

Accounting for Risk Aversion in Derivatives Purchase Timing

Tim Leung*

Mike Ludkovski†

September 15, 2011

Abstract

We study the problem of optimal timing to buy/sell derivatives by a risk-averse agent in incomplete markets. Adopting the exponential utility indifference valuation, we investigate this timing flexibility and the associated delayed purchase premium. This leads to a stochastic control and optimal stopping problem that combines the observed market price dynamics and the agent's risk preferences. Our results extend recent work on indifference valuation of American options, as well as the authors' first paper (Leung and Ludkovski, SIAM J. Fin. Math., 2011). In the case of Markovian models of contracts on non-traded assets, we provide analytical characterizations and numerical studies of the optimal purchase strategies, with applications to both equity and credit derivatives.

Keywords: sequential purchase timing, indifference pricing, exponential utility, stochastic control with optimal stopping

JEL Classifications: G12, G13, C68

1 Introduction

The problems of derivatives pricing and trading in incomplete markets are among the central themes in mathematical finance. Since in incomplete markets not all risks can be hedged away, it is important to model investors' attitudes towards risks. One major approach is the framework of indifference valuation, originally proposed by Hodges and Neuberger (1989). This is an extension of the static certainty equivalence concept that incorporates risk aversion via a utility function and imperfect dynamic hedging into derivative pricing. The investor's subjective price for a derivative, called the indifference price, is derived by comparing the investor's utility maximization problems with and without the claim.

In existing literature, the indifference price is typically used "statically" as a reservation price for risk averse derivative buyers or sellers (see, for example, Carmona (2008) and references therein). From the perspective of a potential buyer, a derivative that costs today more than its indifference price is deemed too expensive, and therefore should not be purchased. In contrast, if the prevailing market ask price is lower than the prospective buyer's indifference price, it is not clear whether the buyer should buy the claim immediately or wait for a potentially better deal in the future. The answer depends on the precise motives of the buyer, but it raises the idea of the *timing option* inherent in this investment decision.

*IEOR Department, Columbia University, New York, NY 10027; email: leung@ieor.columbia.edu. Work partially supported by NSF grant DMS-0908295.

†Department of Statistics & Applied Probability, University of California Santa Barbara, Santa Barbara CA 93106; email: ludkovski@pstat.ucsb.edu.

Motivated by this observation, we study the problem of optimal timing to buy a given derivative from the perspective of a risk-averse investor. To analyze this question, we apply the exponential indifference pricing methodology, which leads to a utility maximization problem with optimal stopping. Intuitively, the purchase timing decision is related to the stochastic spread between the investor’s indifference price h_t and the market price P_t . While the indifference price is formulated under the historical measure, the market price is likely to be computed from a risk-neutral pricing measure, exogenous to the investor. As a result, the purchase timing will also necessarily depend on the interaction between the investor’s and the market’s pricing rules. This phenomenon also arises in our prior work (Leung and Ludkovski, 2011) which investigated the case of a risk-neutral investor (but whose pricing measure was different from the market’s).

To measure the benefit of the timing option, we study the associated *delayed purchase premium* that was introduced in Leung and Ludkovski (2011). In particular, using the duality properties of exponential hedging (see, e.g., Delbaen et al. (2002); Becherer (2003); Rouge and El Karoui (2000)), we derive a probabilistic representation for the delayed purchase premium in a general semimartingale framework. The resulting Theorem 2 (see also Proposition 4) establishes connections with classical indifference pricing and gives insights to the optimal purchase strategy, as well as the extreme cases with large or zero risk aversion.

To illustrate our analysis, in Section 3 we consider the optimal timing problem under a parametric market model with a non-traded underlying asset. This incomplete market setting, sometimes called the basis risk model, has been adopted for utility-based valuation for a number of applications, such as weather derivatives (Davis, 2001), commodities (Davis, 2006), credit derivatives (Leung et al., 2008; Sircar and Zariphopoulou, 2010; Jaimungal and Sigloch, 2010), real options (Henderson, 2007), and employee stock options (Henderson, 2005; Leung and Sircar, 2009a). With basis risk, the delayed purchase premium for a generic contingent claim under exponential utility involves the stochastic bracket between the market price and a density process, plus a quadratic penalization around a benchmark risk premium (see Theorem 7). This allows us to conveniently identify the scenarios where immediate (or never-at-all) purchase is optimal. By numerically solving the corresponding variational inequality, we present the optimal purchase boundaries for digital options, as well as defaultable bonds.

Contrary to risk-neutral pricing, the indifference pricing rule is not linear in quantity. Consequently, if a risk-averse investor wishes to buy multiple contracts of the same option, she will tend to spread her purchases over time (while a risk-neutral investor will buy all at once). To highlight this disparity, we study in Section 5 the problem of sequential option purchase under exponential utility. We introduce the concept of *marginal* delayed purchase premium, which measures the value of optimally waiting to make each incremental purchase. In the non-traded asset model, the investor’s optimal policy is described a series of purchase boundaries along which the marginal delayed purchase premium is zero.

Complementary to our problem of *when* to buy, the more classical question of *“how much?”* can be analyzed by considering the investor’s optimal static position. In particular, since the buyer’s indifference price is increasing concave in quantity, the answer is determined by equating the marginal indifference price with the market price; see İlhan et al. (2005). In another related work, Kramkov and Bank (2010) study dynamic trading among risk-averse market makers and provide a mathematical characterization of Pareto optimal allocations.

The remainder of this paper is organized as follows. Section 2 describes the precise mathematical setup we use to model the timing flexibility in a general semimartingale framework. Section 3 then specializes to the case of Markovian models for non-traded assets where we are able to obtain explicit results by building on existing rich literature on indifference pricing. In particular, we show that under exponential utility the timing flexibility value function admits the same dual characterization

as the risk-neutral case plus a quadratic entropic penalty. Section 4 then presents two illustrative examples with detailed numerical results and figures. Finally, Sections 5 and 6 discuss extensions based on our model and conclude the paper.

2 Model

Throughout, we consider a risk-averse investor whose risk preferences are described by the exponential utility function $U : \mathbb{R} \mapsto \mathbb{R}_-$ defined by

$$U(x) = -e^{-\gamma x}, \quad x \in \mathbb{R},$$

where $\gamma > 0$ is the coefficient of absolute risk aversion. Precisely, $U(x)$ is the investor's utility for having discounted wealth x at the end of the investment horizon T .

In the background, we assume a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$, which satisfies the usual conditions of right continuity and completeness. We shall use the notation $\mathbb{E}_t\{\cdot\} \equiv \mathbb{E}\{\cdot | \mathcal{F}_t\}$ for the conditional expectation given \mathcal{F}_t under \mathbb{P} .

The basic trading assets consist of a riskless bond that pays interest at constant rate $r \geq 0$, and a risky asset whose discounted price process is a non-negative \mathbb{F} -locally bounded semimartingale $(S_t)_{0 \leq t \leq T}$. We denote by $(X_t^\theta)_{0 \leq t \leq T}$ the discounted trading wealth process with a self-financing dynamic trading strategy $(\theta_t)_{0 \leq t \leq T}$ which represents the number of shares held at time t . With initial capital X_t at time $t \in [0, T]$, the discounted wealth at a later date $u \in [t, T]$ is given by

$$X_u^\theta = X_t + G_{t,u}(\theta), \quad \text{with} \quad G_{t,u}(\theta) := \int_t^u \theta_s dS_s. \quad (2.1)$$

The stochastic integral $G_{t,u}(\theta)$ is the discounted capital gains or losses from trading with strategy θ from time t to u .

We first consider the portfolio optimization problem where a static derivative position is incorporated. Specifically, the risk-averse investor dynamically trades in the riskless and risky assets throughout the horizon $[0, T]$. In addition, the investor also holds $\alpha \geq 0$ units of a derivative till expiration, where the terminal payoff is $D \in \mathcal{F}_T$. For an investor with initial wealth X_t at time $t \in [0, T]$, her maximal expected utility from terminal wealth is

$$V_t(X_t; \alpha) := \text{ess sup}_{\theta \in \Theta_{t,T}} \mathbb{E}_t \left\{ U(X_T^\theta + \alpha D) \right\}, \quad (2.2)$$

where the precise definition of admissible trading strategies is given below in (2.11).

When there is no derivative ($\alpha = 0$), the optimization (2.2) reduces to the Merton portfolio optimization problem. We denote the Merton value function by

$$M_t(X_t) := V_t(X_t; 0). \quad (2.3)$$

The investor's indifference price $h_t \equiv h_t(\alpha)$ for holding α units of option D is found from comparing the maximal expected utility with and without the derivative. It satisfies the indifference equation:

$$M_t(X_t + h_t) := V_t(X_t; \alpha). \quad (2.4)$$

The indifference price is the investor's subjective valuation, which may differ from the actual cost of buying the derivative from the market. In this paper, we assume that the investor has no influence over market prices of options and their underlying assets. As is standard in no-arbitrage pricing, the market price of an option is given by the expectation under some equivalent martingale measure (EMM) $Q^* \sim \mathbb{P}$. Therefore, the market (ask) price for claim D is given by

$$P_t = \mathbb{E}^{Q^*} \{ e^{-r(T-t)} D | \mathcal{F}_t \}, \quad 0 \leq t \leq T. \quad (2.5)$$

2.1 Timing Problem

An investor who intends to buy/sell derivatives in the market has the option to time her trade. For clarity of exposition, we henceforth focus on the purchase timing problem; the case of optimally timing derivative sales can be studied similarly. In Section 6, we also discuss an extension to the sequential buying and selling problem.

Denote by \mathcal{T} the set of all stopping times with respect to \mathbb{F} taking values in $[0, T]$. This will be the collection of all admissible purchase times for investor. For any stopping times $s, u \in \mathcal{T}$ with $s \leq u$, we set $\mathcal{T}_{s,u} := \{\tau \in \mathcal{T} : s \leq \tau \leq u\}$. After the purchase, the investor continues to dynamically trade till the expiration date T . At time $t \leq T$, the investor faces the combined stochastic control and optimal stopping problem:

$$J_t(X_t; \alpha) = \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,sup}_{\theta \in \Theta_{t,\tau}} \mathbb{E}_t \left\{ V_\tau(X_\tau^\theta - \alpha P_\tau; \alpha) \right\} \quad (2.6)$$

$$= \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,sup}_{\theta \in \Theta_{t,\tau}} \mathbb{E}_t \{ M_\tau(X_\tau^\theta + h_\tau - \alpha P_\tau) \}, \quad (2.7)$$

where the second equality follows from (2.4). For any choice of purchase date $\tau \leq T$, the investor's trading strategy over the period $[\tau, T]$ (after purchase) is implicitly optimized in the value function V_τ in (2.6).

Alternatively, we can interpret problem (2.7) as if the investor is optimally timing to exercise an American claim with payoff $h - \alpha P$. At expiration $h_T - \alpha P_T = 0$, so the choice of $\tau = T$ reduces J_t to the Merton function M_t , and we have $J_t(X_t; \alpha) \geq M_t(X_t)$. Henceforth, we interpret $\tau = T$ as the investor never purchases the derivative.

In order to quantify the value of optimally timing to buy the derivatives rather than buying them immediately, we compare the value functions J_t with V_t . Precisely, we define the delayed purchase premium $L_t \equiv L_t(\alpha)$ via the equation:

$$V_t(X_t + L_t - \alpha P_t; \alpha) := J_t(X_t; \alpha). \quad (2.8)$$

Since V_t is increasing in wealth and from (2.6) $J_t(X_t; \alpha) \geq V_t(X_t - \alpha P_t; \alpha)$, we infer that $L_t \geq 0$. If $L_t > 0$ for some $t < T$, then it is not optimal to buy at t because there is a strictly positive benefit of delaying the purchase.

On the other hand, we can apply the indifference equation (2.4) to (2.8) and write

$$J_t(X_t; \alpha) = M_t(X_t + h_t - \alpha P_t + L_t). \quad (2.9)$$

In view of (2.9), we define

$$f_t := h_t - \alpha P_t + L_t \geq 0, \quad (2.10)$$

which can be interpreted as the *indifference value* for the opportunity to optimally time to buy and hold α units of option D till maturity. In fact, (2.10) reflects the decomposition of the indifference value f_t into three parts: the indifference price h_t for holding the options, plus the delayed purchase premium L_t , minus the cost of the options αP_t . Also, whenever $h_t < \alpha P_t$ for some $t < T$, then $L_t > 0$ since $f_t \geq 0$. This confirms the intuition that the investor should wait if the market price strictly dominates her own indifference price. Note that f_t , L_t , and h_t all depend on risk aversion γ and are typically not linear in quantity α .

2.2 Duality Representation

To better understand the structure of the indifference value f_t , in this section we establish a duality representation for f_t in terms of entropic penalties. Related studies on exponential hedging in general semimartingale incomplete markets can be found in, among others, Becherer (2001), Delbaen et al. (2002), and Leung and Sircar (2009b).

For any measure Q , the relative entropy of Q with respect to \mathbb{P} is given by

$$H(Q|\mathbb{P}) := \begin{cases} \mathbb{E}^Q \left\{ \log \frac{dQ}{d\mathbb{P}} \right\}, & Q \ll \mathbb{P}, \\ +\infty, & \text{otherwise.} \end{cases}$$

Let \mathcal{P}_f be the set of equivalent local martingale measures with finite relative entropy with respect to \mathbb{P} . We assume that $\mathcal{P}_f \neq \emptyset$ is non-empty, and that the market pricing measure satisfies $Q^* \in \mathcal{P}_f$ (see (2.5)). Our set of admissible self-financing strategies is

$$\Theta \equiv \Theta_{0,T} := \{ \theta \in \mathbb{L}(S) \mid G_{0,T}(\theta) \text{ is a } (Q, \mathbb{F})\text{-martingale for all } Q \in \mathcal{P}_f \}, \quad (2.11)$$

where $\mathbb{L}(S)$ is the set of \mathbb{F} -predictable S -integrable \mathbb{R} -valued processes.

Theorems 2.1 and 2.2 of Frittelli (2000) guarantee that there is a unique minimizer $Q^E \in \mathcal{P}_f$,

$$Q^E := \arg \min_{Q \in \mathcal{P}_f} H(Q|\mathbb{P}). \quad (2.12)$$

This measure is called the minimal entropy martingale measure (MEMM).

Definition 1. For $Q \in \mathcal{P}_f$, let $Z_t^{Q,\mathbb{P}} := \mathbb{E}_t \left\{ \frac{dQ}{d\mathbb{P}} \right\}$ denote the density process of Q with respect to \mathbb{P} . The conditional relative entropy of Q with respect to \mathbb{P} over the time interval $[t, u]$ is defined via

$$H_t^u(Q|\mathbb{P}) := \mathbb{E}_t^Q \left\{ \log \frac{Z_u^{Q,\mathbb{P}}}{Z_t^{Q,\mathbb{P}}} \right\}, \quad 0 \leq t \leq u \leq T. \quad (2.13)$$

For any $t \in [0, T]$ and $Q \in \mathcal{P}_f$, the random variable $\log Z_t^{Q,\mathbb{P}}$ is Q -integrable (see Lemma 3.3 of Delbaen et al. (2002)), so the conditional relative entropy is well-defined. Also, Jensen's Inequality yields that $H_t^T(Q|\mathbb{P}) \geq 0$. By Proposition 4.1 of Kabanov and Stricker (2002), the MEMM Q^E also minimizes the conditional relative entropy $H_\tau^T(Q|\mathbb{P})$ at any $\tau \in \mathcal{T}$. Alternatively, treating Q^E as a prior measure, one can similarly compute the relative entropy $H_t^\tau(Q|Q^E)$ and define the corresponding set $\mathcal{P}_f(Q^E)$; as a mild technical condition, we assume that $\mathcal{P}_f(Q^E) = \mathcal{P}_f$.

The next Theorem gives the dual representations of J_t , f_t and L_t .

Theorem 2. The value function $J_t(X_t; \alpha)$ can be expressed as

$$J_t(X_t; \alpha) = U(X_t) \cdot \exp \left(- \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,inf}_{Q \in \mathcal{P}_f} \left(\gamma \mathbb{E}_t^Q \{ h_\tau - \alpha P_\tau \} + H_t^\tau(Q|\mathbb{P}) + \mathbb{E}_t^Q \{ H_\tau^T(Q^E|\mathbb{P}) \} \right) \right). \quad (2.14)$$

Moreover, the indifference value f_t is given by

$$f_t = \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,inf}_{Q \in \mathcal{P}_f} \left(\mathbb{E}_t^Q \{ h_\tau - \alpha P_\tau \} + \frac{1}{\gamma} H_t^\tau(Q|Q^E) \right), \quad (2.15)$$

and the delayed purchase premium is

$$L_t = \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,inf}_{Q \in \mathcal{P}_f} \left(\mathbb{E}_t^Q \{ h_\tau - \alpha P_\tau \} + \frac{1}{\gamma} H_t^\tau(Q|Q^E) \right) - (h_t - \alpha P_t). \quad (2.16)$$

Finally, the optimal purchase time τ^* is given by

$$\tau_t^* = \inf\{t \leq u \leq T : f_u = h_u - \alpha P_u\} = \inf\{t \leq u \leq T : L_u = 0\}. \quad (2.17)$$

Proof. For exponential utility, the Merton function admits the representation (see e.g. Theorem 1 of Delbaen et al. (2002))

$$M_t(X_t) = -e^{-\gamma X_t} e^{-H_t^T(Q^E|\mathbb{P})}. \quad (2.18)$$

Applying (2.18) to (2.9) we get

$$J_t(X_t; \alpha) = -e^{-\gamma(X_t + f_t)} e^{-H_t^T(Q^E|\mathbb{P})}. \quad (2.19)$$

Combining (2.19) with Propositions 2.4 and 2.8 of Leung and Sircar (2009b), where the early exercisable claim's payoff is now $h_\tau - \alpha P_\tau$ at exercise time τ , we immediately obtain (2.14) and (2.15). Substituting (2.15) into (2.10), the delayed purchase premium can be expressed as $L_t = f_t - (h_t - \alpha P_t)$, which leads to (2.16). \square

Equation (2.17) means that the investor should buy the option as soon as the delayed purchase premium L vanishes. We will further explore the structure of L under a parametric model in Section 3 (see Proposition 4).

2.3 Asymptotic Limits

Theorem 2 allows to obtain the asymptotic values of L_t for extreme values of risk aversion γ . Denote by $L_t(\gamma, \alpha)$ the delayed purchase premium in (2.10) for buying α options when the investor's risk aversion is $\gamma > 0$. Similarly, we use the notations $f_t(\gamma, \alpha)$ and $h_t(\gamma, \alpha)$ to highlight the dependence on γ and α .

By standard arguments (see e.g. Becherer (2001)), the indifference value $f_t(\gamma, \alpha)$ and indifference price $h_t(\gamma, \alpha)$ are decreasing in γ . However, the same may not hold for their difference that constitutes $L_t(\gamma, \alpha)$. In the next proposition, we show that the zero risk aversion limit is in fact *less* than the large risk aversion limit.

Proposition 3. *The delayed purchase premium in (2.10) admits the following limits:*

$$\lim_{\gamma \rightarrow 0} L_t(\gamma, \alpha) = \lim_{\alpha \rightarrow 0} \frac{L_t(\gamma, \alpha)}{\alpha} = \alpha \cdot (P_t - P_t^{E*}) =: \alpha \cdot L_t^E, \quad (2.20)$$

$$\lim_{\gamma \rightarrow \infty} L_t(\gamma, \alpha) = \lim_{\alpha \rightarrow \infty} \frac{L_t(\gamma, \alpha)}{\alpha} = \alpha \cdot (P_t - \underline{h}_t) =: \alpha \cdot \bar{L}_t, \quad (2.21)$$

where

$$P_t^{E*} := \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_t^{Q^E} \{P_\tau\}, \quad \text{and} \quad \underline{h}_t := \operatorname{ess\,inf}_{Q \in \mathcal{P}_f} \mathbb{E}_t^Q \{D\}. \quad (2.22)$$

Proof. As $\gamma \searrow 0$, it follows from Proposition 1.3.4 of Becherer (2001) that $h_t(\gamma, \alpha) \nearrow \alpha \mathbb{E}_t^{Q^E} \{D\} =: \alpha h_t^E$, which is the risk-neutral price of D under the MEMM Q^E . By this and Proposition 2.18 of Leung and Sircar (2009b) with the early exercisable claim payoff being $h_\tau(\gamma, \alpha) - \alpha P_\tau$, we obtain the limit

$$\lim_{\gamma \rightarrow 0} f_t(\gamma, \alpha) = \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_t^{Q^E} \{\alpha h_\tau^E - \alpha P_\tau\} = \alpha (h_t^E - \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_t^{Q^E} \{P_\tau\}),$$

where the last equality holds by iterated expectation under Q^E . Applying these limits to (2.10) gives the limit in (2.20).

By Proposition 11 of Delbaen et al. (2002), as $\gamma \nearrow \infty$, $h_t(\gamma, \alpha) \searrow \alpha h_t$. Also, by Proposition 2.17 of Leung and Sircar (2009b) with payoff $h_\tau(\gamma, \alpha) - \alpha P_\tau$, one can show that

$$\lim_{\gamma \rightarrow \infty} f_t(\gamma, \alpha) = \alpha \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,inf}_{Q \in \mathcal{P}_f} \mathbb{E}_t^Q \{h_\tau - \alpha P_\tau\} = 0. \quad (2.23)$$

For the last equality, note that $h_\tau \leq \alpha P_\tau$ for all τ, Q and $h_T = \alpha P_T = \alpha D$, so the choice of $\tau = T$ (under any $Q \in \mathbb{P}_f$) yields the maximum value zero. Applying this to (2.10) yields (2.21).

It is also well known that the indifference price $h_t(\gamma, \alpha)$ has the scaling property: $h_t(\gamma, \alpha)/\alpha = h_t(\alpha\gamma, 1)$; see Becherer (2001). Applying this to (2.16), we deduce the same property for the delayed purchase premium, namely, $L_t(\gamma, \alpha)/\alpha = L_t(\alpha\gamma, 1)$. With this, the risk-aversion limits (2.20) and (2.21) can be interpreted as the stated large-volume and small-volume limits for the *average* delayed purchase premium $L_t(\gamma, \alpha)/\alpha$. \square

In both risk aversion limits, the investor's indifference prices and delayed purchase premia become linear in quantity. In the zero risk aversion case, the investor's indifference price limit $h_t^E = \mathbb{E}_t^{Q^E} \{D\}$ is also referred to as the Davis price (see Davis (1997)). The investor's optimal purchase timing is found from P^{E*} , which is independent of quantity. In the large risk aversion case, the investor will never buy the option since $\tau = T$ is optimal (see (2.23)).

To better understand (2.20), we use the following equality:

$$P_t^{E*} = \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_t^{Q^E} \{P_\tau\} = \operatorname{ess\,inf}_{\{Q^\tau\}_{\tau \in \mathcal{T}}} \mathbb{E}_t^{Q^\tau} \{D\}, \quad (2.24)$$

where each $Q^\tau \in \mathcal{P}_f$ is a probability measure whose density process with respect to \mathbb{P} is defined by

$$Z_t^{Q^\tau, \mathbb{P}} := Z_t^{Q^E, \mathbb{P}} \mathbb{1}_{[0, \tau)}(t) + Z_t^{Q^*, \mathbb{P}} \frac{Z_\tau^{Q^E, \mathbb{P}}}{Z_\tau^{Q^*, \mathbb{P}}} \mathbb{1}_{[\tau, T]}(t), \quad 0 \leq t \leq T. \quad (2.25)$$

Intuitively, the probability measure Q^τ is identical to the MEMM Q^E up to the \mathcal{F} -stopping time τ and then coincides with the market measure Q^* over $(\tau, T]$. The equality (2.24) reveals that minimization over stopping times under a *single* measure can be cast as minimization over the collection of pricing measures $\{Q^\tau\}_{\tau \in \mathcal{T}}$ parametrized by stopping time τ . This interpretation is referred to as the τ -optimal concatenation of pricing measures (see Proposition 2.2 of Leung and Ludkovski (2011)), while concatenation of the density processes is also used in other financial applications (see, for example, Delbaen (2006) and Riedel (2009)). Given the optimal stopping time τ^* , the right-hand side of (2.24) corresponds to pricing D under the special EMM $Q^{\tau^*} \in \mathcal{P}_f$. Moreover, since $\{Q^\tau\}_{\tau \in \mathcal{T}} \subseteq \mathcal{P}_f$, we deduce that $h_t \leq P_t^{E*}$, and therefore, $L_t^E \leq \bar{L}_t$. Our numerical experiments suggest that L is monotone in γ (*resp.* in α), and a more risk averse agent postpones derivative purchases, i.e. τ^* is increasing in γ (*resp.* in α). We are not able to establish this property in generality, because γ affects both h_t and the optimal stopping problem for L .

Furthermore, the zero risk-aversion limit can be viewed as a special case of the delayed purchase premium in Section 2.3 of Leung and Ludkovski (2011) where the investor's pricing measure is the MEMM Q^E . Proposition 3 provides an intuitive mechanism why the investor and market measures might differ: the market reflects a risk-neutral Q^* while the investor applies utility-based framework under the physical \mathbb{P} to end up with Q^E in the small- γ limit. As in Leung and Ludkovski (2011), P^{E*} can be regarded as the minimized expected cost of acquiring the option D given the prevailing price process P .

In view of our general analysis above, it is clear that tractable results are possible as soon as the investor price h and the market price P are available in closed-form. Consequently, we are able to study any number of models that have obtained explicit expressions for indifference prices. In Sections 4.1-4.2 we consider two such parametric models arising in trading of illiquid assets and defaultable bonds, respectively.

3 Buying Options on a Non-Traded Asset

We first illustrate our previous analysis in the classical setting of a Markovian market with a liquidly traded asset S and a non-traded asset Y . The respective prices are modeled by the stochastic differential equations (SDEs):

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \quad (\text{traded}) \quad (3.1)$$

$$dY_t = b(t, Y_t) dt + c(t, Y_t) (\rho dW_t + \hat{\rho} d\hat{W}_t), \quad (\text{non-traded}) \quad (3.2)$$

where W and \hat{W} are two independent Brownian motions under the measure \mathbb{P} , $\sigma \geq 0$, and $\hat{\rho} := \sqrt{1 - \rho^2}$. The filtration \mathbb{F} is generated by (W, \hat{W}) . The drift and diffusion coefficients b and $c \geq 0$ are chosen so that a unique strong solution exists for SDEs (3.1)-(3.2). The derivative claim in question is a European option with discounted bounded payoff $D(Y_T)$ at expiration date T . Very similar setups have appeared in the indifference pricing literature, including Musiela and Zariphopoulou (2004), Henderson (2005), and Davis (2006). For notational simplicity, we set the interest rate to be zero.

Suppose an investor is holding α contracts of D , and dynamically trades S as a partial hedge. Her trading wealth follows the SDE

$$dX_t^\theta = \sigma \theta_t (\lambda dt + dW_t), \quad (3.3)$$

where

$$\lambda := \mu/\sigma$$

is the Sharpe ratio of S , and $(\theta_t)_{0 \leq t \leq T}$ is the cash amount invested in S satisfying $\mathbb{E}\{\int_0^T \theta_t^2 dt\} < \infty$. The maximal expected utility from terminal wealth is given by

$$V(t, x, y) = \sup_{(\theta_u)_{t \leq u \leq T}} \mathbb{E} \left\{ U(X_T^\theta + \alpha D(Y_T)) \mid X_t = x, Y_t = y \right\}. \quad (3.4)$$

The function V solves a nonlinear PDE of HJB type. As studied in, for example, Musiela and Zariphopoulou (2004), the holder's indifference price $h(t, y)$ is independent of wealth x and satisfies

$$V(t, x, y) = -e^{-\gamma(x+h(t,y)) - \frac{\lambda^2}{2}(T-t)}.$$

It can be determined as the (unique viscosity) solution of the semilinear PDE:

$$h_t + \mathcal{L}^0 h - \frac{\gamma}{2}(1 - \rho^2)c^2(t, y)h_y^2 = 0, \quad (3.5)$$

on $(t, y) \in [0, T) \times \mathbb{R}_+$, with terminal condition $h(T, y) = D(y)$. Here, the differential operator is

$$\mathcal{L}^0 := \frac{c^2(t, y)}{2} \frac{\partial^2}{\partial y^2} + [b(t, y) - \rho \lambda c(t, y)] \frac{\partial}{\partial y}. \quad (3.6)$$

Next, we summarize some results on the dual representation of the indifference price $h(t, y)$. The set of EMMs with respect to \mathbb{P} on \mathcal{F}_T is characterized by the stochastic exponential

$$\frac{dQ^\phi}{d\mathbb{P}} \Big|_{\mathcal{F}_t} = \exp \left(-\frac{1}{2} \int_0^t (\lambda^2 + \phi_s^2) ds - \int_0^t \lambda dW_s - \int_0^t \phi_s d\hat{W}_s \right), \quad (3.7)$$

where $(\phi_t)_{0 \leq t \leq T}$ is a progressively measurable process satisfying $\mathbb{E}\{\int_0^T \phi_s^2 ds\} < \infty$ and $\mathbb{E}\{Z_T^\phi\} = 1$. Under measure Q^ϕ , $W_t^\phi = W_t + \lambda t$ and $\hat{W}_t^\phi = \hat{W}_t + \int_0^t \phi_s ds$ are independent Brownian motions. The process ϕ is premium for the idiosyncratic risk represented by the second Brownian motion \hat{W} . Throughout this section, we shall consider Markovian risk premia of the form $\phi_t = \phi(t, Y_t)$ for some deterministic function $\phi(t, y)$. Under a given EMM Q^ϕ , the associated infinitesimal generator of Y is given by

$$\mathcal{L}^\phi = \frac{c^2(t, y)}{2} \frac{\partial^2}{\partial y^2} + [b(t, y) - \lambda \rho c(t, y) - \phi(t, y) \hat{\rho} c(t, y)] \frac{\partial}{\partial y}. \quad (3.8)$$

In particular, \mathcal{L}^0 in (3.6) corresponds to the risk premium $\phi(t, y) = 0$ and the associated measure Q^0 is called the minimal martingale measure (MMM) (see Föllmer and Schweizer (1990)).

Consequently, the conditional relative entropy of any Q^ϕ with respect to \mathbb{P} is simply a quadratic penalty term, namely

$$H_t^\tau(Q^\phi | \mathbb{P}) = \mathbb{E}_{t, y}^\phi \left\{ \int_t^\tau \frac{\lambda^2 + \phi_s^2}{2} ds \right\}, \quad (3.9)$$

where we use the shorthand $\mathbb{E}_{t, y}^\phi \{\cdot\} \equiv \mathbb{E}^\phi \{\cdot | Y_t = y\}$. Under this model, the minimal entropy martingale measure Q^E with respect to \mathbb{P} on \mathcal{F}_T is simply the MMM Q^0 .

Applying the well-known duality results for exponential indifference prices (see, e.g., Delbaen et al. (2002)), the dual representation for $h(t, y)$ is given by

$$h(t, y) = \inf_\phi \mathbb{E}_{t, y}^\phi \left\{ D(Y_T) + \frac{1}{\gamma} \int_t^T \frac{\phi_s^2}{2} ds \right\}, \quad (3.10)$$

and the associated minimizer ϕ^* is given in feedback form:

$$\phi^*(t, y) = \gamma \hat{\rho} c(t, y) h_y(t, y). \quad (3.11)$$

Now suppose the market prices options with the EMM Q^ψ , with idiosyncratic risk premium $\psi(t, y)$ for the second Brownian motion \hat{W} . Then, the discounted option price is the Q^ψ -martingale $P(t, y) = \mathbb{E}_{t, y}^\psi \{D(Y_T)\}$, solving the linear PDE

$$P_t + \mathcal{L}^\psi P = 0, \quad (3.12)$$

on $(t, y) \in [0, T) \times \mathbb{R}_+$, with $P(T, y) = D(y)$.

3.1 Analytic Representation

Given $h(t, y)$ and $P(t, y)$, we can express the indifference value $f(t, y)$ according to (2.15), namely,

$$f(t, y) = \sup_{t \leq \tau \leq T} \inf_\phi \mathbb{E}_{t, y}^\phi \left\{ h(\tau, Y_\tau) - \alpha P(\tau, Y_\tau) + \frac{1}{\gamma} \int_t^\tau \frac{\phi_s^2}{2} ds \right\}, \quad (3.13)$$

where the last term is the relative entropy with respect to Q^E , $H_t^\tau(Q^\phi | Q^E) = \mathbb{E}_{t, y}^\phi \left\{ \int_t^\tau \frac{\phi_s^2}{2} ds \right\}$. In turn, we derive a new expression for the delayed purchase premium.

Proposition 4. *The delayed purchase premium admits the representation:*

$$L(t, y) = \sup_{t \leq \tau \leq T} \inf_{\phi} \mathbb{E}_{t,y}^{\phi} \left\{ \int_t^{\tau} \frac{1}{2\gamma} (\phi_s - \phi^*(s, Y_s))^2 + \alpha \hat{\rho} c(s, Y_s) P_y(s, Y_s) (\phi_s - \psi(s, Y_s)) ds \right\}, \quad (3.14)$$

where $\phi^*(t, y)$ is given in (3.11).

Proof. Recall from (2.10) that $L(t, y) = f(t, y) - h(t, y) + \alpha P(t, y)$. By Girsanov's Theorem the indifference price and market price follow the SDEs

$$dh(t, Y_t) = \left(\frac{\gamma}{2} (1 - \rho^2) c^2(t, Y_t) h_y^2(t, Y_t) - \phi_t \hat{\rho} c(t, Y_t) h_y(t, Y_t) \right) dt + c(t, Y_t) h_y(t, Y_t) (\rho dW_t^{\phi} + \hat{\rho} d\hat{W}_t^{\phi}), \quad (3.15)$$

$$dP(t, Y_t) = -(\phi_t - \psi(t, Y_t)) c(t, Y_t) \hat{\rho} P_y(t, Y_t) dt + c(t, Y_t) P_y(t, Y_t) (\rho dW_t^{\phi} + \hat{\rho} d\hat{W}_t^{\phi}). \quad (3.16)$$

Substituting (3.15), (3.16) and (3.13) into (2.10) yields that

$$\begin{aligned} L(t, y) &= \sup_{t \leq \tau \leq T} \inf_{\phi} \mathbb{E}^{\phi} \left\{ h(\tau, Y_{\tau}) - \alpha P(\tau, Y_{\tau}) + \frac{1}{\gamma} \int_t^{\tau} \frac{\phi_s^2}{2} ds \mid Y_t = y \right\} - h(t, y) + \alpha P(t, y) \quad (3.17) \\ &= \sup_{t \leq \tau \leq T} \inf_{\phi} \mathbb{E}_{t,y}^{\phi} \left\{ \int_t^{\tau} \frac{\phi_s^2}{2\gamma} + \phi_s \hat{\rho} c(s, Y_s) (\alpha P_y(s, Y_s) - h_y(s, Y_s)) \right. \\ &\quad \left. - \psi(s, Y_s) \hat{\rho} \alpha c(s, Y_s) P_y(s, Y_s) + \frac{\gamma}{2} \hat{\rho}^2 c^2(s, Y_s) h_y^2(s, Y_s) ds \right\}. \end{aligned}$$

Then by completing the square in terms of ϕ and using $\phi^*(t, y)$ from (3.11), we obtain (3.14). \square

Proposition 4 reveals a convenient structure of the delayed purchase premium in terms of the corresponding premia: the optimized generic ϕ , the entropic ϕ^* , and the market ψ . In particular, the first integrand in (3.14) involves the quadratic penalization $(\phi - \phi^*)^2$, while the second term depends on the difference $(\phi - \psi)$. If the overall integrand in (3.14) is positive for all choices of ϕ for all (t, y) , then it is clear that it is optimal to delay the purchase till T .

Moreover, looking at expression (3.14) more carefully, the second term in fact involves

$$dP(s, Y_s) dZ_s = -Z_s \hat{\rho} c(s, Y_s) P_y(s, Y_s) (\phi_s - \psi_s) ds,$$

where $Z_t \equiv Z_t^{\phi, \psi} := \mathbb{E}_t^{\psi} \left\{ \frac{dQ^{\phi}}{dQ^{\psi}} \right\}$ is the density process of Q^{ϕ} with respect to Q^{ψ} . Therefore, the delayed purchase premium can be expressed in terms of the quadratic covariation between the market price and the density process, along with a quadratic penalty scaled by risk aversion. Precisely, we have

$$L(t, y) = \sup_{t \leq \tau \leq T} \inf_{\phi} \mathbb{E}_{t,y}^{\phi} \left\{ \int_t^{\tau} \frac{1}{2\gamma} (\phi_s - \phi^*(s, Y_s))^2 ds - \alpha \int_t^{\tau} Z_s^{-1} dP(s, Y_s) dZ_s \right\}. \quad (3.18)$$

Under the risk-neutral framework in Leung and Ludkovski (2011), the quadratic covariation also appears in the delayed purchase premium. In contrast, the current risk averse case involves an additional quadratic penalty term, and is nonlinear in quantity α . Finally, we remark that the stochastic control problems (3.13) for f and (3.14) and (3.18) for L all admit the same optimal control $(\tau^*, \tilde{\phi}^*)$.

From Theorem 2 or expression (3.13), we see that $f(t, y)$ is equivalent to indifference pricing of an American claim with payoff $h(\tau, Y_{\tau}) - \alpha P(\tau, Y_{\tau})$. We can then employ the analysis of indifference pricing for American options from Oberman and Zariphopoulou (2003) to derive the quasi-variational inequalities for $f(t, y)$ and $L(t, y)$.

Proposition 5. *The indifference value $f(t, y)$ is given by*

$$f(t, y) = \frac{-1}{\gamma(1 - \rho^2)} \log w(t, y),$$

where $w(t, y)$ is the unique bounded viscosity solution of the linear variational inequality (VI)

$$\min \left(w_t + \mathcal{L}^0 w, e^{-\gamma(1-\rho^2)(h-\alpha P)} - w \right) = 0, \quad (3.19)$$

with $w(T, y) = 1$. In turn, the delayed purchase premium $L(t, y)$ solves the semilinear VI:

$$\max \left(L_t + \mathcal{L}^{\phi^*} L - \frac{\gamma}{2} \hat{\rho}^2 c^2 L_y^2 + \hat{\rho} c P_y (\phi^* - \psi), -L \right) = 0, \quad (3.20)$$

with $L(T, y) = 0$.

Proof. The VI and associated existence-uniqueness for $f(t, y)$ follow from Theorem 7 of Oberman and Zariphopoulou (2003) with an American claim $h - \alpha P$. Then, we derive the VI for $L(t, y) = f(t, y) - h(t, y) + \alpha P(t, y)$ using the associated VI for $f(t, y)$, as well as the PDEs (3.5) and (3.12) for $h(t, y)$ and $P(t, y)$, respectively. Direct substitution yields VI (3.20). \square

The nonlinear payoff transformation $e^{-\gamma(1-\rho^2)(h(t,y)-\alpha P(t,y))}$, as well as the logarithmic transform from $f(t, y)$ to $w(t, y)$ precisely correspond to the risk-aversion effects. As $\gamma \rightarrow 0$, one obtains a linear VI for $f(t, y)$ itself (see Proposition 8 of Oberman and Zariphopoulou (2003)). Note that the VIs (3.20) and (3.19) yield the same purchase boundary for the investor. To solve either VI, one needs to first solve for the indifference price $h(t, y)$ and the market price $P(t, y)$. In some cases, both $h(t, y)$ and $P(t, y)$ admit closed-form formulas that facilitate the numerical implementation.

Proposition 5 also offers the opportunity to carry out comparative statics on the optimal purchase time and delayed purchase premium L . For instance, the market premium ψ only affects $P(t, y)$ in (3.19). If $\psi \mapsto P(t, y)$ is monotone, we obtain corresponding monotonicity in $L(t, y)$ and τ^* . The effect of other model parameters is more complicated. The risk-aversion γ , for example, affects both $\exp(\gamma \hat{\rho}^2 P(t, y))$ and $h(t, y)$. Note that in contrast to classical exponential utility cases (see, e.g., Theorem 3 of Musiela and Zariphopoulou (2004)), risk aversion γ and the correlation ρ are no longer coupled together, since ρ also has a direct influence on the diffusion Y under Q^{ϕ^*} .

3.2 Analysis of Purchase Strategies

In this section, we present several properties of the optimal purchase strategy. In particular, we explore the conditions under which immediate purchase or permanent delay is optimal.

To start, we notice that if the market price dominates the investor's indifference price, then it is never optimal to purchase the option from the market.

Lemma 6. *If $\alpha P(t, y) > h(t, y) \forall (t, y) \in [0, T) \times \mathbb{R}_+$, then it is never optimal to purchase the option, $\tau^* = T$. Moreover, $f(t, y) = 0$ and $L(t, y) = \alpha P(t, y) - h(t, y) > 0$.*

Proof. By direct substitution, one can verify that $w(t, y) = 1$ solves VI (3.19), and $L(t, y) = \alpha P(t, y) - h(t, y) > 0$ solves VI (3.20). Then, according to (2.17) the delayed purchase premium never reaches zero prior to expiration date T , so it is never optimal to purchase early. \square

For example, if the market price is always higher than the MEMM/MMM price corresponding to risk premium $\psi = 0$, then it must also dominate the indifference price $h(t, y)$ for any risk aversion level. By Lemma 6, the buyer will then never purchase the option.

More generally, we can study the sign of the integrand of $L(t, y)$ in (3.14) to deduce when the optimal strategy is trivial.

Theorem 7. Define the drift function

$$G(t, y) := \frac{1}{2\gamma} (\tilde{\phi}^*(t, y) - \phi^*(t, y))^2 + \alpha \hat{\rho} c(t, y) P_y(t, y) (\tilde{\phi}^*(t, y) - \psi(t, y)), \quad (3.21)$$

where $\tilde{\phi}^*(t, y) = \gamma \hat{\rho} c(t, y) f_y(t, y)$ is the minimizer in (3.13).

If $G(t, y) \geq 0 \forall (t, y)$, then it is never optimal to purchase. In this case, $f(t, y) = 0$.

If $G(t, y) \leq 0 \forall (t, y)$, then it is optimal to purchase immediately. In this case, $f(t, y) = h(t, y) - \alpha P(t, y)$.

Theorem 7 offers the counterpart of Theorems 3.1 and 4.2 in Leung and Ludkovski (2011) for the drift function G of risk-averse investors. It summarizes the interaction between the market and indifference prices and the optimal timing problem. As stated, (3.21) requires the knowledge of $\tilde{\phi}^*(t, y)$, or equivalently $f_y(t, y)$, in addition to the partial derivatives $h_y(t, y)$ and $P_y(t, y)$. However, there is a similar sufficient condition that does not involve $f_y(t, y)$. In (3.14), the integrand is quadratic in ϕ . Let us minimize the integrand over ϕ while fixing the expectation under an arbitrary measure Q^ϕ . If the resulting integrand is positive a.s. under some measure Q^ϕ , then it is also positive under the optimal measure $Q^{\tilde{\phi}^*}$, which means it is never optimal to purchase. This is the case if

$$g(t, y) := \gamma \hat{\rho}^2 c^2(t, y) P_y(t, y) \left(h_y(t, y) - \frac{P_y(t, y)}{2} \right) - \hat{\rho} c(t, y) P_y(t, y) \psi(t, y) \geq 0. \quad (3.22)$$

It is straightforward to show that $G(t, y) = g(t, y) + \frac{\gamma}{2} \hat{\rho}^2 c^2(t, y) L_y^2(t, y)$ and, therefore, the condition (3.22) implies that $G(t, y) \geq 0$.

Since the first term of the G function in (3.21) is non-negative, we infer the following result:

Corollary 8. If

$$\left(\tilde{\phi}^*(t, y) - \psi(t, y) \right) P_y(t, y) \geq 0, \quad \forall (t, y) \in [0, T] \times \mathbb{R}_+ \quad (3.23)$$

then it is never optimal to purchase the option.

For the most common options, such as Calls and Puts, the sign of P_y is constant. Therefore, checking the inequality (3.23) reduces to the direct comparison between the risk premia $\tilde{\phi}^*(t, y)$ and $\psi(t, y)$.

From (3.14) it is clear that if $G(t, y) > 0$ then the buyer should postpone purchase since an additional infinitesimal premium can be obtained by taking $\tau = t + \epsilon$ for ϵ sufficiently small in (3.14). Hence, for every (t, y) in the purchase region \mathcal{B} (including the purchase boundary), we must have $G(t, y) \leq 0$.

Furthermore, if two drift functions satisfy the dominance condition $G_1(t, y) \geq G_2(t, y)$ for all (t, y) , then the corresponding delayed purchase premia satisfy $L_1(t, y) \geq L_2(t, y)$. As a result, it is always optimal to purchase the derivative associated with G_2 before that associated with G_1 .

Applying the zero risk aversion limit in Proposition 3, $L^E(t, y) = P(t, y) - \inf_{t \leq \tau \leq T} \mathbb{E}_t^{Q^E} \{P(\tau, Y_\tau)\}$. Recall that $Q^E = Q^0$ is the MEMM corresponding to zero risk premium. Considering the SDE of $P(t, Y_t)$ under the measure Q^E , which amounts to setting $\phi = 0$ in (3.16), we obtain the probabilistic representation for $L^E(t, y)$:

$$L^E(t, y) = \sup_{t \leq \tau \leq T} \mathbb{E}_t^{Q^E} \left\{ \int_t^\tau \psi(s, Y_s) c(s, Y_s) \hat{\rho} P_y(s, Y_s) ds \right\}. \quad (3.24)$$

This can be viewed as a special case of the (risk-neutral) delayed purchase premium studied in Leung and Ludkovski (2011) where the investor's and the market pricing measures are Q^E and Q^ψ ,

respectively. Clearly, the delayed purchase premium is zero, $L^E = 0$, if $Q^E = Q^\psi$ or equivalently $\psi = 0$. In general, *price discrepancy* arises when the investor and market disagree on the risk-neutral pricing measure. The zero risk aversion limit serves as an example of how an investor may pick a pricing measure different from the market, $Q^E \neq Q^\psi$.

4 Examples

In this section we present two detailed examples illustrating the framework developed in Section 3.

4.1 Digital Options

The risk-averse investor's purchase timing requires the computation of the indifference price. In the current Markovian model, the indifference price can always be obtained by numerically solving the underlying PDE (3.5). In special cases, explicit computations are also possible. For instance, consider the purchase of a digital Call with Y a geometric Brownian motion, namely,

$$dY_t = bY_t dt + cY_t (\rho dW_t + \hat{\rho} d\hat{W}_t). \quad (4.1)$$

As is well known (see, e.g., Theorem 3 of Musiela and Zariphopoulou (2004)), the indifference price is $h(t, y) = -\frac{1}{\gamma(1-\rho^2)} \log \mathbb{E}_{t,y}^0 \{e^{-\gamma(1-\rho^2)D(Y_T)}\}$. For a digital Call with payoff $D(Y_T) = \mathbb{1}_{\{Y_T \geq K\}}$, the indifference price is explicitly given by the Black-Scholes type formula

$$\begin{aligned} h(t, y) &= -\frac{1}{\gamma(1-\rho^2)} \log \mathbb{E}^{Q^0} \{e^{-\gamma(1-\rho^2)\mathbb{1}_{\{Y_T \geq K\}}} \mid Y_t = y\} \\ &= -\frac{1}{\gamma(1-\rho^2)} \log \left(e^{-\gamma\rho^2} Q^0 \{Y_T \geq K \mid Y_t = y\} + Q^0 \{Y_T < K \mid Y_t = y\} \right) \\ &= -\frac{1}{\gamma(1-\rho^2)} \log \left(1 - \Phi(d_0)(1 - e^{-\gamma\rho^2}) \right), \end{aligned} \quad (4.2)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function and

$$d_0 \equiv d_0(t, y) := \frac{\log(y/K) + (b - \lambda\rho c - c^2/2)(T - t)}{c\sqrt{T - t}}.$$

Since $h(t, y) < Q^0(Y_T \geq K) = \Phi(d_0)$, to have a non-trivial purchase problem, the market must be assigning a larger risk premium compared to the MMM Q^0 . Namely, if the market price of risk for \hat{W} is a constant $\psi \geq 0$ then

$$P(t, y) = \Phi(d_\psi) \quad \text{with} \quad d_\psi := d_0 - \psi\hat{\rho}\sqrt{T - t}.$$

Given the above explicit expressions for $h(t, y)$ and $P(t, y)$, we can use the linearized VI (3.19) to compute the delayed exercise premium and the resulting purchase boundaries; recall that $w(t, y) \leq e^{-\gamma(1-\rho^2)(h(t,y)-P(t,y))} \wedge 1$ in (3.19).

The price spread $h(t, y) - P(t, y)$ changes signs twice, being negative at-the-money $y = K$ and positive when $|\log(y/K)|$ is large, see Figure 2. In the limit $\log(y/K) \rightarrow \pm\infty$, the spread is asymptotically zero. This complex shape ensures that the purchase region \mathcal{B} is non-trivial and in fact exhibits two purchase boundaries $Y_*(t) \leq K \leq Y^*(t)$, namely one should postpone purchase if the option is currently close to at-the-money. As time to maturity increases, the smaller market price of risk under Q^0 begins to dominate the impact of risk aversion, so that $t \mapsto h(t, y) - P(t, y)$ is decreasing. Therefore, for maturity long enough, $h > P$ everywhere and there is no reason to

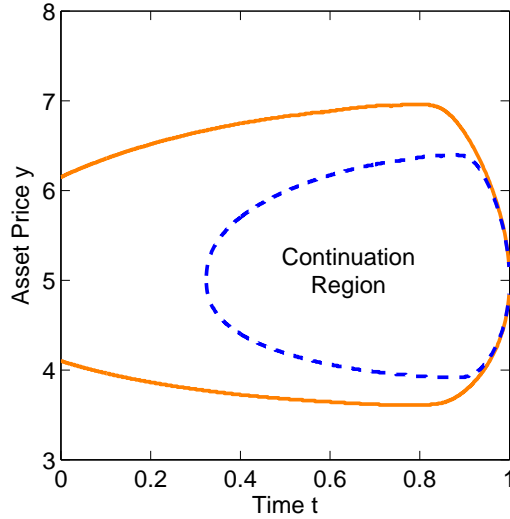


Figure 1: Optimal purchase boundaries for a digital Call. We take $K = 5$, $T = 1$ and $b = 0.1, c = 0.2, \psi = 0.125, \lambda = 0.4$ with $\rho = 0.9$. The plot shows the optimal purchase boundaries $(Y_*(t), Y^*(t))$ as a function of t for $\gamma = 1$ (solid) and $\gamma = 0.75$ (dashed). The continuation region is in the middle. Note that the continuation region is empty for $\gamma = 0.75$ and $t < 0.33$.

postpone purchases, i.e. $Y_*(t) = Y^*(t) = K$. Conversely, as $t \rightarrow T$, both h and P converge to the payoff $\mathbb{1}_{\{y \geq K\}}$ and their difference shrinks to zero. The trade-off between these two effects makes the purchase boundaries non-monotone in t , see Figure 1.

We note that in this example the effect of correlation parameter ρ is distinct from that of the risk-aversion γ , since ρ affects the risk premia spreads $\psi \hat{\rho} c$ and $\lambda \rho c$ separately from γ . Larger risk-aversion lowers the buyer's indifference price $h(t, y)$ and therefore reduces the spread $h(t, y) - P(t, y)$. As a result, the continuation region widens and $f(t, y)$ decreases, with purchases made closer to T . On the other hand, the impact of increasing ρ is *not* monotone on $h(t, y) - P(t, y)$ (as ρ increases, $w(t, y)$ increases at-the-money but decreases deep in-the-money/out-of-the-money), though we observe that $\rho \mapsto L(t, y)$ is still monotone and so is the effect on purchase boundaries.

The right panel in Figure 2 shows the behavior of purchase boundaries as we vary γ . As explained above, for any $\gamma > 0$ there exists $T^*(\gamma)$ such that the continuation region is empty for option time-to-maturity larger than T^* . As $\gamma \rightarrow 0$, the discount due to risk-aversion disappears, and in the limit $h(t, y) > P(t, y)$ for all t, y , so that $T^* \rightarrow 0$. Alternatively, for any fixed t , we find that $\gamma \mapsto Y_*(t)$ is decreasing (resp. $\gamma \mapsto Y^*(t)$ is increasing) and the continuation region widens in γ .

As $\gamma \rightarrow 0$, the investor's indifference price converges to

$$h(t, y) \xrightarrow{\gamma \rightarrow 0} \mathbb{E}^{Q^0} \{ \mathbb{1}_{\{Y_T \geq K\}} | Y_t = y \} = \Phi(d_0).$$

In other words, the investor prices the option with zero risk premium. The corresponding delayed purchase premium is given by (3.24). With a constant market risk premium ψ , the resulting timing problem must be trivial (either buy now or never). For instance, when $\psi > 0$, since $P_y(t, y) \geq 0$ it follows from the positivity of the integrand in (3.24) that $\tau_t^* = T$, for digital Calls with any strike or expiry. Hence, in this example, risk aversion adds a significant level of complexity to the timing decision.




Figure 2: Digital Call with $K = 5$ and $T = 1$. Parameters: $b = 0.1, c = 0.2, \psi = 0.125, \lambda = 0.4$ and default values of $\gamma = 1, \rho = 0.9$. Left panel: price spread $h(t, y) - P(t, y)$ (dashed lines) compared to the purchase indifference value $f(t, y)$ for $t = 0, 0.5, 0.9$. The crosses indicate the purchase boundaries $Y_*(t)$ and $Y^*(t)$. Right panel: the price spread $h(0, y) - P(0, y)$ (dashed lines) compared to the purchase indifference values $f(0, y)$ for $\gamma = 0.8, 1, 1.2$ (note that the continuation region is empty for $\gamma = 0.8$ whereby it is optimal to buy the digital Call immediately for any $Y_0 = y$).

4.2 Defaultable Bonds in a Structural Model

As another example, let us consider optimally timing to buy a defaultable bond that pays \$1 on expiration date T unless the underlying firm defaults. For simplicity, we assume zero recovery. The firm's stock price and its non-traded asset value are modeled by (S, Y) in (3.1) and (4.1). Similar to Section 4.1, the firm's stock is traded but its net asset value is not.

Following the structural default model introduced by Black and Cox (1976), the firm's default is signaled by Y hitting the default boundary $\beta_\epsilon(t) = \beta e^{-\epsilon(T-t)}$, where $\beta, \epsilon \geq 0$. Given that the firm survives through time $t \geq 0$, its default time ζ_t is given by

$$\zeta_t := \inf\{u \geq t : Y_u \leq \beta_\epsilon(u)\}. \quad (4.3)$$

The investor dynamically trades the firm's stock and the money market account over $[0, T]$. Prior to default time ζ , the trading wealth X , with strategy θ , evolves according to (3.3). After default, the firm's stock is no longer tradable, so the investor liquidates holdings in the stock and deposits in the money market account (with zero interest rate). Following Leung et al. (2008) and Sircar and Zariphopoulou (2010), we assume a full pre-default market value on stock holdings upon liquidation. Hence, on $\{\zeta_t \leq T\}$, the wealth is $X_u = X_{\zeta_t}$ for $u \in (\zeta_t, T]$.

When the investor takes a long position on $\alpha \geq 0$ defaultable bonds, she faces the utility maximization problem:

$$V(t, x, y; \alpha) = \sup_{\theta \in \Theta_{t,T}} \mathbb{E} \left\{ U(X_T^\theta + \alpha \mathbb{1}_{\{\zeta_t > T\}}) \mid X_t = x, Y_t = y \right\}.$$

As shown by Leung et al. (2008), the value function is given in closed-form. Let

$$d \equiv d(t, y) = \frac{\log(\beta/y) - \epsilon(T-t)}{c} \quad \text{and}$$

$$a = (1 - \rho^2) \frac{\lambda^2}{2}, \quad k_1 = \frac{b - \epsilon}{c} - \rho\lambda - \frac{c}{2}, \quad k_2 = \sqrt{k_1^2 + 2a}.$$

Theorem 9. *The value function admits the separation of variables:*

$$V(t, x, y; \alpha) = -e^{-\gamma(x+\alpha)} v(t, y; \alpha)^{\frac{1}{1-\rho^2}} \quad (4.4)$$

on $\{(t, x, y) \in [0, T] \times \mathbb{R} \times \mathbb{R}_+ : y \geq \beta_\epsilon(t)\}$ where

$$v(t, y; \alpha) = \mathbb{E}^{Q^0} \left\{ e^{-a(T-t)} \mathbb{1}_{\{\zeta_t > T\}} | Y_t = y \right\} + e^{\alpha\gamma(1-\rho^2)} \mathbb{E}^{Q^0} \left\{ e^{-a(\zeta_t-t)} \mathbb{1}_{\{\zeta_t \leq T\}} | Y_t = y \right\} \quad (4.5)$$

$$= A(t, y) + e^{\alpha\gamma(1-\rho^2)} B(t, y), \quad (4.6)$$

where Q^0 is the minimal martingale measure and

$$A(t, y) = e^{-a(T-t)} \left[\Phi \left(\frac{-d + k_1(T-t)}{\sqrt{T-t}} \right) - e^{2dk_1} \Phi \left(\frac{d + k_1(T-t)}{\sqrt{T-t}} \right) \right], \quad (4.7)$$

$$B(t, y) = \Phi \left(\frac{d - k_2(T-t)}{\sqrt{T-t}} \right) + e^{2dk_2} \Phi \left(\frac{d + k_2(T-t)}{\sqrt{T-t}} \right). \quad (4.8)$$

Note that $v(t, y; \alpha)$ depends on y only through the default-to-asset ratio β/y . With formula (4.4), the investor's indifference price $h(t, y; \alpha)$ for holding $\alpha \geq 0$ units of the defaultable bond is given by

$$h(t, y; \alpha) = \alpha - \frac{1}{\gamma(1-\rho^2)} \log \frac{v(t, y; \alpha)}{v(t, y; 0)} \leq \alpha. \quad (4.9)$$

On the other hand, the market price is given by a risk-neutral expectation under the market pricing measure Q^* , namely,

$$P(t, y) = \mathbb{E}^{Q^*} \{ \mathbb{1}_{\{\zeta_t > T\}} | Y_t = y \} = Q^* \{ \zeta_t > T | Y_t = y \}.$$

In particular, if the market risk premium for \hat{W} is a constant ψ , then the corresponding price can be computed explicitly:

$$P(t, y) = \Phi \left(\frac{-d + k_3(T-t)}{\sqrt{T-t}} \right) - e^{2dk_3} \Phi \left(\frac{d + k_3(T-t)}{\sqrt{T-t}} \right), \quad (4.10)$$

where $k_3 = \frac{b-\epsilon}{c} - \rho\lambda - \hat{\rho}\psi - \frac{c}{2}$.

Using the payoff function $h(t, y; \alpha) - \alpha P(t, y)$, which is explicitly given in (4.9) and (4.10), we find that it is optimal to postpone bond purchase when close to the default level and buy immediately if the default/asset ratio β/y is sufficiently small. Moreover, the spread $h(t, y; \alpha) - \alpha P(t, y)$ grows in time-to-maturity $T-t$. Figure 3 shows that the resulting purchase boundary $t \mapsto \bar{Y}(t)$ with $\tau^* = \inf\{t \geq 0 : Y_t \leq \bar{Y}(t)\}$ is monotone in t . Numerical experiments also indicate that larger risk-aversion γ widens the continuation region, as does the volume α of bonds to purchase (cf. discussion in Section 3.1).

In this situation, as $\gamma \rightarrow 0$, the buyer's purchase problem remains non-trivial. We have

$$h(t, Y; \alpha) \xrightarrow{\gamma \rightarrow 0} \alpha e^{-a(T-t)} \frac{A(t, y)}{A(t, y) + B(t, y)},$$

where A and B are defined in (4.7) and (4.8), respectively. The right-hand-side above may be bigger or smaller than $\alpha P(t, y)$ depending on the value of y . Hence, even with $\gamma = 0$ there is a purchase boundary and it is optimal to buy the defaultable bond only if the default-to-asset ratio is sufficiently low.




Figure 3: Defaultable bond with $\beta = 5, \epsilon = 0$ and $T = 1$. Parameters: $b = 0.08, c = 0.2, \psi = 0, \lambda = 0.3$ and default values of $\gamma = 1, \rho = 0.8$. Left panel: the price spread $[h(t, y; 1) - P(t, y)]$ (dashed lines) versus the purchase indifference value $f(t, y)$ (solid lines) for $t = 0, 0.5, 0.9$ and $\alpha = 1$. The crosses indicate the purchase boundary $\bar{Y}(t)$. Right panel: the purchase boundary $\bar{Y}(t)$ for $\alpha = 1, 2, 3$.

5 Optimal Sequential Option Purchase

If the investor is contemplating buying more than one option, she has the opportunity to spread her trades over time. To illustrate, suppose that the investor needs to purchase several identical option contracts from the market prior to some pre-specified date T_1 . At any moment, the investor can buy one or more contracts by paying the current market price, or wait for a later time. In this section we address the question of her optimal purchasing schedule.

To fix ideas, we again consider a derivative D expiring at time $T \geq T_1$, and its market price P as defined in (2.5). The objective is to optimally accumulate a *pre-specified* n units of D at or before $T_1 \leq T$. For any $t \in [0, T_1]$ and $i \leq n$, we denote by $\tau_t^{(i)*}$ the optimal purchase time of the next contract when $n - i$ units have already been bought at time t . At each purchase time $\tau_t^{(i)*}$, the investor withdraws $P_{\tau_t^{(i)*}}$ to buy another contract, leaving $(i - 1)$ units to buy afterwards. It is possible that multiple units get purchased simultaneously, whereby the corresponding purchase times coincide. As required, at date T_1 the investor will hold n units of option D .

If the investor has purchased all n contracts already and has wealth X_t at time t , then her indirect utility is $J_t^{(0)}(X_t) = V_t(X_t; n)$. Subsequently, if there are $i \in \{1, \dots, n\}$ units of D remaining to be bought, then the investor's indirect utility is given recursively by

$$J_t^{(i)}(X_t) = \text{ess sup}_{\tau_i \in \mathcal{T}_t, T_1} \text{ess sup}_{\theta \in \Theta_t, \tau_i} \mathbb{E}_t \left\{ J_{\tau_i}^{(i-1)}(X_{\tau_i}^\theta - P_{\tau_i}) \right\}. \quad (5.1)$$

For any $t \in [0, T_1]$, the investor's indifference value, $f_t^{(i)}$, for optimally timing to buy the remaining $i \in \{1, \dots, n\}$ units of D is determined from the equation $J_t^{(i)}(X_t) =: M_t(X_t + f_t^{(i)})$, with $f_t^{(0)} = h_t(n)$. Substituting the definition of $f^{(i)}$ into (5.1), we obtain

$$J_t^{(i)}(X_t) = \text{ess sup}_{\tau_i \in \mathcal{T}_t, T_1} \text{ess sup}_{\theta \in \Theta_t, \tau_i} \mathbb{E}_t \left\{ M_{\tau_i} \left(X_{\tau_i} + f_{\tau_i}^{(i-1)} - P_{\tau_i} \right) \right\}. \quad (5.2)$$

Expression (5.2) reveals the similarity between $J^{(i)}$ and J in (2.6). Indeed, the investor with i remaining options to buy can determine the next optimal purchase time by considering the problem of optimally exercising a claim with payoff $f_t^{(i-1)} - P_t$. Accordingly, we can derive the dual for $J_t^{(i)}(X_t)$ and the indifference price $f_t^{(i)}$ by just replacing $h_\tau - \alpha P_\tau$ by $f_\tau^{(i-1)} - P_\tau$ in Theorem 2.

Proposition 10. *The values $(f_t^{(0)}, f_t^{(1)}, \dots, f_t^{(n)})$ can be expressed as $f_t^{(0)} = h_t(n)$, and*

$$f_t^{(i)} = \operatorname{ess\,sup}_{\tau^{(i)} \in \mathcal{T}_{t, T_1}} \operatorname{ess\,inf}_{Q \in \mathcal{P}_f} \left(\mathbb{E}_t^Q \{ f_{\tau^{(i)}}^{(i-1)} - P_{\tau^{(i)}} \} + \frac{1}{\gamma} H_t^{(i)}(Q|Q^E) \right), \quad \text{for } i = 1, 2, \dots, n. \quad (5.3)$$

Since $\tau = t$ is a candidate purchase time, we deduce from (5.2) that $M_t(X_t + f_t^{(i-1)} - P_t) \leq J_t^{(i)}(X_t) = M_t(X_t + f_t^{(i)})$, which in turn implies that

$$f_t^{(i-1)} - f_t^{(i)} \leq P_t. \quad (5.4)$$

The difference $f_t^{(i-1)} - f_t^{(i)}$ can be interpreted as the *marginal indifference value* for timing to buy the next option when i options remain to be purchased. From (5.4), we see that this marginal value is always less than or equal to the market price.

In order to quantify the benefit of optimally waiting to buy, we define the delayed purchase premium $L_t^{(i)}$ by comparing two scenarios faced by the investor who wants to buy i options: a) buy all i units of option D now by paying the prevailing market price iP_t , and b) pay $L_t^{(i)}$ now for the right to optimally wait to buy i unit of D over time. This leads to the indifference equation:

$$J_t^{(i)}(X_t - L_t^{(i)}) = V_t(X_t - iP_t; n). \quad (5.5)$$

Using (2.6) and (5.1), we infer that $J_t^{(i)}(X_t) \geq J_t^{(i-1)}(X_t - P_t) \geq \dots \geq V_t(X_t - iP_t; n)$, which implies that $L_t^{(i)} \geq 0$. In particular, these inequalities become equalities at time T_1 , and therefore, $L_{T_1} = 0$, meaning the delayed purchase premium vanishes at T_1 .

Using the definition of $f^{(i)}$ and (5.5), we can again decompose the investor's indifference price $f_t^{(i)}$ into three parts as in (2.10), namely,

$$f_t^{(i)} = h_t(n) - iP_t + L_t^{(i)}.$$

Hence, by subtraction, the marginal indifference price $f_t^{(i-1)} - f_t^{(i)}$ is given by

$$f_t^{(i-1)} - f_t^{(i)} = P_t - (L_t^{(i)} - L_t^{(i-1)}). \quad (5.6)$$

Note that the difference $L_t^{(i)} - L_t^{(i-1)}$ is non-negative by (5.4). It can be viewed as the *marginal* delayed purchase premium for the next contract when there are i units left to buy. Also, the investor's optimal time to buy the next contracts, with i units of D left to buy, is given by

$$\tau_t^{(i)*} = \inf \left\{ t \leq u \leq T_1 : f_u^{(i-1)} - f_u^{(i)} = P_u \right\} = \inf \left\{ t \leq u \leq T_1 : L_u^{(i)} - L_u^{(i-1)} = 0 \right\}. \quad (5.7)$$

As a result, the investor will purchase the next contract as soon as the marginal delayed purchase premium decreases to zero, and then repeat the same strategy for subsequent units until time T_1 .

By the dual representations of $f^{(i)}$ in Proposition 10, we can formally derive the large and zero risk aversion limits following the proof of Proposition 3. As risk aversion decreases to zero, we have

$$\lim_{\gamma \rightarrow 0} f_t^{(i)} = nh_t^E - i \cdot \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{t, T}} \mathbb{E}_t^{Q^E} \{ P_\tau \}, \quad i = 0, 1, \dots, n. \quad (5.8)$$

In turn, the marginal delayed purchase premium reverts back to the familiar form (see (2.20)):

$$\lim_{\gamma \rightarrow 0} (L_t^{(i)} - L_t^{(i-1)}) = P_t - \operatorname{ess\,inf}_{\tau \in T_t, T} \mathbb{E}_t^{Q^E} \{P_\tau\}, \quad i = 1, 2, \dots, n. \quad (5.9)$$

As is intuitive, when the investor is not risk averse, then the indifference value is linear in quantity and all contracts get purchased simultaneously. Therefore, the incorporation of risk aversion leads to significantly different timing decision in the purchase of multiple options.

5.1 Numerical Example

The optimal purchase timing schedule can be obtained by numerically solving a chain of variational inequalities. To illustrate, let us consider the correlated geometric Brownian motion model (S, Y) in (3.1) and (4.1). The purchase indifference value $f^{(i)}(t, y)$ is given by

$$f^{(i)}(t, y) = \frac{-1}{\gamma(1 - \rho^2)} \log w^{(i)}(t, y),$$

where $w^{(i)}(t, y)$, $(t, y) \in [0, T_1] \times \mathbb{R}_+$, is the unique viscosity solution of the linear variational inequality

$$\min \left(w_t^{(i)} + \mathcal{L}^0 w^{(i)}, w^{(i-1)} \cdot e^{\gamma(1-\rho^2)P} - w^{(i)} \right) = 0, \quad (5.10)$$

with $w^{(i)}(T_1, y) = \exp(-\gamma \hat{\rho}^2 \{h(T_1, y; n) - (n - i)P(T_1, y)\})$ and \mathcal{L}^0 defined in (3.6).




Figure 4: Optimal purchase times on a sample path. We consider purchasing $n = 3$ digital Calls with strike $K = 5$ and maturity $T = 1$ until the deadline $T_1 = 0.75$. The other parameters are $b = 0.1$, $c = 0.2$, $\psi = 0.025$, $\gamma = 1$ and $\rho = \sqrt{0.75}$. We show a sample path of (Y_t) , starting from $Y_0 = 5$ and the three corresponding purchase dates $\tau^{(3)*} \simeq 0.298$, $\tau^{(2)*} \simeq 0.687$, $\tau^{(1)*} \simeq 0.688$, indicated with the diamonds.

The optimal purchase dates are given by

$$\tau^{(i)*} = \inf \left\{ 0 \leq t \leq T_1 : f^{(i-1)}(t, Y_t) - f^{(i)}(t, Y_t) = P(t, Y_t) \right\}, \quad i = 1, 2, \dots, n.$$

Figure 4 illustrates the resulting solution in case of the digital Call of Section 4.1. We again observe two-sided purchase boundaries, cf. Figure 1, and the nesting of the continuation regions which cause the sequential purchases to be spread over time.

6 Extensions and Concluding Remarks

The presented framework for analyzing the timing flexibility in derivative trading is amenable to multiple extensions. Among others, it is natural to consider the optimal single/sequential *buy-and-sell* strategy. In the simplest case, the investor will optimally time to buy, say α units of option D , and then sell it at the market price at or before expiration, with the goal of generating profit while accounting for risk-aversion. She then faces the combined optimal control and stopping problem:

$$\widehat{J}_t(X_t; \alpha) := \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,sup}_{\theta \in \Theta_{t,\tau}} \mathbb{E}_t \left\{ \widehat{V}_\tau(X_\tau^\theta - \alpha P_\tau; \alpha) \right\}, \quad (6.1)$$

where nested is the optimal liquidation problem after purchase

$$\widehat{V}_t(X_t; \alpha) := \operatorname{ess\,sup}_{\nu \in \mathcal{T}_{t,T}} \operatorname{ess\,sup}_{\theta \in \Theta_{t,\nu}} \mathbb{E}_t \left\{ M_\nu(X_\nu^\theta + \alpha P_\nu) \right\}. \quad (6.2)$$

Following the indifference pricing arguments as above, it turns out that the investor's indifference price is the sum of the delayed purchase premium and the delayed liquidation premium; see also (Leung and Ludkovski, 2011, Sec. 5.2) for the treatment of the risk-neutral case. For this extended problem, tractable representations may be straightforwardly obtained using the above methods under exponential utility.

We can also treat many other parametric models, including stochastic volatility models (Sircar and Zariphopoulou, 2005) and credit risk models (Leung et al., 2008; Jaimungal and Sigloch, 2010) where closed-form expressions and dual representations are available for exponential indifference prices. It will certainly be interesting and challenging to consider derivative trading under other risk preferences. We may also consider purchase of American derivatives, which will lead to a compound timing option.

Finally, in our main model we assumed that the investor is risk averse while the market prices in a risk-neutral way via risk premium specification. In principle and in practice, it is possible that both buyers and sellers are risk averse, especially in the over-the-counter market, so indifference pricing mechanisms may apply for both parties. However, at least in the case of exponential utility, trading will be precluded if the buyers and sellers agree on the historical measure \mathbb{P} , even though they may have different risk aversion coefficients. Indeed, one can show that the buyer's indifference price h_t and the seller's indifference price h_t^S are respectively monotonically decreasing and increasing in γ , with the same zero risk-aversion limit h_t^E (priced under the MEMM). This results in the price domination $h_t \leq h_t^E \leq h_t^S$ for all t a.s., leading to no purchase of derivatives. Other non-trivial trade/no-trade conditions may arise in markets with buyers and sellers with different families of utilities, or with heterogeneity in market view, see Leung (2011).

References

Becherer, D. (2001). Rational hedging and valuation with utility-based preferences. Doctoral dissertation, Technical University of Berlin.

- Becherer, D. (2003). Rational hedging and valuation of integrated risks under constant absolute risk aversion. *Insurance: Mathematics and Economics*, 3:1–28.
- Black, F. and Cox, J. (1976). Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance*, 31:351–367.
- Carmona, R., editor (2008). *Indifference Pricing: Theory and Applications*. Princeton University Press, Princeton, NJ.
- Davis, M. (1997). Option pricing in incomplete markets. In Dempster, M. and Pliska, S., editors, *Mathematics of Derivatives Securities*, pages 227–254. Cambridge University Press.
- Davis, M. (2001). Pricing weather derivatives by marginal value. *Quantitative Finance*, 1:1–4.
- Davis, M. (2006). Optimal hedging with basis risk. In Kabanov, Y., Lipster, R., and Stoyanov, J., editors, *From Stochastic Calculus to Mathematical Finance : The Shiryaev Festschrift*, pages 169–188. Springer, Berlin.
- Delbaen, F. (2006). The structure of mstable sets and in particular of the set of risk neutral measures. In *Lecture Notes in Mathematics*, volume 1874, pages 215–258. Springer, Berlin.
- Delbaen, F., Grandits, P., Rheinländer, T., Samperi, D., Schweizer, M., and Stricker, C. (2002). Exponential hedging and entropic penalties. *Mathematical Finance*, 12:99–123.
- Föllmer, H. and Schweizer, M. (1990). Hedging of contingent claims under incomplete information. In Davis, M. and Elliot, R., editors, *Applied Stochastic Analysis, Stochastics Monographs*, volume 5, pages 389 – 414. Gordon and Breach, London/New York.
- Frittelli, M. (2000). The minimal entropy martingale measure and the valuation problem in incomplete markets. *Mathematical Finance*, 10:39–52.
- Henderson, V. (2005). The impact of the market portfolio on the valuation, incentives and optimality of executive stock options. *Quantitative Finance*, 5:1–13.
- Henderson, V. (2007). Valuing the option to invest in an incomplete market. *Mathematics and Financial Economics*, 1(2):103–128.
- Hodges, S. and Neuberger, A. (1989). Optimal replication of contingent claims under transaction costs. *Review of Futures Markets*, 8:222–239.
- İlhan, A., Jonsson, M., and Sircar, R. (2005). Optimal investment with derivative securities. *Finance and Stochastics*, 9(4):585–595.
- Jaimungal, S. and Sigloch, G. (2010). Incorporating risk and ambiguity aversion into a hybrid model of default. *Mathematical Finance*. To appear.
- Kabanov, Y. and Stricker, C. (2002). On the optimal portfolio for the exponential utility maximization: Remarks to the six-author paper. *Mathematical Finance*, 12:125–134.
- Kramkov, D. and Bank, P. (2010). A model for a large investor trading at market indifference prices II: continuous-period case. preprint.
- Leung, T. (2011). Explicit solutions to optimal risk-averse trading of defaultable bonds under heterogeneous beliefs. Working paper.
- Leung, T. and Ludkovski, M. (2011). Optimal timing to purchase options. *SIAM Journal on Financial Mathematics*. forthcoming.
- Leung, T. and Sircar, R. (2009a). Accounting for risk aversion, vesting, job termination risk and multiple exercises in valuation of employee stock options. *Mathematical Finance*, 19(1):99–128.

- Leung, T. and Sircar, R. (2009b). Exponential hedging with optimal stopping and application to ESO valuation. *SIAM Journal of Control and Optimization*, 48(3):1422–1451.
- Leung, T., Sircar, R., and Zariphopoulou, T. (2008). Credit derivatives and risk aversion. In Fomby, T., Fouque, J.-P., and Solna, K., editors, *Advances in Econometrics*, volume 22, pages 275 – 291. Elsevier Science.
- Musiela, M. and Zariphopoulou, T. (2004). An example of indifference pricing under exponential preferences. *Finance and Stochastics*, 8:229–239.
- Oberman, A. and Zariphopoulou, T. (2003). Pricing early exercise contracts in incomplete markets. *Computational Management Science*, 1:75–107.
- Riedel, F. (2009). Optimal stopping with multiple priors. *Econometrica*, 77(3):857 – 908.
- Rouge, R. and El Karoui, N. (2000). Pricing via utility maximization and entropy. *Mathematical Finance*, 10(2):259–276.
- Sircar, R. and Zariphopoulou, T. (2005). Bounds and asymptotic approximations for utility prices when volatility is random. *SIAM Journal on Control and Optimization*, 43(4):1328–1353.
- Sircar, R. and Zariphopoulou, T. (2010). Utility valuation of multiname credit derivatives and application to CDOs. *Quantitative Finance*, 10(2):195–208.