



This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Horizon-unbiased utility functions

Vicky Henderson^a, David Hobson^{b,*}

^a *Warwick Business School, University of Warwick, Coventry, CV4 7AL, UK*

^b *Department of Statistics, University of Warwick, Coventry, CV4 7AL, UK*

Received 22 May 2006; received in revised form 12 March 2007; accepted 15 March 2007

Available online 2 June 2007

Abstract

In this paper we consider a class of mixed optimal control/optimal stopping problems related to the choice of the best time to sell a single unit of an indivisible asset. We assume that in addition to the indivisible asset, the agent has access to a financial market. Investments in the financial market can be used for hedging, but the financial assets are only partially correlated with the indivisible asset, so that the agent faces an incomplete markets problem.

We show how, even in the infinite horizon case, it is possible to express the problem as a maximisation problem with respect to an inter-temporal utility function evaluated at the sale time, but that this objective function must satisfy consistency conditions over time.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Optimal stopping; Stochastic control; Utility maximisation; Real options; Incomplete market; CRRA utility; Horizon-unbiased utility; Backward heat equation

1. Introduction and motivation

Consider the following problem. An agent has a single, indivisible unit of an asset to sell, and the decision to sell is irreversible. At the moment of sale, chosen by the agent, she receives a one-off, lump-sum payment. Her aim is to maximise the expected utility of wealth on exercise, where the objective function depends on both her wealth (including the revenue from the asset sale) and on time. The primary issue is to determine the optimal time to sell this asset, and the purpose of this paper is to argue that in order to have a mathematical problem which is consistent with

* Corresponding author.

E-mail addresses: Vicky.Henderson@wbs.ac.uk (V. Henderson), D.Hobson@warwick.ac.uk (D. Hobson).

the desired economic interpretation, the time dependence of the inter-temporal utility function cannot be arbitrary, and must satisfy certain consistency conditions.

In particular, we want to consider the sense in which we can write her mixed control/stopping problem as

$$\sup_{\tau} \sup_{\theta} \mathbb{E}U(\tau, W_{\tau}^{\theta} + Y_{\tau}) \quad (1)$$

for some inter-temporal utility or objective function $U(t, x)$. Here, τ is chosen from a suitable class of stopping times; W , parameterised by θ , is an element of the set of admissible wealth processes for the agent; and Y_{τ} is the one-off lump-sum payment received at τ from the asset sale.

Suppose first that the planning horizon over which we consider (1) is finite (for example, the set of allowed τ satisfy $\tau \leq T$ for some $T < \infty$). The idea is to show that if the ultimate objective of the agent is to maximise expected utility at T , where utility is specified via a numerical representation \bar{U} of preferences at time T , then this forces a particular choice of U . This choice corresponds to the implied or indirect utility function. In particular, we cannot write down an arbitrary specification for U in (1), but instead $U(t, x)$ must satisfy a dynamic consistency condition, for otherwise biases are introduced into the problem. This idea is entirely natural, and has been used previously in specific settings by Davis and Zariphopoulou [4] for pricing American options under transaction costs, and Oberman and Zariphopoulou [17] for pricing finite horizon American options in an incomplete market.

In this sense, the main novelty in our study arises from consideration of (1) over an infinite horizon, and the second aim of this paper is to show how the ideas which underpin the finite horizon problem can be extended to the infinite horizon case. In this case, it is not possible to define an indirect utility via preferences specified at a terminal horizon, and an alternative criterion is needed. Instead we can postulate a problem of the form (1), but again we are forced to choose U such that it satisfies a dynamic consistency condition. This condition can only be satisfied for certain families of objective functions, and even in these cases, the choice of discount parameter is determined by the other parameters in the problem.

The fundamental modelling principle which drives our conclusions is the idea that in the absence of Y , the agent solving (1) should have no preference over the choice of τ . Thus, when the asset Y is introduced, we can be sure that the choice of optimising τ is motivated by the existence of the right to sell, and not by any prior biases in the set-up of the problem.

In mathematical terms, our problem is one of mixed control/optimal stopping. Such mixed problems have occurred previously in mathematical finance in the work of Davis and Zariphopoulou [4], Karatzas and Kou [13] and Karatzas and Wang [14], among others. Further, and as described in the previous paragraph, one of our principles is that underlying the set-up is an optimisation problem which is robust over choices of stopping times. Although the approach is completely different, in this sense there is a link to the work of Choulli and Stricker [1] on the minimal entropy-Hellinger (MEH) martingale measure. Choulli and Stricker show that the MEH measure has the property that it is robust over stopping.

In financial terms, the basic asset sale problem falls into the class of problems considered in real options, see Dixit and Pindyck [5]. Typically in real options, investment decisions are interpreted as American call options where the exercise time of the option is the time of investment. The canonical models consider the investment decision over an infinite time horizon, since this is a reasonable assumption in practice, and because it leads to tractable solutions. The standard assumption made in real options (Dixit and Pindyck [5], McDonald and Siegel [15]) is

that the real asset is traded, or perfectly correlated with a traded asset, resulting in a complete market problem. In contrast, we consider this problem in an incomplete market, since we do not assume the payoff from the real asset to be a traded variable.

The papers of Henderson [8] and Evans et al. [6] considered versions of our real asset sale problem and motivated this study of horizon-unbiased utilities. Both treat the infinite horizon (or perpetual) problem so, as we will see, the formulation of such utilities is important. Henderson [8] considered the investment timing problem where a lump-sum investment payoff is received for an investment cost. The investment cost represents the strike of the call option. She solved the problem in closed-form for exponential utility. For this choice, wealth factors out, reducing the number of variables by one. Evans et al. [6] treat power or CRRA utility, for which the solution to the problem depends on wealth. Issues relating to these works are discussed further in Section 6.1.

In the next section, we introduce the simple model which motivated our study. We follow this section by a considering the features of the set-up which are required in general. We then consider the problem with a fixed and finite terminal horizon, and then in Section 5 we use the same ideas to guide our analysis of the perpetual problem. A final section contains a more detailed discussion of the applications of the results of the paper, and some general conclusions.

2. Horizon-unbiased utilities: A preview

2.1. The non-traded assets model

In later sections, we will work in a more general setting, but for the present we will restrict our attention to the asset sale problem in a simple market model. To make the problem both more interesting and more realistic, we assume that the asset does not exist in isolation, and there are other investment and hedging opportunities available to the agent. These opportunities are represented by a financial market including both a riskless bank account and risky asset. In contrast to the indivisible real asset, the financial market is characterised by the fact that assets are infinitely divisible, and asset sales are reversible. The financial market is complete.

As a concrete example, suppose Y follows an exponential Brownian motion

$$\frac{dY_t}{Y_t} = \sigma dB_t^Y + \mu dt \quad Y_0 = y;$$

and suppose that the (frictionless) financial market consists of a riskless bond paying rate of interest r , and a risky asset with price process P given by

$$\frac{dP_t}{P_t} = \eta dB_t^P + \nu dt \tag{2}$$

where $dB_t^Y dB_t^P = \rho dt$, for $\rho^2 < 1$. Denote by $\lambda = (\nu - r)/\eta$ the instantaneous Sharpe ratio of the risky asset. All the parameters are assumed to be constants. Note that although we assume there is a single traded asset in the financial market, this is for notational convenience, and the theory extends directly to the complete market case with n assets.

This model is often called the non-traded assets model (see the survey of Henderson and Hobson [10]), or the model with basis risk (Davis [2]). It has been used by Davis [3], Henderson and Hobson [9], Henderson [7], and others in the context of the pricing and hedging of options in incomplete markets via utility indifference pricing, where the main source of incompleteness is the fact that some assets are not liquidly traded.

The canonical model has several key features which simplify calculations — the underlying assets are Markovian; the financial market consisting of the bond and risky asset alone is in itself complete; there is a unique minimal-distance equivalent martingale measure in the larger model containing the pair of risky assets Y and P . The key feature for our purposes is that the underlying optimal investment problem in the financial market is simple to solve, and the solution depends only on the current wealth of the agent. Such optimal investment problems were first solved by Merton [16].

2.2. Power utilities and no preferred horizon

The agent's problem is to find the optimal time τ at which to sell one unit of Y , and the optimal investment strategy in risky asset P . The random time τ in (1) belongs to some prescribed set \mathcal{T} of stopping times — for example, the set of all (possibly infinite) stopping times, or the set of stopping times bounded by a finite horizon T . However we are particularly interested in the infinite horizon case, since this is a standard assumption in the real options applications that are our main focus.

The trading strategies are parameterised by θ . Under the (very reasonable, but see Section 6.3 for extensions) assumption that trading strategies are self-financing, for $t < \tau$, the agent's wealth W_t^θ follows

$$dW_t^\theta = \theta_t dP_t + r(W_t^\theta - \theta_t P_t) dt.$$

We are now in a position to describe what we mean by an appropriate choice of function $U(t, x)$. Often in the literature, utility functions are adjusted by an arbitrary subjective discount factor.

Example 2.1. Constant relative risk aversion (CRRA) preferences can be represented by

$$U(t, x) = e^{-\delta t} \frac{x^{1-R}}{1-R}; \quad R > 0, R \neq 1 \tag{3}$$

where δ is a subjective discount factor.

We will argue that for problems such as (1), this subjective discount factor cannot be arbitrary (and, in general, $U(t, x)$ cannot be chosen in an arbitrary fashion) in order for the problem to be internally consistent. More precisely, we say the problem (1) has *no preferred horizon* if the solution of

$$\sup_{\theta} \mathbb{E}[U(\tau, W_\tau^\theta)] \tag{4}$$

does not depend on τ . This requires the agent to be indifferent to the choice of horizon in the underlying optimal investment problem. That is, when the agent faces the problem without the sale of the real asset Y (or equivalently, if $Y_t \equiv 0$), she should not have any preference for one horizon over another. Choosing such a function $U(t, x)$ ensures that when the original problem (1) is considered, conclusions about the optimal sale time τ are not influenced by any in-built incentive for the agent to prefer early or late horizons. From a mathematical standpoint, the problem in (1) for U of the form (3) makes perfect sense for any δ , but if U is such that there is a preferred horizon, then there are distortions to the sale timing choice arising from the underlying investment problem.

As we shall see, in the model of Section 2.1 with CRRA utility, the requirement that the problem has no preferred horizon forces the choice $\delta = \beta$, where

$$\beta = \beta(R) = (1 - R)r + (1 - R)\lambda^2/2R.$$

For problems where the stopping time must be chosen to be smaller than some finite expiry date T , we can make the following definition:

Definition 2.2. The utility function $U(t, x)$ of the CRRA family is horizon-unbiased for the non-traded assets model if and only if $U(t, x) = e^{-\beta t} x^{1-R}/1 - R$.

Remark 2.3. Consider the optimal stopping/control problem (1) with U as given in Definition 2.2. At first sight, it appears that the problem which is being solved is one where all wealth is being liquidated at time τ , and then this wealth is evaluated using a utility function which depends on market parameters. However, this is not the most useful interpretation. Instead, τ should be considered as a time at which the nature of the problem facing the agent changes (thereafter she solves a standard optimal investment problem). Furthermore, U does not represent the preferences of the agent directly, but rather represents the induced utility function which arises when the preferences are combined with the solution of the optimal investment problem. The market parameters enter through this second element. The agent should act *as if* all wealth is liquidated at τ and preferences depend on market parameters, but this objective function reflects the optimal behaviour she will follow after τ .

Remark 2.4. The discount factor β has the following interpretation. The first component $(1 - R)r$ discounts future wealth into current values. The second component $(1 - R)\lambda^2/2R$ is chosen to exactly compensate for the opportunity cost of holding onto the real asset Y . This opportunity cost is related to the fact that the money realized from the sale of Y can be invested (to solve the Merton-style investment problem) in the financial asset.

Remark 2.5. Note that even in the situation without the risky financial asset (so $\lambda \equiv 0$), we have $\beta = (1 - R)r$. We still require that the discount factor is not arbitrary, but takes this precise form to ensure there are no biases in the conclusions on sale timing. If the discount factor were taken to be arbitrary, conclusions about the sale timing would be primarily driven by this discount factor. In their model of real options, Kadam et al. [12] allow for a subjective discount factor and this discount factor is the main determinant of the decision to sell. There is no hedging asset in the model of Kadam et al. [12], but one sensible interpretation of their model would be to relate their discount factor to an opportunity cost associated with alternative investments which are outside the model. The alternative investments are explicit in our model, and this fixes the appropriate discount factor.

3. The general case

We now turn to a more general setting, but with the asset sale problem of and the basis risk model of Section 2.1 as motivating examples.

3.1. The space of admissible strategies

Let $T_\infty \leq \infty$ be the largest time-horizon of interest, and let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T_\infty}, \mathbb{P})$ be a filtered probability space.

We begin by defining the class of admissible wealth processes for the agent. For our purposes, a wealth process is given by a triple (τ, W^-, W^+) , where $\tau \leq T_\infty$ is the time at which a certain action is taken, and the adapted processes W_t^- (defined for $0 \leq t \leq \tau$) and W_t^+ (defined for $t \leq \tau \leq T_\infty$) describe the wealth of the agent before and after this action. We assume the action is taken exactly once, although this assumption can be relaxed.

We will denote the class of admissible action times by \mathcal{T} , and suppose it is a subset of the set of stopping times. We suppose that $0 \in \mathcal{T}$.

Example 3.1. Let \mathcal{T}_T be the set of all stopping times with $\tau \leq T$.

Let $\mathcal{A}_{0,\tau}^-(w)$ denote the class of admissible wealth processes over the time period $[0, \tau]$ which satisfy $W_0^- = w$.

Similarly, for F_τ an \mathcal{F}_τ -measurable random variable, set $\mathcal{A}_{\tau,T_\infty}^+(F_\tau)$ to be the class of admissible wealth processes over the time period $[\tau, T_\infty]$ with initial value $W_\tau^+ = F_\tau$. Here, F_τ should be interpreted as the wealth of the agent at time τ .

Example 3.2. As a simplest example, suppose $\mathcal{A}_{0,\tau}^-(w) = \{W_t^- = w; 0 \leq t \leq \tau\}$ and $\mathcal{A}_{\tau,T_\infty}^+(F_\tau) = \{W_t^+ = F_\tau; \tau \leq t \leq T_\infty\}$. The idea behind this example is that there are no investment opportunities, so that wealth is constant, except for the potential change at the stopping time τ .

At the action time τ , we need an algorithm relating the pre- and post- τ wealth. In general, this relationship will depend on an auxiliary process Y_t . We suppose the pre- and post- τ wealths are related by the update rule $W_\tau^+ = H(\tau, W_\tau^-, Y_\tau)$.

Example 3.3. The problem of selling the real asset corresponds to the choice $H(t, w, y) = w + y$. More generally, we could consider options on the real asset, so, for example, $H(t, w, y) = w + (y - k)^+$. In Section 6.2 we will discuss an example of Hugonnier and Morellec [11] where $H(t, x, y) = wh(y)$ for h a positive function with $h(y) \leq 1$.

We are now able to put the components together to define the class of admissible wealth processes.

Definition 3.4. The wealth process $W = (\tau, W_t^-, W_t^+)$ with initial wealth w is admissible (we write $W \in \mathcal{A} = \mathcal{A}(w)$) if

- (i) $\tau \in \mathcal{T}$,
- (ii) W_t^- is an element of $\mathcal{A}_{0,\tau}^-(w)$,
- (iii) $W_\tau^+ = H(\tau, W_\tau^-, Y_\tau)$, and
- (iv) W_t^+ is an element of $\mathcal{A}_{\tau,T_\infty}^+(W_\tau^+)$.

Where it is clear from the context, we sometimes abbreviate $\mathcal{A}_{0,\tau}^-(w)$ by omitting either the time-subscripts, or the argument describing initial wealth, and similarly for \mathcal{A}^+ . Also, although \mathcal{A}^- and \mathcal{A}^+ are the primitive objects, we will talk about the horizon-unbiasedness relative to the derived space \mathcal{A} .

The fundamental problem facing the agent is to choose an optimal admissible wealth process (including the action time τ), so as to maximise the expected utility of wealth. We want to consider the sense in which the problem can be expressed as a mixed control/stopping problem

$$\sup_{\tau \in \mathcal{T}} \sup_{W^- \in \mathcal{A}_{0,\tau}^-} \mathbb{E}[U(\tau, H(\tau, W_\tau^-, Y_\tau))], \tag{5}$$

for an appropriately defined inter-temporal objective function $U(t, x)$. Our key contribution is to argue that this objective function cannot be defined in an arbitrary fashion, but must satisfy some quite restrictive conditions for the problem to be self-consistent. The main novelty comes from considering the perpetual versions of the problem, where $T_\infty = \infty$.

We close this section with a discussion of some consistency conditions on \mathcal{A} and $\mathcal{A}_{0, T_\infty}^+$. The latter space plays a dual role, first as the set of admissible strategies for an agent who took action at $\tau = 0$, and second as the set of admissible strategies for an agent who never had the possibility of choosing to take action — for example the agent who never had the real asset to sell. We use the notation X to denote the wealth of an agent in this second situation. The role of the consistency conditions is to allow us to divide admissible processes into their pre- and post-action parts, and conversely, to combine pre- and post- τ admissible processes into an admissible wealth process.

Given $X \in \mathcal{A}_{0, T_\infty}^+$ and $\sigma \in \mathcal{T}$, we can define the triple $(\sigma, X^{\sigma, -}, X^{\sigma, +})$ via $X_t^{\sigma, -} = X_t$ for $t < \sigma$, and $X_t^{\sigma, +} = X_t$ for $t > \sigma$. Here $X^{\sigma, -}$ is defined on $[0, \sigma]$, and $X^{\sigma, +}$ is defined on $[\sigma, T_\infty]$; also note that $X_\sigma^{\sigma, -} = X_\sigma^{\sigma, +}$. Conversely, given $(\sigma, X^{\sigma, -}, X^{\sigma, +})$, we can reconstruct X by $X_t = X_t^{\sigma, -} I_{\{t \leq \sigma\}} + X_t^{\sigma, +} I_{\{t > \sigma\}}$. Recall that $\mathcal{A}_{\sigma, T_\infty}^+$ is assumed as given, and for $\sigma \in \mathcal{T}$ define $\mathcal{A}_{0, \sigma}^+(x) = \cup_{X \in \mathcal{A}_{0, T_\infty}^+(x)} \{X^{\sigma, -}\}$.

Definition 3.5. We say that \mathcal{A} has the concatenation property if:

- (A1) Suppose the family $\{X(t, x)\}$ satisfies $X(t, x) \in \mathcal{A}_{t, T_\infty}^+(x)$. Suppose $\sigma \in \mathcal{T}$ and that F_σ is \mathcal{F}_σ measurable. Then $X^{\sigma, +}$ defined on $[\sigma, T_\infty]$ by $X_s^{\sigma, +} = X_s(\sigma, F_\sigma)$ is such that $X^{\sigma, +} \in \mathcal{A}_{\sigma, T_\infty}^+(F_\sigma)$.
- (A2) Suppose $\sigma \in \mathcal{T}$, $X^{\sigma, -} \in \mathcal{A}_{0, \sigma}^+(x)$ and $X^{\sigma, +} \in \mathcal{A}_{\sigma, T_\infty}^+(X_\sigma^{\sigma, -})$. Then $X \in \mathcal{A}_{0, T_\infty}^+(x)$.
- (A3) Suppose $\sigma \in \mathcal{T}$, $W^{\sigma, -} \in \mathcal{A}_{0, \sigma}^-(x)$ and $W^{\sigma, +} \in \mathcal{A}_{\sigma, T_\infty}^+(H(\sigma, W_\sigma^{\sigma, -}, Y_\sigma))$. Then $W \in \mathcal{A}_{0, T_\infty}(x)$.

The first part of the definition says that it is possible to define an element of \mathcal{A}^+ by conditioning on the initial value at σ . The second and third parts of the definition say that if the first part of an admissible strategy is followed by the second part of a (potentially different) admissible strategy, then the conjoined wealth process is also admissible, provided there is an appropriate change in wealth at the instant where the wealth processes are combined.

Assumption 3.6. \mathcal{A} has the concatenation property.

3.2. Admissibility in the non-traded assets model

We return to the model of Section 2.1 and the asset sale problem for a non-dividend paying asset. The natural candidate for the class $\mathcal{A}_{\tau, T_\infty}^+$ is the set of self-financing wealth processes for which $W_t^+ \geq 0$ for all $t \in [\tau, T_\infty]$. The non-negativity ensures that the discounted wealth process is a supermartingale under any equivalent martingale measure for the traded asset. This rules out Ponzi schemes. Observe that the definition of \mathcal{A}^+ does not depend on Y , and recall that the space \mathcal{A}^+ plays a dual role, as both the set of admissible strategies available to the agent who has sold the real asset and the set of admissible strategies available to the agent who never had access to the real asset.

For $\mathcal{A}_{0, \tau}^-$, there are two natural candidates. In both cases, the agent's wealth is self-financing. This criterion is then augmented by a non-negativity constraint for each $t \leq \tau$, either of the form

$W_t^- \geq 0$, or of the form $H(t, W_t^-, Y_t) \geq 0$. In the former case, the requirement is that the agent must keep her wealth in terms of liquid (i.e. financial) assets non-negative. In the latter case, the restriction is weaker, and the agent is permitted to allow her financial wealth to go negative, provided these debts are secured against the real asset.

In the non-traded assets model of Section 2.1, we can be even more explicit about the space of admissible strategies. The wealth process $W = (\tau, W_t^-, W_t^+)$ can be re-parameterised in terms of a trading strategy θ and the exercise rule τ . We write $W \equiv (\tau, W_t^{\theta, \tau})_{0 \leq t \leq T_\infty}$ for the wealth process in this parameterisation. In particular, $W_t^{\theta, \tau} = W_t^-$ for $t < \tau$ and $W_t^{\theta, \tau} = W_t^+$ for $t \geq \tau$. The idea is that, except at the time $t = \tau$, $W^{\theta, \tau}$ is a self-financing wealth process, and indeed, for $t < \tau$ or $t > \tau$, we have

$$dW_t^{\theta, \tau} = \theta_t dP_t + r(W_t^{\theta, \tau} - \theta_t P_t) dt. \tag{6}$$

At time $t = \tau$, the wealth process $W^{\theta, \tau}$ receives a lump-sum boost of size Y_τ :

$$W_\tau^{\theta, \tau} = W_{\tau-}^{\theta, \tau} + Y_\tau \equiv H(\tau, W_{\tau-}^{\theta, \tau}, Y_\tau). \tag{7}$$

This last condition is the appropriate specification of Definition 3.4(iii) in the asset sale problem.

In the light of the discussion of the two natural candidates for \mathcal{A}^- , it is convenient to define some natural classes of admissible strategies for the model.

- Definition 3.7.** (i) The wealth process W is L -admissible (or liquid-admissible) if W satisfies (6) and (7) and $W_t^{\theta, \tau} \geq 0$. For this family of admissible wealths, the pre- τ constraint on wealth is $W_t^- \geq 0$, so that liquid wealth must be kept non-negative prior to exercise.
- (ii) The wealth process W is C -admissible (or collateral-admissible) if W satisfies (6) and (7), if $W_t^{\theta, \tau} + Y_t \geq 0$ for $t < \tau$ and $W_t^{\theta, \tau} \geq 0$ for $t \geq \tau$, so that the solvency requirement imposed on the agent allows her to sell the real asset at its instantaneous value. Effectively she is allowed to borrow against her holding in the real asset.

Note that in both cases, $\mathcal{A}_{0, T_\infty}^+(x) = \{X : X_0 = x, X \text{ satisfies (6) and } X_t \geq 0\}$. It immediately follows that \mathcal{A} has the concatenation property.

In this model, we can define an admissible portfolio strategy as one for which the associated wealth process is admissible. Given $\tau \in \mathcal{T}$, we write $\theta \in \Theta(x)$ if $(\tau, W_t^{\theta, \tau}) \in \mathcal{A}(x)$ with similar definitions for $\Theta_{0, \tau}^-$, $\Theta_{0, \tau}^+$ and $\Theta_{\tau, T_\infty}^+$.

- Definition 3.8.** (i) The class Θ^L of liquid-admissible strategies is such that W is given by (6) and (7) and $W_t^{\theta, \tau} \geq 0$.
- (ii) The class Θ^C of collateral-admissible strategies is such that W is given by (6) and (7) and $W_t^{\theta, \tau} + Y_t \geq 0$ for $t < \tau$ and $W_t^{\theta, \tau} \geq 0$ for $t \geq \tau$.

4. The terminal horizon case

4.1. Horizon-unbiased inter-temporal utility functions

Consider an agent with initial wealth w who aims to maximise her expected utility of terminal wealth, where the terminal horizon is fixed at $T_\infty = T < \infty$, and the utility function at time T is given by the (concave, increasing) function \bar{U} . The problem facing the agent is to find the optimal admissible stopping rule and wealth process:

$$\sup_{W \in \mathcal{A}(w)} \mathbb{E}[\bar{U}(W_T^+)]. \tag{8}$$

where, recall Definition 3.4, $\mathcal{A}(w)$ is the set of admissible wealth processes.

Consider the (Merton-style) optimal investment problem facing the comparable agent whose wealth is not affected by any choice of action:

$$\sup_{X \in \mathcal{A}_{0,T}^+(x)} \mathbb{E}[\bar{U}(X_T)] \equiv \sup_{W^+ \in \mathcal{A}_{0,T}^+(x)} \mathbb{E}[\bar{U}(W_T^+)].$$

(Here, the roles of X and W^+ are interchangeable, but we shall generally use the former in circumstances which do not involve the choice of stopping rule.) We denote by \tilde{U} the solution of this problem at an intermediate time t (the indirect utility function):

$$\tilde{U}(t, X_t) = \sup_{X \in \mathcal{A}_{t,T}^+(X_t)} \mathbb{E}[\bar{U}(X_T) | \mathcal{F}_t]. \tag{9}$$

Assumption 4.1. We suppose that the solution to the Merton-style optimal investment problem exists, and can be written in this form.

This assumption will certainly be satisfied in the set-up of Section 3.2, but will not hold, for example, in the case of stochastic volatility models for the financial assets.

Now return to the problem in (8). From (9) and the identification of X with W^+ in $\mathcal{A}_{\tau,T}^+(x)$, we have

$$\tilde{U}(\tau, H(\tau, W_\tau^-, Y_\tau)) = \sup_{W^+ \in \mathcal{A}_{\tau,T}^+(H(\tau, W_\tau^-, Y_\tau))} \mathbb{E}[\bar{U}(W_T^+) | \mathcal{F}_\tau]. \tag{10}$$

Then, taking expectations and a supremum over τ and W_τ^- , we obtain:

Theorem 4.2. Suppose Assumptions 3.6 and 4.1 hold. Let $\mathcal{T} \subseteq \mathcal{T}_T$ be a set of admissible stopping times. In the set-up of Section 3.1, the problems

$$\sup_{W \in \mathcal{A}(w)} \mathbb{E}[\bar{U}(W_T^+)] \tag{11}$$

and

$$\sup_{\tau \in \mathcal{T}} \sup_{W^- \in \mathcal{A}^-(w)} \mathbb{E}[\tilde{U}(\tau, H(\tau, W_\tau^-, Y_\tau))] \tag{12}$$

are equivalent, where the inter-temporal utility function \tilde{U} is given by (9).

Proof. The result is essentially tautological, and the only issue is to check that the set of admissible wealth processes has the relevant properties.

Fix τ and $W^- \in \mathcal{A}^-(w)$. By Assumption 3.6, and by implication Definition 3.5(A3), which says that the concatenation of pre- and post- τ admissible wealth processes is itself admissible, and the result will follow if

$$\sup_{W^+ \in \mathcal{A}_{\tau,T}^+(H(\tau, W_\tau^-, Y_\tau))} \mathbb{E}[\bar{U}(W_T^+)] = \mathbb{E}[\tilde{U}(\tau, H(\tau, W_\tau^-, Y_\tau))], \tag{13}$$

where, by definition,

$$\tilde{U}(\tau, H(\tau, W_\tau^-, Y_\tau)) = \sup_{W^+ \in \mathcal{A}_{\tau, T}^+(H(\tau, W_\tau^-, Y_\tau))} \mathbb{E}[\bar{U}(W_T^+) | \mathcal{F}_\tau].$$

For $W^+ \in \mathcal{A}_{\tau, T}^+(H(\tau, W_\tau^-, Y_\tau))$, we have

$$\mathbb{E}[\bar{U}(W_T^+) | \mathcal{F}_\tau] \leq \sup_{W^+ \in \mathcal{A}_{\tau, T}^+(H(\tau, W_\tau^-, Y_\tau))} \mathbb{E}[\bar{U}(W_T^+) | \mathcal{F}_\tau]$$

and we have \leq in (13). For the converse, given $\epsilon > 0$, for each (t, x) , choose a process $W^\epsilon(t, x) \in \mathcal{A}_{t, T}^+(x)$ such that $\mathbb{E}[\bar{U}(W_T^\epsilon(t, x))] \geq \tilde{U}(t, x) - \epsilon$, and construct the process $W^\epsilon(\tau, H(\tau, W_\tau^-, Y_\tau))$ by conditioning on the value of W^- and Y at τ . By the concatenation property (A1), $W^\epsilon \in \mathcal{A}_{\tau, T}^+(H(\tau, W_\tau^-, Y_\tau))$ and

$$\mathbb{E}[\bar{U}(W_T^\epsilon) | \mathcal{F}_\tau] \geq \tilde{U}(\tau, H(\tau, W_\tau^-, Y_\tau)) - \epsilon.$$

The required inequality follows on taking expectations. \square

We have shown that the utility maximisation problem (11) can be reduced to an optimal stopping/portfolio choice problem (12), provided the inter-temporal utility function is given by (9). In particular, (12) is well defined provided the inter-temporal utility function satisfies certain consistency conditions.

One can also approach this problem from the opposite direction by asking, if we write down an inter-temporal utility function $U(t, x)$ when, and in what sense, we have a unbiased problem.

Definition 4.3. The optimal stopping problem for U , \mathcal{T} and \mathcal{A} has no preferred horizon if

$$\sup_{X \in \mathcal{A}_{0, \tau}^+} \mathbb{E}[U(\tau, X_\tau)]$$

does not depend on $\tau \in \mathcal{T}$. Otherwise, we say the optimal stopping problem has a preferred horizon.

Note that the definition only depends on \mathcal{A} via \mathcal{A}^+ .

The idea is that if the exercise problem has a preferred horizon, then artificial incentives are introduced which encourage one stopping time to be preferred over another. Any conclusions about the optimal stopping rule for the full problem (11) or (12) are biased by these incentives. From a mathematical standpoint, the optimal stopping and control problem in (12) is well defined for a general function $U(t, x)$, but if there is a preferred horizon, then the economic interpretation is distorted.

Proposition 4.4. \tilde{U} given by (9) results in a problem with no preferred horizon for $(\mathcal{A}, \mathcal{T})$.

Proof. Fix $\tau \in \mathcal{T}$. By (9) and the concatenation property of \mathcal{A}^+ (the proof mirrors the proof of Theorem 4.2, except that it uses (A2) rather than (A3) of Definition 3.5), we have

$$\begin{aligned} \sup_{X^- \in \mathcal{A}_{0, \tau}^+} \mathbb{E}[\tilde{U}(\tau, X_\tau)] &= \sup_{X^- \in \mathcal{A}_{0, \tau}^+} \mathbb{E} \left[\sup_{X^+ \in \mathcal{A}_{\tau, T}^+(X_\tau)} \mathbb{E}[\bar{U}(X_T) | \mathcal{F}_\tau] \right] \\ &= \sup_{X \in \mathcal{A}_{0, T}^+} \mathbb{E}[\bar{U}(X_T)], \end{aligned}$$

and this last expression is independent of τ . \square

Remark 4.5. Observe that if we take $\tau = 0$ in the above argument, then we find $\tilde{U}(0, x) = \sup_{X \in \mathcal{A}_{0,T}^+(x)} \mathbb{E}[\tilde{U}(X_T)]$.

Remark 4.6. The key quantities of interest in the problem are the optimal decision rule and hedging strategy for the agent, together with the associated utility-indifference price of the real asset. The solution of the utility maximisation problem, in itself, is less important, except as a tool in deriving the other quantities.

Given **Remark 4.5**, the utility-indifference price p is the solution to

$$\tilde{U}(0, x + p) = \sup_{\tau \in \mathcal{T}} \sup_{W^- \in \mathcal{A}^-(w)} \mathbb{E}[\tilde{U}(\tau, H(\tau, W_\tau^-, Y_\tau))]. \tag{14}$$

In particular, the values of any economically important variables are not affected if the terminal utility function \tilde{U} is multiplied by a constant.

4.2. Constant Relative Risk Aversion

In this section, we consider an agent with a Constant Relative Risk Aversion (CRRA) utility function of the form ($R > 0, R \neq 1$)

$$\bar{U}(x) = A \frac{x^{1-R}}{1-R}$$

in the non-traded assets model of Section 2.1.

It is a standard exercise dating back to Merton [16] to show that \tilde{U} , as defined in (9), is given by

$$\tilde{U}(t, X_t) = B e^{-\beta t} \frac{X_t^{1-R}}{1-R},$$

where $\beta \equiv \beta_R = (1-R)r + (1-R)\lambda^2/2R$ and $B = A e^{\beta T}$.

The optimal sale problem becomes to find

$$\sup_{\tau \in \mathcal{T}} \sup_{\theta \in \Theta} \mathbb{E}[\bar{U}(W_T^{\theta, \tau})].$$

By **Theorem 4.2** and **Remark 4.6**, up to a constant this problem is equivalent to finding

$$\sup_{\tau \in \mathcal{T}} \sup_{\theta \in \Theta^-} \mathbb{E}[\hat{U}(\tau, W_\tau^{\theta, \tau})]$$

where $\hat{U}(t, x) = e^{-\beta t} x^{1-R}/(1-R)$. This motivates the choice of discount factor in **Definition 2.2**.

Remark 4.7. Suppose that we aim to solve

$$\sup_{\tau \in \mathcal{T}} \sup_{\theta \in \Theta^-} \mathbb{E}[U(\tau, W_{\tau^-}^{\theta, \tau} + cY_\tau)]$$

for a power law inter-temporal utility function

$$U(t, x) = e^{-\delta t} \frac{x^{1-R}}{1-R}$$

where δ is an arbitrary discount factor. Suppose we consider this problem without the real asset, or equivalently we set $c = 0$. Then we are left to solve the Merton problem over $[0, T]$. We find that the optimal stopping rule is $\tau = 0$ if $\delta > \beta$, and $\tau = T$ if $\delta < \beta$. Only if $\delta = \beta$ is the agent indifferent to the choice of stopping time, and only in that case can we be sure that conclusions about the stopping rule when $c = 1$ are driven by the presence of the real option rather than structural features of the optimal stopping problem.

4.3. Duality arguments

In many settings, the dual approach has been very successful in solving and providing insight into the solutions of utility maximisation problems. One such setting is the non-traded assets model of Section 3.2. In this model, an element $X \in \mathcal{A}_{0,t}^+(x)$ can be identified with its terminal value, which must satisfy

$$\mathbb{E}[X_t \zeta_t] \leq x$$

for all state price densities ζ_t .

Let U be an inter-temporal utility function and define the convex dual $V(t, y) = \sup_x \{U(t, x) - xy\}$, and note that this can be inverted to give $U(t, x) = \inf_y \{V(t, y) + xy\}$. We have that

$$\mathbb{E}[U(t, X_t)] \leq \mathbb{E}[U(t, X_t) - \eta(X_t \zeta_t - x)] \leq \mathbb{E}V(t, \eta \zeta_t) + \eta x$$

so that

$$\sup_{X \in \mathcal{A}_{0,t}^+} \mathbb{E}[U(t, X_t)] \leq \inf_{\eta} \{H(t, \eta) + \eta x\} \tag{15}$$

where $H(t, \eta) = \inf_{\zeta_t} \mathbb{E}V(t, \eta \zeta_t)$.

In the non-traded assets model, there is equality in (15), and the infimum over ζ_t is attained by $\zeta_t^* = e^{-rt} \mathcal{E}(-\lambda \cdot B)$. (The same state-price-density is minimal for all convex functions $V(t, \cdot)$.) Suppose also that U is such that the problem has no preferred horizon. Then $H(t, \eta)$ is independent of t . In particular, V must satisfy

$$V(0, y) = \inf_{\zeta_t} \mathbb{E}V(t, y \zeta_t) = \mathbb{E}V(t, y \zeta_t^*).$$

For example, if $V(t, y) = (R/(1 - R))e^{-\gamma t} y^{1-1/R}$, then $\gamma = \beta/R$. This is consistent with the results in Section 4.2.

5. The infinite horizon problem

5.1. Horizon-unbiased utility functions

The analysis so far depended crucially on the use of a fixed terminal horizon T , at which the utility of wealth is ultimately determined. However, many real asset sale problems take place without an upper bound on the sale time. Moreover, there is the hope that for certain problems, the perpetual version will be more tractable than the finite horizon version. (For example, this is true for the standard American put in a complete market.) For this reason, we wish to extend the above analysis to the infinite horizon, and rather than starting with the pair T, \bar{U} , we aim to write down a problem of the form (12) directly, for a suitably defined inter-temporal utility function \tilde{U} . The key observation is that (modulo technical conditions)

$\tilde{U}(t, X_t)$ as defined in (9) is a supermartingale, and a martingale for the optimal wealth process $X_t \in \mathcal{A}^+$.

Thus it is natural to insist that the agent with inter-temporal utility function $U(t, x)$, solving the problem

$$\sup_{\tau} \sup_{W^- \in \mathcal{A}^-} \mathbb{E}[U(\tau, H(\tau, W_{\tau}^-, Y_{\tau}))]$$

over the infinite horizon, should be required to use an objective function for which $U(t, X_t)$ is a supermartingale for any element of \mathcal{A}^+ and a martingale for some element.

Definition 5.1. The inter-temporal utility function $U(t, x)$ is *horizon-unbiased* for $(\mathcal{A}, \mathcal{T})$ if for every admissible τ and for every admissible wealth processes X in $\mathcal{A}_{0,\tau}^+$, we have that $U(t \wedge \tau, X_{t \wedge \tau})$ is a supermartingale, and if for each $\tau \in \mathcal{T}$, there exists $X \in \mathcal{A}_{0,\tau}^+$ such that $U(t \wedge \tau, X_{t \wedge \tau})$ is a (uniformly integrable) martingale.

Theorem 5.2. If $U(t, x)$ is horizon-unbiased, then for every admissible τ , we have

$$\sup_{X \in \mathcal{A}_{0,\tau}^+(x)} \mathbb{E}[U(\tau, X_{\tau})] = U(0, x). \tag{16}$$

In particular, the left-hand-side is independent of $\tau \in \mathcal{T}$, and the optimal stopping problem has no preferred horizon for $U, \mathcal{T}, \mathcal{A}$.

Thus, for a horizon-unbiased inter-temporal utility function, and for the problem in (16), the agent is indifferent over the choice of τ . When the additional claim Y is introduced, we can be sure that the choice of optimal stopping rule is not biased by the solution of the problem without Y . However, the uniform integrability requirement of horizon-unbiasedness is a severe requirement over the infinite horizon, and is not achieved in practice for many natural examples, as we shall see in the next section. Hence, we introduce some modified classes of stopping times. In particular, let \mathcal{T}_{∞} denote the set of all finite-valued stopping times, and let $\mathcal{T}_B = \cup_{T < \infty} \mathcal{T}_T$ denote the set of bounded stopping times.

Definition 5.3. (i) If U is horizon-unbiased for $(\mathcal{A}, \mathcal{T}_{\infty})$, then we say that U is *horizon-unbiased over the infinite horizon*.

(ii) If U is horizon-unbiased for $(\mathcal{A}, \mathcal{T}_B)$, then we say that U is *horizon-unbiased over every finite horizon*.

(iii) If U is horizon-unbiased over every finite horizon and

$$\sup_{\tau \in \mathcal{T}_{\infty}} \sup_{X \in \mathcal{A}_{0,\tau}^+} \mathbb{E}[U(\tau, X_{\tau})] = U(0, x)$$

then we say that U is *weakly horizon-unbiased over the infinite horizon*.

5.2. The infinite horizon problem and CRRA preferences

In this section, we work with the model of Section 3.2 and with CRRA preferences.

In Section 4.2, we saw that the finite horizon problem could be rewritten as

$$\sup_{\tau \in \mathcal{T}} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[\frac{(X_{\tau}^{\theta} + Y_{\tau})^{1-R}}{1-R} e^{-\beta\tau} \right] \tag{17}$$

where $X_t^\theta \equiv W_t^{\theta, \tau}$, and where, following Remark 4.6 we have ignored any pre-factor. The key observation is that this problem makes sense over the infinite horizon. Thus, even if it is not possible to relate (17) to a terminal horizon problem, the optimal stopping/portfolio choice problem implicit in (17) is in itself well defined.

Fix R , and let $M_t^\theta = e^{-\beta t} (X_t^\theta)^{1-R} / (1 - R)$. Then

$$\begin{aligned} \frac{dM_t^\theta}{M_t^\theta} &= -\beta dt + (1 - R) \frac{dX_t^\theta}{X_t^\theta} - \frac{R(1 - R)}{2} \left(\frac{dX_t^\theta}{X_t^\theta} \right)^2 \\ &= (1 - R) \frac{\theta_t \eta P_t}{X_t^\theta} dB_t^P - \frac{R(1 - R)}{2} \left(\frac{\theta_t \eta P_t}{X_t^\theta} - \frac{\lambda}{R} \right)^2 dt. \end{aligned} \tag{18}$$

Let $\psi_t = \lambda X_t^\psi / \eta R P_t$. Then M_t^ψ is a local martingale, and moreover

$$M_t^\psi = \frac{x^{1-R}}{1 - R} e^{((1-R)\lambda/R)B_t^P - ((1-R)^2\lambda^2/2R^2)t}.$$

Define $\mathcal{T}_{UI} = \{\tau : M_{t \wedge \tau}^\psi \text{ is uniformly integrable}\}$.

Lemma 5.4. For $\tau \in \mathcal{T}_{UI}$, (16) holds. Further, if $\mathcal{T} \equiv \mathcal{T}_{UI}$, then $U(t, x) = e^{-\beta t} x^{1-R} / (1 - R)$ is a horizon-unbiased utility function.

Proof. For any $\tau \in \mathcal{T}_{UI}$, we can define $\hat{\mathbb{P}}$, equivalent to \mathbb{P} via

$$\left. \frac{d\hat{\mathbb{P}}}{d\mathbb{P}} \right|_{\mathcal{F}_{t \wedge \tau}} = \frac{(1 - R)}{x^{1-R}} M_{t \wedge \tau}^\psi.$$

By Girsanov's theorem, under $\hat{\mathbb{P}}$, for $t \leq \tau$ $\hat{X}_t^\theta := X_t^\theta / X_t^\psi$ is a non-negative local martingale, and hence a supermartingale. Then, for any θ , and any $G \in \mathcal{F}_{s \wedge \tau}$, we get

$$\begin{aligned} \mathbb{E} \left[\frac{e^{-\beta(t \wedge \tau)} (X_{t \wedge \tau}^\theta)^{1-R}}{1 - R} I_G \right] &= \mathbb{E} \left[M_{t \wedge \tau}^\psi (\hat{X}_{t \wedge \tau}^\theta)^{1-R} I_G \right] \\ &= \frac{x^{1-R}}{1 - R} \hat{\mathbb{E}} \left[(\hat{X}_{t \wedge \tau}^\theta)^{1-R} I_G \right]. \end{aligned}$$

By Jensen's inequality, and the fact that \hat{X}_t^θ is a supermartingale under $\hat{\mathbb{P}}$, this last expression can be bounded by

$$\frac{x^{1-R}}{1 - R} \hat{\mathbb{E}} \left[(\hat{X}_{s \wedge \tau}^\theta)^{1-R} I_G \right] = \mathbb{E} \left[M_{s \wedge \tau}^\psi (\hat{X}_{s \wedge \tau}^\theta)^{1-R} I_G \right] = \mathbb{E} \left[e^{-\beta(s \wedge \tau)} \frac{(X_{s \wedge \tau}^\theta)^{1-R}}{1 - R} I_G \right].$$

Hence, $U(t \wedge \tau, X_{t \wedge \tau}^\theta)$ is a submartingale. Further, there is equality throughout for the choice $\theta = \psi$, and for this choice of θ , \hat{X}_t^ψ is constant, and $(\hat{X}_t^\psi)^{1-R}$ is trivially a uniformly integrable martingale. \square

Theorem 5.5. Let $U(t, x) = e^{-\beta t} x^{1-R} / (1 - R)$, and suppose $\Theta = \Theta^L$ or $\Theta = \Theta^C$. Then $U(t, x)$ is horizon-unbiased over every finite horizon. Moreover,

- (i) Suppose $\lambda = 0$. Then $U(t, x)$ is horizon-unbiased over the infinite horizon.
- (ii) Suppose $R < 1$. Then $U(t, x)$ is weakly horizon-unbiased over the infinite horizon.

Proof. Since $\mathcal{T}_T \subseteq \mathcal{T}_{UI}$, the first part follows easily. Moreover, if $\lambda = 0$, then M^ψ is constant, and $\mathcal{T}_{UI} = \mathcal{T}_\infty$.

For (ii), we have from (18) that for all stopping times τ , $M_{t \wedge \tau}^\theta$ is a non-negative local supermartingale, and hence a supermartingale. Hence, for all τ ,

$$\mathbb{E} \left[e^{-\beta\tau} \frac{(X_\tau^\theta)^{1-R}}{1-R} \right] \leq U(0, x). \quad \square$$

Remark 5.6. When $R > 1$, and $\lambda \neq 0$, $U(t, x)$ is not weakly horizon-unbiased over the infinite horizon. In this case, M_t^ψ is negative but is not bounded below; indeed $M_t^\psi \rightarrow 0$ almost surely, and there is a sequence of (non-uniformly integrable) stopping times such that

$$\sup_{\theta \in \Theta^+} \mathbb{E}U(\tau_n, X_{\tau_n}^\theta) \rightarrow 0 > U(0, x).$$

For $R > 1$, some restriction on the set of admissible stopping times is necessary to get a non-degenerate problem.

5.3. Other utilities and the infinite horizon

We continue to investigate the optimal time to sell the real asset in the constant parameter non-traded assets model, but now we consider more general utility functions, and the question of whether it is possible to define other inter-temporal utility functions which possess the horizon-unbiasedness property. In particular, following the remarks in Section 5, we wish to find smooth functions $U(t, x)$ such that $\sup_\theta U(t, X_t^\theta)$ is a supermartingale in general, and a (local) martingale for the optimal strategy.

Under the assumption that U is sufficiently differentiable, we can apply Itô's formula. It follows that U must satisfy

$$0 = \sup_\theta \left\{ U_t + \frac{\theta^2 p^2 \eta^2}{2} U_{xx} + (\theta v p + r(x - \theta p)) U_x \right\}$$

(where the subscript t now refers to a time derivative) which simplifies to

$$0 = U_t - \frac{\lambda^2}{2} \frac{U_x^2}{U_{xx}} + r x U_x \tag{19}$$

where $\lambda = (v - r)/\eta$ is the Sharpe ratio of the financial asset. This is a non-linear equation, but it simplifies greatly under the Legendre transformation: set $y = U_x$, $s = \lambda^2 t$ and $v(s, y) = x U_x - U$, and then

$$v_s = -\frac{y^2}{2} v_{yy} + \frac{r}{\lambda^2} y v_y$$

where v is a concave function in y . If trading wealth is restricted to be positive ($X \geq 0$), then v is increasing in y . Subject to a rescaling of time, the function v is the negative of the dual function V introduced in Section 4.3.

Set $y = e^{z - ((r/\lambda^2) - 1/2)s}$, $w(s, z) = e^{-(r/\lambda^2)s} v_y(s, y)$ and $w(s, z) = v(s, y) = v(s, e^{z - ((r/\lambda^2) - 1/2)s})$. Then w solves the backward heat equation

$$w_s = -\frac{1}{2} w_{zz}. \tag{20}$$

We want positive, increasing solutions of this equation.

Note that the backward heat equation is ill-posed. For most initial data, the solution will cease to exist after a finite time-period. However, we know one family of initial data for which the solution is defined for all time, namely the family associated with the CRRA utility of the previous section.

The CRRA example takes the form $U(t, x) = e^{-\beta t} x^{1-R} / 1 - R$, together with

$$v(s, y) = -\frac{R}{1-R} y^{1-1/R} e^{-(\beta/\lambda^2 R)s}, \quad w(s, z) = e^{-z/R - s/2R^2}.$$

These are the travelling wave solutions to (20). Essentially these correspond to the only simple horizon-unbiased utility functions.

Theorem 5.7. *Suppose $\lambda \neq 0$. Let $U(t, x)$ be a horizon-unbiased utility function in the sense of Definition 5.1 for the non-traded assets model over the infinite horizon. Suppose U is separable in the sense that $U(t, x) = F(xe^{-rt})G(t)$ for the functions F and G , such that F is strictly increasing and concave, and $F'(0) = \infty$, and such that G is differentiable. Then U is of the CRRA family.*

Proof. It follows from (19) that

$$\frac{G'}{G} = c \frac{\lambda^2}{2} = \frac{\lambda^2}{2} \frac{F'^2}{FF''}$$

for some constant c . If we set $F = H^\alpha / \alpha$ for $\alpha = c / (c - 1)$, then H solves $H'' = 0$ with solution $H(x) = x$ (up to a multiplicative constant). If c is such that $\alpha = (1 - R)$, then we recover U . \square

Remark 5.8. The case where $c = 1$ leads to $F(z) = e^{-\gamma z}$ and $G(t) = e^{\lambda^2 t / 2}$, which is the horizon-unbiased utility function for exponential utility — see Henderson [8]. Horizon-unbiased logarithmic utility is associated with a solution of the form $U(t, x) = F(xe^{-rt}) + G(t)$.

Since (20) is linear, a more general family of solutions can be obtained by taking positive combinations.

Theorem 5.9. *Let w be a positive superposition of decreasing travelling wave solutions of the backward heat equation:*

$$w(s, z) = \int_0^\infty e^{-z/R} e^{-s/R^2} \mu(dR)$$

where μ is a finite measure with support in some interval contained in $(0, \infty)$. Set

$$v(s, y) = -\int_0^\infty \frac{R}{1-R} y^{1-1/R} e^{-(\beta R/\lambda^2 R)s} \mu(dR)$$

and let $y = F_R^{-1}(s, x)$ be the solution to

$$x = F_R(y, s) = \int_0^\infty y^{-1/R} e^{-(\beta R/\lambda^2 R)s} \mu(dR).$$

Then $U(x, t) = (F_R^{-1}(\lambda^2 t, x))v_y(\lambda^2 t, F_R^{-1}(\lambda^2 t, x)) - v(\lambda^2 t, F_R^{-1}(\lambda^2 t, x))$ is a horizon-unbiased utility function over every finite horizon.

Example 5.10. Let $\mu(dR) = (\delta_{(R-1/3)} + \delta_{(R-2/3)})dR$, where δ is the Dirac-measure. Then

$$U(t, x) = \left(\frac{3}{2} \left[\left(x e^{-(r+9\lambda^2/4)t} + \frac{1}{4} \right)^{1/2} - \frac{1}{2} \right]^{4/3} + 3 \left[\left(x e^{-(r+9\lambda^2/4)t} + \frac{1}{4} \right)^{1/2} - \frac{1}{2} \right]^{1/3} \right) e^{\lambda^2 t/2}$$

is weakly horizon-unbiased over the infinite horizon.

6. Further examples and conclusions

6.1. Discussion of Henderson [8] and Evans et al. [6]

As mentioned earlier, Henderson [8] and Evans et al. [6] consider problems of real options or assets in an incomplete setting. Both [8] and [6] consider the perpetual version of the problem, and hence it is important that they use a horizon-unbiased utility function. An in-depth discussion of such utilities was deferred to this paper.

As in this paper, it is assumed that there is a correlated asset which may be used for hedging the idiosyncratic risk, which is implicit in continued investment in the real asset. Henderson [8] considered the investment timing problem where a lump-sum investment payoff is received for an investment cost, represented by the option strike. She solved the problem in closed-form for exponential utility. Evans et al. [6] treated the power or CRRA utility for which the solution to the problem depends on wealth, and specialised to the asset sale problem (corresponding to a zero strike).

The aim in each case is to characterise the optimal strategy of the agent, both in terms of the optimal time to sell the real asset or invest, and the optimal hedging strategy in the non-traded asset, and to determine the utility-indifference value to the agent of selling/investing.

The utility maximisation problem of [6] is to find (at least when stopping times are constrained to be finite)

$$\sup_{\tau \in \mathcal{T}} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[e^{-\beta\tau} \frac{(X_\tau^\theta + Y_\tau)^{1-R}}{1-R} \right]. \tag{21}$$

The dynamics of the price processes are as given in Section 2.1, and the class of admissible strategies is as specified in Definition 3.8. (Thus there are two variants of the problem, depending on whether we use the liquid-admissible constraint, or the collateral-admissible constraint for \mathcal{A}^- .) In [6], the optimal time to sell the real asset is characterised as the first time that the ratio of the value of the real asset to the wealth of the agent exceeds a critical level, where this critical level is the solution of a transcendental equation. In particular, this means that if τ^* is the optimal stopping rule, then $\mathbb{P}(\tau^* = \infty) > 0$. For this reason, it makes sense to consider the expanded class of stopping rules $\mathcal{T}_\infty^{\leq} = \{\tau : \tau \leq \infty\}$.

Note that for $R < 1$, $e^{-\beta t} (X_t^\theta)^{1-R} / (1-R)$ is a non-negative submartingale, and hence converges almost surely. Thus, if we want to allow infinite stopping times, we can replace (21) with

$$\sup_{\tau \in \mathcal{T}_\infty^{\leq}} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[e^{-\beta\tau} \frac{(X_\tau^\theta + Y_\tau I_{\{\tau < \infty\}})^{1-R}}{1-R} \right] \tag{22}$$

where, on $\tau = \infty$, $e^{-\beta\tau} (X_\tau^\theta)^{1-R} / (1-R)$ is replaced by its limiting value.

Remark 6.1. Note that for some θ , $\lim_{t \uparrow \infty} \mathbb{E}[e^{-\beta t} (X_t^\theta)^{1-R}; \tau = \infty] > 0$. In particular, it is not appropriate to apply a transversality condition.

Let p_T be the solution to

$$\frac{(x + p_T)^{1-R}}{1 - R} = \sup_{\tau \in \mathcal{T}_T} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[e^{-\beta \tau} \frac{(X_\tau^\theta + Y_\tau)^{1-R}}{1 - R} \right]$$

and define $p_{\lim} = \lim_{T \uparrow \infty} p_T$. (Note that p_T is increasing, so the limit exists.) Define also p_B (respectively p_{UI} , p_∞ , p_∞^\leq) as the solution to

$$\frac{(x + p_B)^{1-R}}{1 - R} = \sup_{\tau \in \mathcal{T}_B} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[e^{-\beta \tau} \frac{(X_\tau^\theta + Y_\tau I_{\{\tau < \infty\}})^{1-R}}{1 - R} \right]$$

(where \mathcal{T}_B is replaced by \mathcal{T}_{UI} , \mathcal{T}_∞ , \mathcal{T}_∞^\leq respectively).

It is easy to see that $p_{\lim} = p_B \leq p_{UI} \leq p_\infty \leq p_\infty^\leq$.

Theorem 6.2. Suppose $R < 1$. Then $p_{\lim} = p_B = p_{UI} = p_\infty = p_\infty^\leq$.

Proof. It is sufficient to show that

$$\begin{aligned} & \lim_{T \uparrow \infty} \sup_{\tau \in \mathcal{T}_T} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[e^{-\beta \tau} \frac{(X_\tau^\theta + Y_\tau I_{\{\tau < \infty\}})^{1-R}}{1 - R} \right] \\ & \geq \sup_{\tau \in \mathcal{T}_\infty^\leq} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[e^{-\beta \tau} \frac{(X_\tau^\theta + Y_\tau I_{\{\tau < \infty\}})^{1-R}}{1 - R} \right]. \end{aligned}$$

Fix $\tau \in \mathcal{T}_\infty^\leq$ and $\theta \in \Theta_{0,\tau}^-$, and define $N_t = e^{-\beta t} (X_t^\theta + Y_t I_{\{t < \infty\}})^{1-R} / (1 - R)$, and $M_t = e^{-\beta t} (X_t^\theta)^{1-R} / (1 - R)$. Then, M is a non-negative supermartingale, which converges almost surely, $N_t \geq M_t$, and $N_\infty = M_\infty$. We want to show that $\lim_{T \uparrow \infty} \mathbb{E}[N_{T \wedge \tau}] \geq \mathbb{E}[N_\tau]$.

Suppose first that $\mathbb{E}[N_\tau] = \infty$. Then, for all t , we must have that $\mathbb{E}[N_\tau; t < \tau < \infty] = \infty$ and, fixing t and taking $T > t$, $\mathbb{E}[N_{T \wedge \tau}] \geq \mathbb{E}[N_\tau; t < \tau < T] \uparrow \infty$.

Otherwise, we have $\mathbb{E}[N_\tau; t < \tau < \infty] < \infty$ and

$$\mathbb{E}[N_\tau] = \mathbb{E}[N_\tau; \tau \leq T] + \mathbb{E}[N_\tau; T < \tau < \infty] + \mathbb{E}[N_\tau; \tau = \infty]. \tag{23}$$

For the last term, we have

$$\mathbb{E}[N_\tau; \tau = \infty] = \mathbb{E}[M_\infty; \tau = \infty] \leq \mathbb{E}[M_\tau; \tau > T] \leq \mathbb{E}[M_T; \tau > T] \leq \mathbb{E}[N_T; \tau > T].$$

Combining this with the first term on the right hand side of (23), we conclude

$$\mathbb{E}[N_\tau] \leq \mathbb{E}[N_{\tau \wedge T}] + \mathbb{E}[N_\tau; T < \tau < \infty],$$

and taking limits, we have $\lim_{T \uparrow \infty} \mathbb{E}[N_{\tau \wedge T}] \geq \mathbb{E}[N_\tau]$. \square

Theorem 6.3. Suppose $R > 1$, $\lambda \neq 0$ and $Y \equiv 0$. Then $p_B < p_\infty$.

Proof. This is immediate from the fact that for $R > 1$, U is horizon-unbiased over every finite horizon, but not weakly horizon-unbiased over the infinite horizon. See Theorem 5.5 and Remark 5.6. \square

6.2. A model of corporate control of Hugonnier and Morellec

In its simplest version, the problem in Hugonnier and Morellec [11] is to find

$$\sup_{\tau \in \mathcal{T}} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E}U(\tau, H(\tau, W_\tau^{\theta,\tau}, Y_\tau)) \tag{24}$$

in the non-traded assets model of Section 3.2, with $\mathcal{T} = \mathcal{T}_\infty$, and $\Theta = \Theta^L$. Hugonnier and Morellec use the CRRA utility $U(t, x) = e^{-\delta t} x^{1-R} / (1 - R)$ for $R < 1$ and the update function $H(t, w, y) = wh(y)$, where h is a positive function with $h(y) \leq 1$ and $h(y^*) = 1$ for some y^* . The idea is that the manager chooses the time of investment for the firm, and his investment decision is compared with the optimal decision for the shareholders. If, from the standpoint of the shareholders, his behaviour is suboptimal, then he faces the risk of dismissal. He is dismissed with a probability $p(y)$ which depends on the value of Y at the moment he exercises the option. If y^* is the optimal threshold for the shareholders, then p is assumed to be decreasing for $y < y^*$ and increasing for $y > y^*$, with $p(y^*) = 0$. Finally, if the manager is dismissed, the effect is equivalent to his wealth being scaled down by a multiplicative factor c , $c \in (0, 1)$. (Note that the manager's wealth is not affected directly by Y , except through the possibility of dismissal.) It follows that with $h(y) = (c^{1-R} p(y) + (1 - p(y)))^{\frac{1}{1-R}}$, the problem becomes to find

$$\sup_{\tau \in \mathcal{T}} \sup_{\theta \in \Theta_{0,\tau}^-} \mathbb{E} \left[e^{-\delta \tau} \frac{(X_\tau^\theta)^{1-R} h(Y_\tau)^{1-R}}{1 - R} \right].$$

Hugonnier and Morellec [11] solve (24) in this setting under the assumption that $\delta > \beta$. They conclude that the incompleteness and the risk of control challenges induces the manager to invest early versus what the shareholders would optimally choose. However, the use of $\delta > \beta$ means that the problem they solve has a bias towards stopping times τ which are small. In particular, if $h(y) \equiv 1$ (either because $p(y) \equiv 0$, so that there is no risk of dismissal, or because $c \equiv 1$, so that dismissal has no effect on his wealth), then $\tau = 0$ is the optimal choice of action time. The primary reason that Hugonnier and Morellec find that the agent acts “early” is that they do not use a horizon-unbiased utility function, and not because of the incompleteness or control challenges.

6.3. Extensions

The lump-sum asset sale problem is the simplest of many closely related problems. Variants would be for the agent to receive a dividend income from the real asset before the sale and zero income after the sale, or for the agent to face running costs in order to keep the real asset in a productive and saleable form. In each of these cases, the definitions of admissible wealth processes would change. Alternatively we could assume that the wealth of the agent is augmented by a stochastic income, both before and after the sale of the real asset, and again the self-financing condition would be inappropriate. However, the problems would still remain in the general framework of Section 3. We could also allow for cases where the real asset was divisible, but the decision to sell any part was still irreversible. In this case the unique action time τ would have to be replaced by a selling strategy represented by a function increasing from zero to one.

6.4. Conclusion

This paper has as its subject problems of optimal stopping in which an investor has an asset to sell, but also has outside investment opportunities. In this setting, delaying the sale of the real asset incurs an opportunity cost from the foregone investment opportunities. If the aim is to analyse the agent's behaviour in terms of the optimal choice for the time to sell the real asset, then it is important to design the problem such that if we consider the corresponding investment problem without the real asset, then the agent would have no preference over the choice of time at which her utility was measured. If this is not the case, then any results will be biased towards these preferred stopping times.

Over a finite horizon, it is straightforward to construct horizon-unbiased utility functions by backward induction from the terminal condition. Over the infinite horizon, the issue is much more delicate — even in simple models, horizon-unbiased utilities are associated with solutions of the backward heat equation. However, there exists a family of CRRA horizon-unbiased utility functions. For these utilities, the discount factor is not subjective, but in order to give an unbiased problem, must be chosen to reflect the time-value of money, and the risk-aversion dependent opportunity cost of delaying sale.

Acknowledgements

Partial support from the NSF via grant DMI 0447990 is acknowledged by Vicky Henderson. David Hobson is supported by an EPSRC Advanced Fellowship.

References

- [1] T. Choulli, C. Stricker, Minimal entropy-Hellinger martingale measures in incomplete markets, *Mathematical Finance* 15 (2005) 465–490.
- [2] M.H.A. Davis, Option valuation and basis risk, in: T.E. Djaferis, L.C. Shick. (Eds.), *System Theory: Modelling, Analysis and Control*, Academic Publishers, 1999.
- [3] M.H.A. Davis, Optimal hedging with basis risk, in: Y. Kabanov, R. Liptser, J. Stoyanov (Eds.), *From Stochastic Calculus to Mathematical Finance*, Springer, 2006.
- [4] M.H.A. Davis, T. Zariphopoulou, American options and transactions fees, in: *Mathematical Finance IMA Volumes in Mathematics and its Applications*, Springer-Verlag, 1995.
- [5] A.K. Dixit, R.S. Pindyck, *Investment under Uncertainty*, Princeton University Press, 1994.
- [6] J. Evans, V. Henderson, D. Hobson, Optimal timing for an indivisible asset sale, *Mathematical Finance* 2007 (in press).
- [7] V. Henderson, Valuation of claims on non-traded assets using utility maximization, *Mathematical Finance* 12 (4) (2002) 351–373.
- [8] V. Henderson, Valuing the option to invest in an incomplete market, 2004 (Preprint).
- [9] V. Henderson, D. Hobson, Real options with constant relative risk aversion, *Journal of Economic Dynamics and Control* 27 (2002) 329–355.
- [10] V. Henderson, D. Hobson, Utility indifference pricing: An overview, in: Rene Carmona (Ed.), *Volume on Indifference Pricing*, Princeton University Press, 2007 (in press).
- [11] J. Hugonnier, E. Morellec, Corporate control and real investment in incomplete markets, *Journal of Economic Dynamics and Control* 31 (5) (2007) 1781–1800.
- [12] A. Kadam, P. Lakner, A. Srinivasan, Perpetual call options with non-tradeability, *Optimal Control Applications and Methods* 26 (3) (2005) 107–127.
- [13] I. Karatzas, S.G. Kou, Hedging american contingent claims with constrained portfolios, *Finance and Stochastics* 2 (1998) 215–258.
- [14] I. Karatzas, H. Wang, Utility maximization with discretionary stopping, *SIAM Journal on Control and Optimization* 39 (2001) 306–329.
- [15] R. McDonald, D.R. Siegel, The value of waiting to invest, *Quarterly Journal of Economics* 101 (1986) 707–727.

- [16] R.C. Merton, Optimum consumption and portfolio rules in a continuous time model, *Journal of Economic Theory* 3 (1971) 373–413.
- [17] A. Oberman, T. Zariphopoulou, Pricing early exercise contracts in incomplete markets, *Computational Management Science* 1 (2003) 75–107.