *Dynamic exponential utility indifference valuation*, Ann. Appl. Probab., 15 (2005), pp. 2113–2143.

[21] M. Monoyios, *Performance of utility-based strategies for hedging basis risk*, Quant. Finance, 4 (2004), pp. 245–255.

[22] ———, *Optimal hedging and parameter uncertainty*, IMA J. Manag. Math., 18 (2007), pp. 331–351.

[23] M. Musiela and T. Zariphopoulou, *An example of indifference prices under exponential preferences*, Finance Stoch., 8 (2004), pp. 229–239.

[24] H. Nagai and S. Peng, *Risk-sensitive dynamic portfolio optimization with partial information on infinite time horizon*, Ann. Appl. Probab., 12 (2002), pp. 173–195.

[25] H. Pham, *Mean-variance hedging for partially observed drift processes*, Int. J. Theor. Appl. Finance, 4 (2001), pp. 263–284. Information modeling in finance (Évry, 2000).

[26] ———, *Smooth solutions to optimal investment models with stochastic volatilities and portfolio constraints*, Appl. Math. Optim., 46 (2002), pp. 55–78.

[27] L. C. G. Rogers, *The relaxed investor and parameter uncertainty*, Finance Stoch., 5 (2001), pp. 131–154.

[28] J. Xiong and X. Y. Zhou, *Mean-variance portfolio selection under partial information*, SIAM J. Control Optim., 46 (2007), pp. 156–175.