


March 2007

Optimal Stopping under Ambiguity

Frank Riedel



IMW · Bielefeld University
Postfach 100131
33501 Bielefeld · Germany



email: imw@wiwi.uni-bielefeld.de
<http://www.wiwi.uni-bielefeld.de/~imw/Papers/showpaper.php?390>
ISSN: 0931-6558

Optimal Stopping under Ambiguity

FRANK RIEDEL*

Institute for Mathematical Economics
Bielefeld University

March 29, 2007

Abstract

We consider optimal stopping problems for ambiguity averse decision makers with multiple priors. In general, backward induction fails. If, however, the class of priors is time-consistent, we establish a generalization of the classical theory of optimal stopping. To this end, we develop first steps of a martingale theory for multiple priors. We define minimax (super)martingales, provide a Doob–Meyer decomposition, and characterize minimax martingales. This allows us to extend the standard backward induction procedure to ambiguous, time-consistent preferences. The value function is the smallest process that is a minimax supermartingale and dominates the payoff process. It is optimal to stop when the current payoff is equal to the value function. Moving on, we study the infinite horizon case. We show that the value process satisfies the same backward recursion (Bellman equation) as in the finite horizon case. The finite horizon solutions converge to the infinite horizon solution. Finally, we characterize completely the set of time-consistent multiple priors in the binomial tree. We solve two classes of examples: the so-called independent and indistinguishable case (the parking problem) and the case of American Options (Cox–Ross–Rubinstein model).

Key words and phrases: Optimal Stopping, Ambiguity, Uncertainty Aversion

JEL subject classification: D81, C61, G11

*I thank Daniel Engelage and Larry Epstein for comments. Financial Support through the German Research Foundation, Grant Ri 1128–3–1 is gratefully acknowledged.

1 Introduction

Good timing is a key decision in economic environments; whether it is the right time for entering a new market, launching a new brand, going public, exercising an option and so on — in all these cases economic agents have to determine an optimal time for some action. Formally, all these problems can be cast as optimal stopping problems. For a given sequence of possible stochastic payoffs, a decision maker has to choose a random stopping time that maximizes the expected reward.

Traditionally, it has been assumed that the distribution of payoffs is perfectly known to the agent. In such a formulation, the agent faces *risk* as she does not know the possible payoffs ex ante, but does know its probabilistic properties. Now in many real life situations that involve optimal stopping, this assumption seems restrictive. As an example, think of one shot decisions, when, e.g., an entry decision is made for the first time, and no data are available to estimate the distribution of possible profits. Alternatively, one might want to check the robustness of the results derived under the assumption of a unique prior. In that case, we ask if the optimal solution is still approximately optimal if one varies slightly the shape of the distribution.

It seems thus important to study optimal stopping in models that impose less stringent assumptions on the ex ante probabilistic knowledge of agents. In this paper, we adopt the framework of *ambiguity* that allows to distinguish between risk and uncertainty as it has been developed by Gilboa and Schmeidler (1989) and been extended to dynamic settings by Epstein and Schneider (2003b). Accordingly, we assume that the agent has a set of possible prior distributions and evaluates a random payoff by computing the minimal expected value over this class of priors. We thus leave the realm of the Bayesian world. As is well known, one easily runs into dynamic inconsistencies if one does so (Sarin and Wakker (1998), Machina (1989), Yoo (1991), Eichberger and Kelsey (1996)); in the current setting, we also give an example where the naive choice of two priors leads to dynamically inconsistent decisions of stopping (Example 3.1 below). The work of Epstein and Schneider (2003b) shows how to overcome this difficulty. The set of priors must satisfy a certain dynamic consistency condition that they call *rectangularity*. This property appears in other decision-theoretic contexts as well, see, e.g. Delbaen (2002b), Riedel (2004), or Föllmer and Schied (2004). It has also been called *stability under pasting* or *time-consistency*. We go for the last name here.

This paper develops a general theory of optimal stopping under time-consistent ambiguity. We show that much of the classical results are still valid provided that the class of priors is chosen in a time-consistent way. When the horizon is finite, backward induction leads to the optimal solution as in the Bayesian case. One can thus compute easily the value function of the problem, the generalized *Snell envelope* of the payoff sequence. An optimal stopping rule is, as in the classical case, to stop when the payoff from stopping is equal to the Snell envelope.

The proof of these results is not completely straightforward, though. To this end, we develop first steps of a theory of *minimax martingales* in Section 3.2. A minimax martingale (M_t) satisfies the usual martingale property for the nonlinear minimax expectation operator, or

$$M_t = \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [M_{t+1} | \mathcal{F}_t] ,$$

where \mathcal{Q} is the set of time-consistent priors. Similarly, a minimax supermartingale (S_t) satisfies

$$S_t \geq \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [S_{t+1} | \mathcal{F}_t] .$$

A minimax martingale is a *submartingale* for all probability measures $P \in \mathcal{Q}$ and a martingale for some (worst-case) probability measure $P^* \in \mathcal{Q}$. The existence of this worst-case measure requires time-consistency of the set of priors. Intuitively, a minimax martingale models a game against nature that an uncertainty-averse agent would consider as "fair": in the worst case, it is a fair game (martingale) — as a consequence, it must be a favorable game (submartingale) in all other cases. We show that two key results from classical martingale theory hold true for minimax supermartingales: the Doob decomposition and the Optional Sampling Theorem.

The Doob decomposition states that a minimax supermartingale can be written as a minimax martingale minus a predictable, increasing process that starts at 0. While the proof is a copy of the original proof, it is noteworthy that we do not have here a *uniform* decomposition for the class of priors as one obtains it in the *Optional Decomposition Theorem* (Kramkov (1996)). There, one aims to write a uniform supermartingale as a uniform martingale minus some optional increasing process. Minimax martingales are not uniform martingales, in general; thus, the type of decomposition is quite different.

The second theorem that plays a key role here is the preservation of the minimax supermartingale property under stopping, the so-called *Optional Sampling Theorem*. This theorem is a typical *systems theorem*. It says that if you play an unfair game against nature, then you cannot "beat the system" whatever stopping rule you use. Formally, this means that if we start with a minimax supermartingale, then also the stopped process remains a minimax supermartingale. The validity of this theorem hinges critically on the time-consistency of the set of priors.

Having established these two key theorems, one can proceed as in the classical literature (Chow, Robbins, and Siegmund (1971), Snell (1952)). We show that the value function which one defines by backward induction (Bellman principle) is the smallest minimax supermartingale that dominates the payoff process. As long as stopping is not optimal, the value function is a minimax martingale. In discrete time, optimal stopping times are usually not unique. The mentioned rule to stop whenever the payoff equals the value process, is the smallest optimal stopping time. The minimax Doob decomposition allows to determine the largest optimal stopping time as well. We also obtain a duality result that has first been obtained by Föllmer and Schied (2004) and Karatzas and Kou (1998) with different methods. The minimax value function is the lower envelope of all Bayesian value functions. Under our assumptions, the infimum is also attained by some probability measure. As a consequence, the smallest optimal stopping rule in the minimax case is equal to the smallest optimal stopping rule a Bayesian decision maker would choose under some probability measure $P^* \in \mathcal{Q}$.

In Section 5, we extend the theory to infinite horizon where backward induction is not feasible. We show that the value process is still the smallest minimax supermartingale that dominates the payoff process. It satisfies the same recursive (or Bellman) equation as the value process does in the finite horizon case. Moreover, we show that the finite horizon solutions converge to the infinite horizon solution as the horizon tends to infinity. This is important for applications as the finite horizon solution can easily be computed by backward induction. The convergence theorem then allows to approximate the infinite horizon numerically.

In Section 6, we study important classes of economic optimal stopping problems in the binomial tree. Time-consistency of priors is a strong requirement in the binomial tree. We show that a set of priors is time-consistent if and only if the conditional probabilities of moving up at some node stay

between two predictable bounds (\underline{p}_t) and (\bar{p}_t) . If we impose the indistinguishability condition of Epstein and Schneider (2003a), then the bounds (\underline{p}_t) and (\bar{p}_t) must be constant over time.

We then solve two classes of optimal stopping problems in the binomial tree. When the payoff is an increasing function of the current up- or downward move only, then it is optimal to behave as if the lower bound \underline{p} was the probability for moving up. This solves for example the class of *parking problems* where the aim is to stop as close as possible to a given target without knowing whether spots are occupied or not. Similarly, we solve the class of problems where the payoff is a monotone function of the *ambiguous random walk*. This includes the exercise of American Options as well as entry and exit decisions as important special cases.

Decisions under ambiguity are being studied by a number of authors currently. The present paper relies heavily on the fundamental work by Epstein and Schneider (2003b), Delbaen (2002a), and Föllmer and Schied (2004). The duality theorem 3.9 appears in Föllmer and Schied (2004) (derived by other arguments). The notion of a generalized Snell envelope for *maxmax* expected utility (which is easier than the minimax case treated here) appears also in the theory of dynamic coherent risk measures in Artzner, Delbaen, Eber, Heath, and Ku (2002). Another approach can be found in Karatzas and Zamfirescu (2003) (see also Zamfirescu (2003)) who discuss both the maxmax- and the minmax-case and characterize saddle-points. However, they do not assume time-consistency. In the framework of Brownian motion, the concept of g-expectation introduced by Peng (1997) is closely related to minimax expectations. In that framework, Coquet, Hu, Mémin, and Peng (2002) derive a nonlinear Doob-Meyer decomposition. A first application of optimal stopping in continuous time can be found in Nishimura and Ozaki (2007) who solve the optimal stopping problem for an American Option when the drift term is unknown. The corresponding discrete-time result follows from our examples in Section 6.

The aim of the present paper is to present the theory of optimal stopping under ambiguity in a unified and closed form. We have tried to work as closely as possible along the classical lines.

2 Ambiguity and Optimal Stopping

Let $(\Omega, \mathcal{F}, P_0, (\mathcal{F}_t)_{t \in \mathbb{N}})$ be a probability space with a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$. We assume that \mathcal{F}_0 is the trivial σ -field and that \mathcal{F} is the σ -field generated by the union of all $\mathcal{F}_t, t \in \mathbb{N}$.

Let $(X_t)_{t \in \mathbb{N}}$ be an adapted process that describes the payoff from stopping. We assume throughout that (X_t) is bounded.

Assumption 2.1 *The payoff process X is bounded.*

The decision maker chooses a stopping time τ with values in $\mathbb{N} \cup \{\infty\}$ of the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$. From stopping she obtains a payoff $X_\tau(\omega) = X_{\tau(\omega)}(\omega)$ for $\omega \in \Omega$, and we set $X_\tau(\omega) = 0$ if $\tau(\omega) = \infty$. She aims to maximize the expected reward; as she is uncertain about the distribution of X , she uses a class \mathcal{Q} of probability measures on (Ω, \mathcal{F}) . The (minimax) expected reward is thus given by

$$\inf_{P \in \mathcal{Q}} \mathbb{E}^P X_\tau. \quad (1)$$

Without loss of generality, we can assume \mathcal{Q} to be convex.

We impose the following assumption.

Assumption 2.2 *All $P \in \mathcal{Q}$ are locally equivalent to the reference measure P_0 , i.e. for all $t \in \mathbb{N}$, and $A \in \mathcal{F}_t$ we have $P(A) = 0$ if and only if $P_0(A) = 0$.*

The reference measure P_0 just serves the role of fixing the sets of measure zero. Economically, this means that the decision maker has perfect knowledge about sure events. In a discrete model, one can take P_0 to be the uniform distribution without loss of generality. More generally, if the measurable space (Ω, \mathcal{F}) has a nice topological structure, and the minimal expectation as in (1) is continuous from below, one can always construct P_0 from the set of priors \mathcal{Q} , see Tutsch (2006).

Mathematically, it might be possible to extend the theory to classes \mathcal{Q} that are only absolutely continuous with respect to the reference measure \mathbb{P}_0 . As this comes at a high technical cost, I prefer to assume equivalence. Economically, the assumption just excludes the case in which some prior assigns a probability of zero to an event that can occur with positive probability under the reference measure P_0 . We think that it is plausible to exclude this degenerate case. A behavioral foundation for this assumption can be found in Epstein and Marinacci (2006).

Note that we assume *local* equivalence of the priors only. It would not be reasonable to assume that all measures in \mathcal{Q} are equivalent to the reference measure P_0 on the information up to ∞ given by the σ -field $\mathcal{F} = \mathcal{F}_\infty = \sigma(\bigcup_{t=0}^\infty \mathcal{F}_t)$. As an example, assume that the sequence (X_t) is independent and identically distributed with different mean values $m^P \neq m^Q$ under two measures $P, Q \in \mathcal{Q}$. By the Law of Large Numbers, the arithmetic mean converges P -almost surely to m^P , and Q -almost surely to m^Q . Hence, the measures are even singular on \mathcal{F}_∞ . Therefore, we only impose the assumption of equivalence locally for finite times t .

Due to our assumptions and the Radon–Nikodym theorem, the set of priors \mathcal{Q} can be identified with the set of density processes $\mathcal{D} = \left\{ \left(\frac{dP}{dP_0} \Big|_{\mathcal{F}_t} \right)_{t \in \mathbb{N}} \mid P \in \mathcal{Q} \right\}$. We impose the following technical condition that ensures that the infimum in (1) is always attained for bounded stopping times, see Lemma B.2.

Assumption 2.3 *For every $t \in \mathbb{N}$, the family of densities*

$$\mathcal{D}_t = \left\{ \frac{dP}{dP_0} \Big|_{\mathcal{F}_t} \mid P \in \mathcal{Q} \right\}$$

is weakly compact in $L^1(\Omega, \mathcal{F}, P_0)$.

Note that the assumption is satisfied without loss of generality when the densities in \mathcal{D}_t are bounded by a P_0 -integrable random variable. In particular, the assumption is satisfied whenever the state space Ω is finite.

The assumption is equivalent to certain monotone continuity conditions, see Corollary 4.35 in Föllmer and Schied (2004) or Chateauneuf, Maccheroni, Marinacci, and Tallon (2005), and also Lemma B.1 in the Appendix. A behavioral description for such kind of continuity has been given by Arrow (1971).

3 Finite Horizon: Backward Induction and Time Consistency

The problem we consider in this section is

$$\text{maximize } \inf_{P \in \mathcal{Q}} \mathbb{E}^P X_\tau \text{ over all stopping times } \tau \leq T$$

for a finite horizon $T < \infty$.

For standard expectations, i.e. $\mathcal{Q} = \{P\}$ a singleton, the general solution to the above problem is well known¹. One proceeds by backward induction and defines $U_T^P = X_T$, the value in the last period. By backward induction, set for $t < T$

$$U_t^P := \max \{X_t, \mathbb{E}^P [U_{t+1} | \mathcal{F}_t]\} .$$

Then the value process (U_t^P) is the smallest P -supermartingale that dominates the payoff process (X_t) , and an optimal stopping time is given by

$$\tau^* = \inf \{t \geq 0 : U_t^P = X_t\} .$$

The process U^P is called the Snell envelope of X under P .

3.1 Time–Consistency

The following example shows that backward induction fails in general for minimax expected utility.

Example 3.1 *Consider a two–period binomial tree as in Figure 3.1. Let X_0, X_1, X_2 be the sequence of payoffs. We take $X_0 = x$,*

$$X_1 = \begin{cases} 3 & \text{after up} \\ 1 & \text{after down} \end{cases}$$

and

$$X_2 = \begin{cases} 0 & \text{after up, up} \\ 6 & \text{after up, down} \\ 6 & \text{after down, up} \\ 0 & \text{after down, down} \end{cases} .$$

The decision maker believes that the up and down moves are independent and identically distributed. She uses two priors. Under the first prior, one moves up with probability 1/3 in all nodes, whereas under the second prior, one moves up with probability 2/3 in all nodes, see Figure 3.1. If we use backward

induction, the value at time 2 is $U_2 = \begin{pmatrix} 0 \\ 6 \\ 6 \\ 0 \end{pmatrix}$. At time 1, the minimal

¹The theory starts with Snell (1952); for textbook accounts, see Chow, Robbins, and Siegmund (1971), or Dixit and Pindyck (1994) that contains many important economic applications.

conditional expected payoff in the upper node is achieved for the probability $2/3$ with a value of 2. From stopping, we get 3. Hence, backward induction prescribes to stop in this node. Similarly, in the lower node, we obtain a value of 2. Finally, at time 0, the value deduced by backward induction is $U_0 = \max\{x, 7/3\}$. Hence, if $x \geq 7/3$, backward induction prescribes to stop immediately and one obtains a value of x . On the other hand, consider what happens if one does not stop at all. Then the ex ante minimax expected payoff is

$$\min\{1/9 \cdot 0 + 2/9 \cdot 6 + 2/9 \cdot 6 + 4/9 \cdot 0, 4/9 \cdot 0 + 2/9 \cdot 6 + 2/9 \cdot 6 + 1/9 \cdot 0\} = 8/3.$$

Hence, if $7/3 \leq x < 8/3$, we conclude that backward induction does not lead to the ex ante optimal solution. One checks easily that the ex ante optimal decision is to wait until time 2 while backward induction would prescribe to stop immediately.

Backward induction fails in the above example because the preferences of the agent are not time-consistent. This issue has recently received much attention in the decision theory literature and also in Mathematical Finance, see Epstein and Schneider (2003b), Artzner, Delbaen, Eber, Heath, and Ku (2002), Riedel (2004), Detlefsen and Scandolo (2005). There, it is shown that minimax EU preferences (and coherent dynamic risk measures) are time-consistent if and only if the set of priors satisfies a certain condition that has been called rectangularity, stability under pasting, or time-consistency. We are going to impose this property in the following.

Assumption 3.2 *The set of priors \mathcal{Q} is time-consistent in the following sense. For P and Q in \mathcal{Q} , let (p_t) and (q_t) be the density processes of P resp. Q with respect to P_0 , i.e.*

$$p_t = \left. \frac{dP}{dP_0} \right|_{\mathcal{F}_t},$$

and analogously for Q . Fix some stopping time τ . Define a new probability measure R by setting for all $t \in \mathbb{N}$

$$\left. \frac{dR}{dP_0} \right|_{\mathcal{F}_t} = \begin{cases} p_t & \text{if } t \leq \tau \\ \frac{p_\tau q_t}{q_\tau} & \text{else} \end{cases}. \quad (2)$$

Then R belongs to \mathcal{Q} as well.

Note that the measure R above is well defined as the density process q is strictly positive by Assumption 2.2. The above definition of time-consistency is taken from Delbaen (2002b). It may look different to the definition of *rectangularity* used in Epstein and Schneider (2003b). They are equivalent, though. The appendix discusses another equivalent definition given by Föllmer and Schied (2004),

The assumption ensures that the set of priors is closed under the operation of pasting together marginal and conditional distributions. In fact, if the decision maker uses the measure Q until time τ and evaluates expectations after τ according to \hat{Q} , then the newly constructed expectation is still in her set of priors \mathcal{Q} . In the binomial example above, the set of priors is not time-consistent as it does not contain the probability measure under which we first go up with probability $1/3$ and in the second period the probability of an upward move is $2/3$. It has been shown (see Epstein and Schneider (2003b), more generally Delbaen (2002b), Theorem 6.2. and 8.2.) that under Assumption 2.3 \mathcal{Q} is time-consistent if and only if we have for all bounded random variables Z the following version of the law of iterated expectations:

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [Z | \mathcal{F}_t] = \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P \left[\operatorname{ess\,inf}_{Q \in \mathcal{Q}} \mathbb{E}^Q [Z | \mathcal{F}_{t+1}] \middle| \mathcal{F}_t \right] \quad (t \in \mathbb{N}). \quad (3)$$

3.2 Minimax Martingale Theory

This section sketches the beginning of a martingale theory for time-consistent multiple priors that we are going to need in the following. The material might be useful in other contexts as well. To facilitate reading, we have put all proofs into the appendix. Remember that we impose throughout the paper the Assumptions 2.2, 2.3, and 3.2.

Definition 3.3 *Let \mathcal{Q} be a set of priors. Let $(M_t)_{t \in \mathbb{N}}$ be an adapted process with $\mathbb{E}^P |M_t| < \infty$ for all $P \in \mathcal{Q}$ and $t \in \mathbb{N}$. (M_t) is called a minimax (sub-, super-)martingale with respect to \mathcal{Q} if we have for $t \in \mathbb{N}$*

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [M_{t+1} | \mathcal{F}_t] = (\geq, \leq) M_t.$$

It is important to distinguish this concept from the different notion of a \mathcal{Q} -martingale. A \mathcal{Q} -(sub- or super-)martingale is a process that is a (sub- or super-)martingale for all priors in \mathcal{Q} simultaneously. The concepts are related as follows.

Lemma 3.4 *Let (M_t) be a bounded, adapted process.*

1. *M is a minimax submartingale if and only if it is a \mathcal{Q} -submartingale,*
2. *M is a minimax supermartingale if and only there exists $P^* \in \mathcal{Q}$ such that M is a P^* -supermartingale,*
3. *M is a minimax martingale with respect to \mathcal{Q} if and only if*
 - (a) *there exists $P^* \in \mathcal{Q}$ such that M is a P^* -martingale and*
 - (b) *M is a \mathcal{Q} -submartingale.*

Note the big difference between minimax sub- and supermartingales. While a minimax submartingale is a submartingale for all $Q \in \mathcal{Q}$ uniformly, a minimax supermartingale is a supermartingale for *some* $Q \in \mathcal{Q}$ only. This is due, of course, to the fact that we take always the essential infimum over a class of probability measures.

We are now going to extend two fundamental theorems from martingale theory to minimax martingales. We start with the famous Doob decomposition.

Theorem 3.5 (Doob Decomposition) *Let S be a bounded minimax supermartingale (submartingale) with respect to \mathcal{Q} . Then there exists a minimax martingale M and a predictable, nondecreasing process A with $A_0 = 0$ such that $S = M - A$ ($S = M + A$). Such a decomposition is unique.*

In other words, every game against nature that is regarded as unfair under ambiguity (a minimax supermartingale) can be written as a fair game under ambiguity (a minimax martingale) minus some cumulative payments.

The second fundamental theorem concerns the preservation of the (super)-martingale property under optimal stopping. It is a version of the famous folk theorem on unfair games: if you play an unfair game against nature, then you cannot obtain a positive payoff even if you use a fancy exit strategy. More formally, it means that a minimax supermartingale stays a minimax supermartingale when it is stopped at some random stopping time. The validity of this theorem relies on time-consistency of the set of priors.

Theorem 3.6 (Optional Sampling Theorem) *Let Z be a bounded minimax supermartingale with respect to \mathcal{Q} . Let $\sigma \leq \tau$ be stopping times. Assume*

that τ is universally finite in the sense that $P[\tau < \infty] = 1$ for all $P \in \mathcal{Q}$. Then

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [Z_\tau | \mathcal{F}_\sigma] \leq Z_\sigma.$$

3.3 Generalized Snell Envelope, Optimal Stopping Times, and Duality

We show now that backward induction solves the optimal stopping problem for time-consistent sets of priors.

Theorem 3.7 *Define the minimax Snell envelope of X with respect to \mathcal{Q} recursively by $U_T = X_T$ and*

$$U_t = \max \left\{ X_t, \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_{t+1} | \mathcal{F}_t] \right\} \quad (t = 0, \dots, T-1). \quad (4)$$

Then

- (i) U is the smallest minimax supermartingale with respect to \mathcal{Q} that dominates X ,
- (ii) U is the value process of the optimal stopping problem under ambiguity, i.e.

$$U_t = \operatorname{ess\,sup}_{\tau \geq t} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t],$$

- (iii) an optimal stopping rule is given by

$$\tau^* = \inf \{t \geq 0 : U_t = X_t\}.$$

PROOF: U is a minimax supermartingale by definition. Let V be another minimax supermartingale with $V \geq X$. Then we have $V_T \geq X_T = U_T$. Now assume that $V_{t+1} \geq U_{t+1}$; as V is a minimax supermartingale,

$$V_t \geq \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [V_{t+1} | \mathcal{F}_t] \geq \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [U_{t+1} | \mathcal{F}_t].$$

We also have $V_t \geq X_t$ by assumption. Hence

$$V_t \geq \max \left\{ X_t, \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_{t+1} | \mathcal{F}_t] \right\} = U_t.$$

Thus, U is the smallest minimax supermartingale that dominates X .

Now let $W_t = \text{ess sup}_{\tau \geq t} \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t]$. From the minimax supermartingale property of U and $U \geq X$, we conclude with the help of the Optional Sampling Theorem 3.6 that for every stopping time τ with values in $\{t, \dots, T\}$,

$$\text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t] \leq \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_\tau | \mathcal{F}_t] \leq U_t.$$

$W_t \leq U_t$ follows.

It remains to be shown that $U_t \leq W_t$. To this end, we define the stopping time

$$\tau_t^* = \inf \{s \geq t : U_s = X_s\}.$$

We claim that $(U_{s \wedge \tau_t^*})_{s=t, \dots, T}$ is a minimax martingale. It then follows that

$$U_t = \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_{\tau_t^*} | \mathcal{F}_t] = \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_t^*} | \mathcal{F}_t] \leq W_t,$$

and we are done. To check the minimax martingale property, fix $s \in \{t, \dots, T\}$. Note that on the set $\{\tau_t^* \leq s\}$, we have $U_{(s+1) \wedge \tau_t^*} = U_{\tau_t^*} = U_{s \wedge \tau_t^*}$. Hence,

$$\text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_{(s+1) \wedge \tau_t^*} | \mathcal{F}_t] = U_{s \wedge \tau_t^*}$$

on the set $\{\tau_t^* \leq s\}$. On the complement $\{\tau_t^* > s\}$, we have $U_s > X_s$. The definition of U implies that

$$\begin{aligned} U_{s \wedge \tau_t^*} &= U_s = \max \left\{ X_s, \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_{s+1} | \mathcal{F}_s] \right\} \\ &= \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_{s+1} | \mathcal{F}_s] = \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [U_{(s+1) \wedge \tau_t^*} | \mathcal{F}_s]. \end{aligned}$$

Hence, $(U_{s \wedge \tau_t^*})_{s=t, \dots, T}$ is a minimax martingale, and the above claim is proved.

As a further consequence, we obtain for $t = 0$ that $(U_{s \wedge \tau^*})$ is a minimax martingale for

$$\tau^* = \inf \{t \geq 0 : U_t = X_t\}.$$

Hence, the Optional Sampling Theorem 3.6 yields $U_0 = \inf_{P \in \mathcal{Q}} \mathbb{E}^P X_{\tau^*}$. This shows that τ^* is optimal. \square

Remark 3.8 (i) *One might wonder where time consistency of \mathcal{Q} was used in the proof. We need it when using the Optional Sampling Theorem. This theorem does not hold true without time consistency.*

(ii) *Optimal stopping times are usually not unique. By using the Doob decomposition of the Snell envelope $U = M - A$, one can show that the largest optimal stopping time is*

$$\tau^{max} = \inf \{t \geq 0 : A_{t+1} > 0\} .$$

The above theorem gives a complete solution to the optimal stopping problem under ambiguity. One might next wish to study the relationship between the minimax Snell envelope U and the usual Snell envelopes U^P for the individual priors $P \in \mathcal{Q}$. At time $T - 1$, there is a worst prior $P^{T-1} \in \mathcal{Q}$ such that

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [X_T | \mathcal{F}_{T-1}] = E^{P^{T-1}} [X_T | \mathcal{F}_{T-1}] ,$$

and therefore $U_{T-1} = U_{T-1}^{P^{T-1}}$. At time $T - 2$, there is a measure $P^{T-2} \in \mathcal{Q}$ such that

$$U_{T-2} = \max\{X_{T-2}, \mathbb{E}^{P^{T-2}} [U_{T-1} | \mathcal{F}_{T-2}]\} = \max\{X_{T-2}, \mathbb{E}^{P^{T-2}} [U_{T-1}^{P^{T-1}} | \mathcal{F}_{T-2}]\} .$$

Now time consistency allows us to paste P^{T-1} and P^{T-2} together to obtain a new measure $Q^{T-2} \in \mathcal{Q}$ in such a way that $Q^{T-2} = P^{T-2}$ on \mathcal{F}_{T-1} and the conditional probability of Q^{T-2} given \mathcal{F}_{T-1} is equal to the conditional probability of P^{T-1} given \mathcal{F}_{T-1} . We then get

$$U_{T-1}^{P^{T-1}} = U_{T-1}^{Q^{T-2}}$$

and also

$$U_{T-2} = \max\{X_{T-2}, \mathbb{E}^{Q^{T-2}} \max\{X_{T-2}, \mathbb{E}^{Q^{T-2}} [U_{T-1}^{Q^{T-2}} | \mathcal{F}_{T-2}]\}\} = U_{T-2}^{Q^{T-2}} .$$

Continuing in this manner by backward induction, we conclude that there exists a *worst case measure* $\underline{P} \in \mathcal{Q}$ such that $U = U^{\underline{P}}$, see the proof of Lemma 3.4 for the rigorous construction. We have thus derived the following minimax theorem originally obtained by Föllmer and Schied (2004), and by Karatzas and Kou (1998) for American Options in continuous time.

Theorem 3.9 (Duality) *The minimax Snell envelope U constructed in Theorem 3.7 is the lower envelope of the individual Snell envelopes U^P :*

$$U_t = \operatorname{ess\,inf}_{P \in \mathcal{Q}} U_t^P.$$

The essential infimum is attained by some measure $\underline{P} \in \mathcal{Q}$, i.e. $U = U^{\underline{P}}$. We have the minimax identity

$$\operatorname{ess\,sup}_{\tau \geq t} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t] = \operatorname{ess\,inf}_{P \in \mathcal{Q}} \operatorname{ess\,sup}_{\tau \geq t} \mathbb{E}^P [X_\tau | \mathcal{F}_t].$$

The preceding theorem can be viewed as an *equivalence theorem*: for a given payoff process X , the ambiguity averse decision maker behaves like an expected utility maximizer for a certain *worst case measure* \underline{P} . This does not imply, however, that optimal stopping under ambiguity aversion is behaviorally indistinguishable from optimal stopping under expected utility. In general, the worst case measure \underline{P} depends on the payoff process X . For suitably constructed different payoff processes, the ambiguity averse decision maker behaves like two distinct expected utility maximizers. This makes it possible to distinguish behaviorally between ambiguity averse and ambiguity neutral (EU) decision makers.

4 Infinite Time Horizon

Many optimal stopping problems are naturally formulated without imposing a finite time horizon. Also, the infinite horizon case frequently leads to simpler closed form solutions that are usually not available in the finite horizon case. We thus extend the analysis of the preceding section to $T = \infty$. We show that the value function satisfies the same Bellman-type backward recursion as in the finite case. Again, it is optimal to stop when the current payoff is equal to the value function. Moreover, we establish that the solutions of the finite time horizon converge to the infinite horizon solution. This is important as it allows to approximate the general solution by using the constructive algorithm available in the finite horizon case.

The problem we consider in this section is

$$\begin{aligned} &\text{maximize } \inf_{P \in \mathcal{Q}} \mathbb{E}^P X_\tau \text{ over all stopping times } \tau \text{ that are universally finite,} \\ &\text{i.e. } \inf_{P \in \mathcal{Q}} P[\tau < \infty] = 1. \end{aligned}$$

As we cannot use backward induction as in the finite horizon case, we define the value function at time t as

$$V_t = \operatorname{ess\,sup}_{\tau \geq t} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t] .$$

Note that V_t is well defined and finite because X is bounded.

Theorem 4.1 (i) V is the smallest minimax supermartingale with respect to \mathcal{Q} that dominates X ,

(ii) The value process (V_t) satisfies the Bellman principle

$$V_t = \max \left\{ X_t, \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [V_{t+1} | \mathcal{F}_t] \right\}$$

for all $t \geq 0$,

(iii) an optimal stopping rule is given by

$$\tau^* = \inf \{t \geq 0 : U_t = X_t\} ,$$

provided that τ^* is universally finite.

PROOF: We start with (ii). By Lemma B.3, there exists a sequence (τ_k) of stopping times such that

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_k} | \mathcal{F}_{t+1}] \uparrow V_{t+1} .$$

Continuity from below (Lemma B.1) and time-consistency (3) imply that

$$\begin{aligned} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [V_{t+1} | \mathcal{F}_t] &= \lim_{k \rightarrow \infty} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P \left[\operatorname{ess\,inf}_{Q \in \mathcal{Q}} \mathbb{E}^Q [X_{\tau_k} | \mathcal{F}_{t+1}] \middle| \mathcal{F}_t \right] \\ &= \lim_{k \rightarrow \infty} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_k} | \mathcal{F}_t] \leq V_t . \end{aligned}$$

As $X_t \leq V_t$ is clear, we obtain

$$\max \left\{ X_t, \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [V_{t+1} | \mathcal{F}_t] \right\} \leq V_t .$$

For the converse inequality, take some stopping time $\tau \geq t$ and define a new stopping time $\sigma = \max\{\tau, t+1\} \geq t+1$. Then

$$\begin{aligned}
\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t] &= X_t 1_{\{\tau=t\}} + \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau 1_{\{\tau \geq t+1\}} | \mathcal{F}_t] \\
&= X_t 1_{\{\tau=t\}} + \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\sigma 1_{\{\tau \geq t+1\}} | \mathcal{F}_t] \\
&= X_t 1_{\{\tau=t\}} + \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\sigma | \mathcal{F}_t] 1_{\{\tau \geq t+1\}} \\
&= X_t 1_{\{\tau=t\}} + \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P \left[\operatorname{ess\,inf}_{Q \in \mathcal{Q}} E^Q [X_\sigma | \mathcal{F}_{t+1}] \middle| \mathcal{F}_t \right] 1_{\{\tau \geq t+1\}} \\
&\leq \max\{X_t, \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [V_{t+1} | \mathcal{F}_t]\}.
\end{aligned}$$

This proves (ii). As a consequence, (V_t) is a minimax supermartingale. Now suppose that (W_t) is another minimax supermartingale that dominates X . As X is bounded, we can assume without loss of generality that W is bounded. (Else consider $(\min\{W_t, K\})$ for a sufficiently large number $K > 0$.) Then for every stopping time $\tau \geq t$, the Optional Sampling Theorem 3.6 implies that

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t] \leq \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [W_\tau | \mathcal{F}_t] \leq W_t.$$

By taking the supremum over all such stopping times, $V_t \leq W_t$ follows. This shows (i).

For (iii), one shows first that $(U_{s \wedge \tau^*})$ is a minimax martingale, see the proof of Theorem 3.7. If τ^* is universally finite, bounded convergence (Lemma B.1, 4.) gives

$$\inf_{P \in \mathcal{Q}} \mathbb{E}^P U_{\tau^*} = \lim_{T \rightarrow \infty} \inf_{P \in \mathcal{Q}} \mathbb{E}^P U_{T \wedge \tau^*} = U_0.$$

Hence, τ^* is optimal. \square

The above theorem characterizes nicely the value process for an infinite horizon stopping problem. In contrast to the finite time horizon, it does not provide a constructive algorithm to compute the value process, though. It is thus important to know that the Snell envelopes of the finite horizon models converge to the infinite horizon value.

Theorem 4.2 (Finite Horizon Approximation) *Denote by U^T the minimax Snell envelope of X with time horizon T . Then $\lim_{T \rightarrow \infty} U_t^T = V_t$ for all $t \geq 0$.*

PROOF: Note that U_t^T is bounded and increasing in T . Hence, we can define $U_t^\infty = \lim_{T \rightarrow \infty} U_t^T$. By continuity from below (Lemma B.1) and the definition of the Snell envelope, we obtain

$$U_t^\infty = \lim_{T \rightarrow \infty} \max\{X_t, \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[U_{t+1}^T | \mathcal{F}_t]\} = \max\{X_t, \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[U_{t+1}^\infty | \mathcal{F}_t]\}.$$

Hence, U^∞ is a minimax supermartingale that dominates X . By Theorem 4.1, we have $U^\infty \geq V$. On the other hand, by Theorem 3.7, $U_t^T = \operatorname{ess\,sup}_{t \leq \tau \leq T} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[X_\tau | \mathcal{F}_t] \leq V_t$. As a consequence, $V = U^\infty$, and the proof is complete. \square

5 Examples

We now apply our previous theory to two important classes of examples. The first example has a payoff structure that is independent and indistinguishably distributed. This generalizes the well-known case of independent and identically distributed payoffs to the ambiguous framework. The second class of examples is concerned with sums of independent random variables as they occur typically in financial models. Both classes of models use a binomial tree. Hence, we first characterize time-consistent sets of priors in these trees.

5.1 Time-Consistency in the Binomial Tree

In this section, we completely characterize the sets of priors that satisfy our assumptions for the benchmark model of the binomial tree.

We model the binomial tree by fixing a probability space $(\Omega, \mathcal{F}, P_0)$ on which we have a sequence $(X_t)_{t=0,1,2,\dots}$ of random variables that are independent and identically distributed under P_0 with $P_0(X_t = 1) = P_0(X_t = 0) = 1/2$. The uniform measure P_0 serves only the role of a reference measure here. Let (\mathcal{F}_t) be the natural filtration of (X_t) . We now characterize the sets of priors which satisfy our assumptions².

Theorem 5.1 *In the binomial model, a set of priors \mathcal{Q} satisfies Assumptions 2.2, 2.3, and 3.2 if and only if there exist two predictable processes (\underline{p}_t) and*

²A similar result is mentioned without proof in Artzner, Delbaen, Eber, Heath, and Ku (2002).

(\bar{p}_t) with $0 < \underline{p}_t \leq \bar{p}_t < 1$ such that

$\mathcal{Q} = \{P \mid P \text{ has a density process with respect to } P_0 \text{ of the form}$

$$\left. \frac{dP}{dP_0} \right|_{\mathcal{F}_t} = 2^t \prod_{s=1}^t \alpha_s^{X_s} (1 - \alpha_s)^{1-X_s} \quad (5)$$

$\text{for a predictable process } (\alpha_t) \text{ with } \underline{p}_t \leq \alpha_t \leq \bar{p}_t \}$

The proof is given in the appendix.

Time-consistency in the binomial tree means that we specify an interval $[\underline{p}_t, \bar{p}_t]$ at time $t - 1$ in which the probability of moving up at time t has to lie. Note that these bounds can be stochastic in general. If we impose stationarity, the bounds become constant. This concept has been introduced by Epstein and Schneider (2003a) — they call it independent and indistinguishably distributed random variables and prove a strong Law of Large Numbers. In the binomial tree, the following corollary characterizes all \mathcal{Q} that satisfy the Epstein–Schneider requirement. In the terminology introduced in Epstein and Schneider (2003b), we have then an *ambiguous random walk*.

Corollary 5.2 *\mathcal{Q} is a model for an independently and indistinguishably distributed random walk if and only if the bounds (\underline{p}_t) and (\bar{p}_t) of Theorem 5.1 are constant numbers in $(0, 1)$.*

PROOF: Indistinguishability is defined as time stationarity of the conditional distributions in Epstein and Schneider (2003a). Formally, we must have for all $t \in \mathbb{N}$ $P[X_{t+1} = 1 | \mathcal{F}_t] \in [\underline{p}, \bar{p}]$ for fixed numbers $\underline{p}, \bar{p} \in [0, 1]$. \square

5.2 The IID Case and the Cox–Ross–Rubinstein Model

We now discuss two classes of optimal stopping problems in the binomial tree.

Let (Z_t) be a sequence of binary random variables and \mathcal{Q} a family of priors as in Corollary 5.2. In particular, \underline{p} and \bar{p} are the lower resp. upper bound for the conditional probabilities $P[Z_t = 1 | Z_1, \dots, Z_{t-1}] \in [\underline{p}, \bar{p}]$ for all $P \in \mathcal{Q}$. We denote by \underline{P} and \bar{P} the probability measures under which

(Z_t) is identically and independently distributed with $\underline{P}[Z_t = 1] = \underline{p}$ and $\overline{P}[Z_t = 1] = \overline{p}$ resp.

5.2.1 The Indistinguishable Case and the Parking Problem

We consider the case where the payoff can be written as a (possibly time-dependent) increasing function of Z_t , i.e. $X_t = g(t, Z_t)$, and g is increasing in the second variable. When g does not depend on time t , the payoff sequence is independently and indistinguishably distributed (Epstein and Schneider (2003a)). A famous special case is given by the following example.

Example 5.3 *The Parking Problem (see Chow, Robbins, and Siegmund (1971), and Lerche, Keener, and Woodroffe (1994) for a generalization). You are driving along the Rhine. Your aim is to park your car as close as possible to the place where the ship leaves for a sightseeing tour. When a spot is empty, you face the decision whether to stop and park, or to continue hoping to find a spot closer to the departure point. Formally, let $N \in \mathbb{N}$ be the desired parking spot. The spot k is empty when $Z_k = 1$. The payoff from parking at an empty spot is $-|N - k|$. If you stop at an occupied spot, you pay a fee K (assumed to be that large that it is never optimal to stop at an occupied spot). Traditionally, it has been assumed that the probability $p = P[Z_t = 1]$ is known to the driver. We allow for some ambiguity here.*

Theorem 5.4 *Let (\underline{U}_t^T) be the Snell envelope of $X_t = g(t, Z_t)$ under \underline{P} for a horizon $T > 0$. Then the minimax Snell envelope is $U^T = \underline{U}^T$, and an optimal stopping rule under ambiguity is given by*

$$\underline{\tau}^T = \inf\{t \geq 0 : g(t, Z_t) = \underline{U}_t^T\}.$$

The same holds true for an infinite time horizon provided that $\underline{\tau}^\infty = \inf\{t \geq 0 : g(t, Z_t) = \underline{U}_t^\infty\}$ is universally finite.

PROOF: The infinite horizon result follows from the approximation theorem 4.2 once we have established the result for finite horizon. So let $T > 0$. We prove by backward induction that $U_t^T = u(t, Z_t)$ for a function $u(t, z)$ that is increasing in z and equal to the Snell envelope under the measure \underline{P} . We clearly have $U_T = g(T, Z_T) = \underline{U}_T^T$, and the claim is thus valid for $t = T$. We have for $t < T$

$$U_t^T = \max \left\{ g(t, Z_t), \min_{p_{t+1} \in [\underline{p}, \overline{p}]} (p_{t+1} u(t+1, 1) + (1 - p_{t+1}) u(t+1, 0)) \right\}$$

By induction hypothesis, $u(t+1, 1) \geq u(t+1, 0)$, thus

$$U_t^T = \max \{g(t, Z_t), (\underline{p}u(t+1, 1) + (1 - \underline{p})u(t+1, 0))\} = \underline{U}_t^T.$$

□

Example 5.5 *The Parking Problem ctd.* The previous theorem tells us that an ambiguity averse driver should behave as if the lowest probability \underline{p} was the correct one. The solution to this Bayesian problem is well known (see, e.g., Ferguson (2006), Chapter 2.11). Let $r \in \mathbb{N}$ be the smallest number such that $(1 - \underline{p})^{r+1} \leq 1/2$. The optimal rule is to start looking when you are r places away from the desired location and to take the first available spot. If, e.g., you think that in the worst case one out of one hundred places is empty, you should start looking when you are 68 places from your target.

5.2.2 Ambiguous Asset Markets and Optimal Exercise of American Options

In the binomial model of asset markets (Cox, Ross, and Rubinstein (1979)), there is a riskless asset with price $B_t = (1 + r)^t$ for an interest rate $r > -1$, and a risky asset (S_t) given by $S_0 = 1$ and

$$S_{t+1} = S_t \cdot \begin{cases} (1 + b) & \text{if } Z_{t+1} = 1 \\ (1 + a) & \text{if } Z_{t+1} = 0 \end{cases}.$$

To preclude arbitrage opportunities, we assume $-1 < a < r < b$. We consider an investor who exercises an American Option that pays off $A(t, S_t)$ when exercised at time t . We assume that $A(t, \cdot)$ is increasing and bounded. In our model, the investor perceives the risky asset as ambiguous as he does not know the exact distribution of S_t .

Example 5.6 *A risk-neutral buyer of an American Put has $A(t, s) = e^{-rt} \max\{K - s, 0\}$. Our model allows also to include risk aversion. For example, a buyer of an American Call with constant absolute risk aversion maximizes the expected payoff $A(t, s) = -\exp(-\rho t - \alpha \max\{s - K, 0\})$ for some subjective discount rate ρ and risk aversion $\alpha > 0$.*

Theorem 5.7 *Let (\underline{U}_t) be the Snell envelope of $X_t = A(t, S_t)$ under \underline{P} for a horizon $T > 0$. Then the minimax Snell envelope is $U^T = \underline{U}^T$, and an optimal*

stopping rule under ambiguity is given by $\underline{\tau}^T = \inf\{t \geq 0 : A(t, S_t) = \underline{U}_t^T\}$. The same holds true for an infinite time horizon provided that $\underline{\tau}^\infty = \inf\{t \geq 0 : A(t, S_t) = \underline{U}_t^\infty\}$ is universally finite.

PROOF: The infinite horizon result follows from the approximation theorem 4.2 once we have established the result for finite horizon. So let $T > 0$. We prove by backward induction that $U_t^T = u(t, Z_t)$ for a function $u(t, z)$ that is increasing in z ; moreover $u(t, Z_t) = \underline{U}_t^T$. We clearly have $U_T = A(T, Z_T) = \underline{U}_T^T$, and the claim is thus valid for $t = T$. We have for $t < T$

$$U_t^T = \max \left\{ A(t, Z_t), \min_{p_{t+1} \in [\underline{p}, \bar{p}]} (p_{t+1}u(t+1, S_t(1+b)) + (1-p_{t+1})u(t+1, S_t(1+a))) \right\}.$$

By induction hypothesis, $u(t+1, S_t(1+b)) \geq u(t+1, S_t(1+a))$, thus

$$\begin{aligned} U_t^T &= \max \left\{ g(t, Z_t), (\underline{p}u(t+1, S_t(1+b)) + (1-\underline{p})u(t+1, S_t(1+a))) \right\} \\ &= \underline{U}_t^T. \end{aligned}$$

□

From inspection of the proofs, one sees that both Theorem 5.4 and 5.4 rely on the fact that \underline{P} is the worst probability measure in the sense of first-order stochastic dominance.

6 Conclusion

We present a unified and general theory of optimal stopping under ambiguity in discrete time. Much of the received theory can be translated to the multiple priors framework provided the priors satisfy the time consistency criterion. In this case, it seems also possible to generalize much of classical martingale theory. A natural next step is, of course, to extend these results to continuous time. Recent work also shows that one might generalize our results to the more general class of dynamic variational preferences (Maccheroni, Marinacci, and Rustichini (2006)) or convex risk measures (Föllmer and Penner (2007)).

A Equivalent Descriptions of Time-Consistency

Several notions of time-consistency have been introduced in the literature. For the sake of the reader and our own convenience, we gather them here, and prove that they are equivalent to each other. In this section, we fix a *finite* time horizon $T < \infty$. All stopping times τ are thus bounded by T . Moreover, we write $\frac{dP}{dQ}$ etc. for the densities on \mathcal{F}_T .

Epstein and Schneider (2003b) call \mathcal{Q} *rectangular* if for all stopping times τ and all $P, Q \in \mathcal{Q}$ the measure R given by

$$R(B) = \mathbb{E}^Q P(B|\mathcal{F}_\tau) \quad (B \in \mathcal{F})$$

belongs to \mathcal{Q} as well. Föllmer and Schied (2002) call \mathcal{Q} *stable* if for all stopping times τ , sets $A \in \mathcal{F}_\tau$, and priors $P, Q \in \mathcal{Q}$, there exists a unique measure $R \in \mathcal{Q}$ such that $R = P$ on \mathcal{F}_τ and for all random variables $Z \geq 0$ one has

$$\mathbb{E}^R [Z|\mathcal{F}_\tau] = \mathbb{E}^P [Z|\mathcal{F}_\tau] 1_{A^c} + \mathbb{E}^Q [Z|\mathcal{F}_\tau] 1_A. \quad (6)$$

Lemma A.1 *The following assertions are equivalent:*

1. \mathcal{Q} is time-consistent,
2. \mathcal{Q} is stable,
3. \mathcal{Q} is rectangular.

PROOF: Time-consistency implies stability: Suppose that \mathcal{Q} is time-consistent. Fix a stopping time τ , sets $A \in \mathcal{F}_\tau$, and priors $P, Q \in \mathcal{Q}$. Let (p_t) and (q_t) be the density processes of P and Q with respect to P_0 . Define a new stopping time $\sigma = \tau 1_A + T 1_{A^c}$. By time-consistency, the measure R given by

$$\frac{dR}{dP_0} = \frac{p_\sigma}{q_\sigma} \frac{dQ}{dP_0} \in \mathcal{Q}.$$

Note that

$$\frac{dR}{dP_0} = \frac{p_\tau}{q_\tau} \frac{dQ}{dP_0} 1_A + \frac{dP}{dP_0} 1_{A^c}.$$

Taking conditional expectations, we get

$$\left. \frac{dR}{dP_0} \right|_{\mathcal{F}_\tau} = p_\tau.$$

Hence, $R = P$ on \mathcal{F}_τ . Application of Bayes' formula yields (6).

Stability implies Rectangularity: Fix a stopping time τ and $P, Q \in \mathcal{Q}$. Take $A = \Omega$. By stability, there exists a measure $R \in \mathcal{Q}$ with $R = P$ on \mathcal{F}_τ and (6). Take $Z = 1_B$ for $B \in \mathcal{F}$. (6) yields $R(B|\mathcal{F}_\tau) = Q(B|\mathcal{F}_\tau)$. As $R = P$ on \mathcal{F}_τ , we obtain

$$R(B) = \mathbb{E}^R R(B|\mathcal{F}_\tau) = \mathbb{E}^P R(B|\mathcal{F}_\tau) = \mathbb{E}^P Q(B|\mathcal{F}_\tau).$$

Rectangularity implies Time-Consistency: Let $P, Q \in \mathcal{Q}$ and τ be a stopping time. Define R by setting

$$\frac{dR}{dP_0} = \frac{p_\tau}{q_\tau} \frac{dQ}{dP_0}.$$

For $B \in \mathcal{F}$, we obtain by conditioning and using Bayes' formula

$$\begin{aligned} R(B) &= \mathbb{E}^{P_0} \left[1_B \frac{p_\tau}{q_\tau} \frac{dQ}{dP_0} \right] \\ &= \mathbb{E}^{P_0} \left[\frac{p_\tau}{q_\tau} \mathbb{E}^{P_0} \left[1_B \frac{dQ}{dP_0} \middle| \mathcal{F}_\tau \right] \right] \\ &= \mathbb{E}^{P_0} [p_\tau Q(B|\mathcal{F}_\tau)] = \mathbb{E}^P Q(B|\mathcal{F}_\tau). \end{aligned}$$

Rectangularity yields $R \in \mathcal{Q}$. □

B Properties of Minimax Expected Values

For the sake of the reader, we list here some properties of minimax expected values that are known in the literature and used frequently in the arguments of the main text.

Let \mathcal{Q} be a set of probability measures equivalent to the reference measure P_0 . For random variables $Z \in L^\infty(\Omega, \mathcal{F}, P_0)$, we define the conditional minimax expected value $\pi_t(Z) = \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [Z|\mathcal{F}_t]$. From the properties of conditional expectations and the essential infimum, it follows immediately that π_t is

- monotone: for $Z \geq Z'$ in $L^\infty(\Omega, \mathcal{F}, P_0)$ we have $\pi_t(Z) \geq \pi_t(Z')$,
- conditionally homogeneous of degree 1: for \mathcal{F}_t -measurable random variables $\lambda \geq 0$, we have $\pi_t(\lambda Z) = \lambda \pi_t(Z)$ for all bounded random variables Z ,

- superadditive: for $Z, Z' \in L^\infty(\Omega, \mathcal{F}, P_0)$ we have $\pi_t(Z + Z') \geq \pi_t(Z) + \pi_t(Z')$,
- additive with respect to \mathcal{F}_t : for \mathcal{F}_t -measurable, bounded Z and all Z' we have $\pi_t(Z + Z') = Z + \pi_t(Z')$.

We need the following continuity properties.

- Lemma B.1** 1. π_t is Lipschitz-continuous with respect to the sup-norm on $L^\infty(\Omega, \mathcal{F}, P_0)$,
2. π_t is continuous from above in the following sense. If $X_k \downarrow X$ in $L^\infty(\Omega, \mathcal{F}, P_0)$, then $\pi_t(X_k) \downarrow \pi_t(X)$,
3. under Assumption 2.3, π_t is continuous from below in the following sense. For all $T > n$, if $X_k \uparrow X$ in $L^\infty(\Omega, \mathcal{F}_T, P_0)$, then $\pi_t(X_k) \uparrow \pi_t(X)$,
4. under Assumption 2.3, π_t satisfies bounded convergence in the following sense. For all $T > n$, if $X_k \rightarrow X$ in $L^\infty(\Omega, \mathcal{F}_T, P_0)$, and (X_k) is bounded by some number $K > 0$, then $\pi_t(X_k) \rightarrow \pi_t(X)$.

PROOF: The unconditional version of these results is in Delbaen (2002a), see Theorem 3.2. and Theorem 3.6. They carry over easily to the conditional case. \square

Lemma B.2 Let $T > 0$, $Z \in L^\infty(\Omega, \mathcal{F}_T, P_0)$ and $\tau \leq T$ a stopping time. Under Assumption 2.3, there exists a measure $P^{Z, \tau} \in \mathcal{Q}$ that coincides with P_0 on the σ -field \mathcal{F}_τ and

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[Z|\mathcal{F}_\tau] = \mathbb{E}^{P^{Z, \tau}}[Z|\mathcal{F}_\tau] .$$

PROOF: We show below that there exists a sequence $(P^m) \subset \mathcal{Q}$ with $P^m = P_0$ on \mathcal{F}_τ such that

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[Z|\mathcal{F}_\tau] = \lim_{m \rightarrow \infty} \mathbb{E}^{P^m}[Z|\mathcal{F}_\tau] .$$

By Assumption 2.3, the sequence has a weak limit point $P^{Z, \tau} \in \mathcal{Q}$ and

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[Z|\mathcal{F}_\tau] = \mathbb{E}^{P^{Z, \tau}}[Z|\mathcal{F}_\tau]$$

follows.

It remains to establish the existence of the minimizing sequence $(P^m) \subset \mathcal{Q}$. Note first that one can restrict attention to the set $\Phi = \{\mathbb{E}^P[Z|\mathcal{F}_\tau] | P \in \mathcal{Q} \text{ and } P = P_0 \text{ on } \mathcal{F}_\tau\}$. This is so because for arbitrary $P \in \mathcal{Q}$, we can define a new measure R with density

$$\frac{dR}{dP_0} = \frac{\frac{dP}{dP_0}}{\frac{dP}{dP_0} \Big|_{\mathcal{F}_\tau}}.$$

Then $R = P_0$ on \mathcal{F}_τ . As \mathcal{Q} is time-consistent, $R \in \mathcal{Q}$. By Bayes' formula,

$$\mathbb{E}^P[Z|\mathcal{F}_\tau] = \mathbb{E}^R[Z|\mathcal{F}_\tau].$$

We conclude that

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[Z|\mathcal{F}_\tau] = \operatorname{ess\,inf} \Phi.$$

The existence of the sequence $(P^m) \subset \mathcal{Q}$ with the desired properties follows if we can show that Φ is downward directed. Hence, let $P, \hat{P} \in \mathcal{Q}$ with $P = \hat{P} = P_0$ on \mathcal{F}_τ . Then

$$\min \left\{ \mathbb{E}^P[Z|\mathcal{F}_\tau], \mathbb{E}^{\hat{P}}[Z|\mathcal{F}_\tau] \right\} = \mathbb{E}^P[Z|\mathcal{F}_\tau] 1_A + \mathbb{E}^{\hat{P}}[Z|\mathcal{F}_\tau] 1_{A^c}$$

for $A = \left\{ \mathbb{E}^P[Z|\mathcal{F}_\tau] < \mathbb{E}^{\hat{P}}[Z|\mathcal{F}_\tau] \right\}$. We have to show that there exists $R \in \mathcal{Q}$ with $R = P_0$ on \mathcal{F}_τ and

$$\mathbb{E}^P[Z|\mathcal{F}_\tau] 1_A + \mathbb{E}^{\hat{P}}[Z|\mathcal{F}_\tau] 1_{A^c} = \mathbb{E}^R[Z|\mathcal{F}_\tau].$$

This follows from the equivalent characterization of time-consistency in Lemma A.1, 3. \square

Lemma B.3 *Let $Z \in L^\infty(\Omega, \mathcal{F}, P_0)$. Set*

$$V_t = \operatorname{ess\,sup}_{\tau \geq t} \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[X_\tau | \mathcal{F}_t].$$

There exists a sequence of stopping times (τ_k) with $\tau_k \geq t$ and

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P[X_{\tau_k} | \mathcal{F}_t] \uparrow V_t.$$

PROOF: By the usual properties of the essential supremum (see, e.g., Föllmer and Schied (2004), Appendix A.5), it is enough to show that the set $\{\text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_\tau | \mathcal{F}_t] \mid \tau \geq t\}$ is upward directed. Choose two stopping times $\tau_0, \tau_1 \geq t$. Set

$$A = \left\{ \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_0} | \mathcal{F}_t] > \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_1} | \mathcal{F}_t] \right\}.$$

Set $\tau_2 = \tau_0 1_A + \tau_1 1_{A^c}$. Then τ_2 is a stopping time that is greater or equal t . The proof is complete if we can show that

$$\text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_2} | \mathcal{F}_t] = \max \left\{ \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_0} | \mathcal{F}_t], \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_1} | \mathcal{F}_t] \right\}.$$

It is obvious from the definition of τ_2 that the left hand side is smaller or equal the right hand side. Let us show the other inequality. By Lemma B.2, there exist measures P_0 and P_1 such that

$$\text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_i} | \mathcal{F}_t] = \mathbb{E}^{P_i} [X_{\tau_i} | \mathcal{F}_t], \quad i = 0, 1.$$

Then we have

$$\begin{aligned} \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_0} 1_A + X_{\tau_1} 1_{A^c} | \mathcal{F}_t] &\geq \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_0} | \mathcal{F}_t] 1_A + \text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_1} | \mathcal{F}_t] 1_{A^c} \\ &= \mathbb{E}^{P_0} [X_{\tau_0} | \mathcal{F}_t] 1_A + \mathbb{E}^{P_1} [X_{\tau_1} | \mathcal{F}_t] 1_{A^c} \\ &= \mathbb{E}^{P_0} [X_{\tau_2} | \mathcal{F}_t] 1_A + \mathbb{E}^{P_1} [X_{\tau_2} | \mathcal{F}_t] 1_{A^c}. \end{aligned}$$

Time-consistency of \mathcal{Q} and Lemma A.1 imply that there exists a measure $P_2 \in \mathcal{Q}$ such that

$$\mathbb{E}^{P_0} [X_{\tau_2} | \mathcal{F}_t] 1_A + \mathbb{E}^{P_1} [X_{\tau_2} | \mathcal{F}_t] 1_{A^c} = \mathbb{E}^{P_2} [X_{\tau_2} | \mathcal{F}_t].$$

Altogether, we obtain

$$\text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_2} | \mathcal{F}_t] \geq \mathbb{E}^{P_2} [X_{\tau_2} | \mathcal{F}_t],$$

and as $P_2 \in \mathcal{Q}$

$$\text{ess inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_2} | \mathcal{F}_t] = \mathbb{E}^{P_2} [X_{\tau_2} | \mathcal{F}_t].$$

Now our claim follows as we have

$$\begin{aligned}
\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_2} | \mathcal{F}_t] &= \mathbb{E}^{P_2} [X_{\tau_2} | \mathcal{F}_t] \\
&= \mathbb{E}^{P_2} [X_{\tau_0} | \mathcal{F}_t] 1_A + \mathbb{E}^{P_2} [X_{\tau_1} | \mathcal{F}_t] 1_{A^c} \\
&\geq \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_0} | \mathcal{F}_t] 1_A + \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_1} | \mathcal{F}_t] 1_{A^c} \\
&= \max \left\{ \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_0} | \mathcal{F}_t], \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [X_{\tau_1} | \mathcal{F}_t] \right\},
\end{aligned}$$

where we have used the definition of A in the last line. □

C Minimax Martingale Theory

This section gathers the material of Subsection 3.2 with all proofs.

Definition C.1 *Let \mathcal{Q} be a set of priors. Let $(M_t)_{t \in \mathbb{N}}$ be an adapted process with $\mathbb{E}^P |M_t| < \infty$ for all $P \in \mathcal{Q}$ and $t \in \mathbb{N}$. (M_t) is called a minimax (sub-, super-)martingale with respect to \mathcal{Q} if we have for $n \in \mathbb{N}$*

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [M_{t+1} | \mathcal{F}_t] = (\geq, \leq) M_t.$$

Lemma C.2 *Let (M_t) be a bounded, adapted process.*

1. *M is a minimax submartingale if and only if it is a \mathcal{Q} -submartingale,*
2. *M is a minimax supermartingale if and only there exists $P^* \in \mathcal{Q}$ such that M is a P^* -supermartingale,*
3. *M is a minimax martingale with respect to \mathcal{Q} if and only if*
 - (a) *there exists $P^* \in \mathcal{Q}$ such that M is a P^* -martingale and*
 - (b) *M is a \mathcal{Q} -submartingale.*

PROOF: For (1.), note that

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [M_{t+1} | \mathcal{F}_t] \geq M_t$$

is equivalent to

$$\text{for all } P \in \mathcal{Q} \quad \mathbb{E}^P [M_{t+1} | \mathcal{F}_t] \geq M_t.$$

We proceed with (2.). Suppose that M is a P^* -supermartingale for some $P^* \in \mathcal{Q}$ -submartingale. Then we have for $t \in \mathbb{N}$

$$M_t \geq E^{P^*} [M_{t+1} | \mathcal{F}_t] \geq \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [M_{t+1} | \mathcal{F}_t],$$

and M is a minimax supermartingale.

For the converse, we need the assumption of time-consistency. By Lemma B.2, there exist measures $P^{t+1} \in \mathcal{Q}$ for $t \in \mathbb{N}$ that coincide with P_0 on \mathcal{F}_t and satisfy

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [M_{t+1} | \mathcal{F}_t] = \mathbb{E}^{P^{t+1}} [M_{t+1} | \mathcal{F}_t].$$

Therefore, the density of P^{t+1} with respect to P_0 is 1 on \mathcal{F}_t . Let z^{t+1} be the density of P^{t+1} with respect to P_0 on \mathcal{F}_{t+1} . By Bayes' formula, we have

$$M_t \geq \operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [M_{t+1} | \mathcal{F}_t] = \mathbb{E}^{P_0} [M_{t+1} z^{t+1} | \mathcal{F}_t]. \quad (7)$$

Construct a new measure P^* by setting

$$\left. \frac{dP^*}{dP_0} \right|_{\mathcal{F}_T} = z^1 z^2 \cdots z^T \quad (T \in \mathbb{N}).$$

By time-consistency, $P^* \in \mathcal{Q}$. We claim that M is a P^* -supermartingale. To see this, use Bayes' formula and Eqn. (7) to get

$$\begin{aligned} \mathbb{E}^{P^*} [M_{t+1} | \mathcal{F}_t] &= \mathbb{E}^{P_0} \left[M_{t+1} \left. \frac{dP^*}{dP_0} \right|_{\mathcal{F}_{t+1}} \middle| \mathcal{F}_t \right] \left(\left. \frac{dP^*}{dP_0} \right|_{\mathcal{F}_t} \right)^{-1} \\ &= \mathbb{E}^{P_0} [M_{t+1} z^{t+1} | \mathcal{F}_t] \leq M_t. \end{aligned}$$

Therefore, M is a P^* -supermartingale.

For (3.), combine (1.) and (2.). □

Note the big difference between minimax sub- and supermartingales. While a minimax submartingale is a submartingale for all $Q \in \mathcal{Q}$ uniformly, a minimax supermartingale is a supermartingale for *some* $Q \in \mathcal{Q}$ only. This is due, of course, to the fact that we take always the essential infimum over a class of probability measures.

We are now going to extend two fundamental theorems from martingale theory to minimax martingales. We start with the famous Doob decomposition.

Theorem C.3 (Doob Decomposition) *Let S be a bounded minimax supermartingale (submartingale) with respect to \mathcal{Q} . Then there exists a minimax martingale M and a predictable, nondecreasing process A with $A_0 = 0$ such that $S = M - A$ ($S = M + A$). Such a decomposition is unique.*

PROOF: For uniqueness, note that from $S = M - A$ with the stated properties, we obtain

$$0 = \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [S_{t+1} - S_t + A_{t+1} - A_t | \mathcal{F}_t] ,$$

and predictability of A yields the recursive relation

$$A_{t+1} = A_t - \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [S_{t+1} - S_t | \mathcal{F}_t] . \quad (8)$$

In conjunction with $A_0 = 0$, this determines A , and then M , uniquely.

Now let A be given by (8) and $A_0 = 0$. Note that A is predictable and nondecreasing as S is a minimax supermartingale. Let $M_t = S_t + A_t$. We have to show that M is a minimax martingale. But the predictability of A implies

$$\begin{aligned} \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [M_{t+1} - M_t | \mathcal{F}_t] &= \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [S_{t+1} - S_t + A_{t+1} - A_t | \mathcal{F}_t] \\ &= A_{t+1} - A_t + \operatorname{ess\,inf}_{P \in \mathcal{Q}} E^P [S_{t+1} - S_t | \mathcal{F}_t] = 0 . \end{aligned}$$

This completes the proof. \square

Remark C.4 *As minimax submartingales are nothing but \mathcal{Q} -submartingales, it is worthwhile to compare the preceding Doob decomposition with the so-called optional or uniform Doob decompositions for \mathcal{Q} -submartingales used in the theory of hedging, where \mathcal{Q} is given by the (time-consistent) set of equivalent martingale measures for some financial market (see El Karoui and Quenez (1995), Föllmer and Kabanov (1998), Kramkov (1996)). Here, we decompose a \mathcal{Q} -submartingale into a minimax martingale M and a predictable, increasing process A starting at 0. This, however, is not a uniform Doob decomposition as M is usually only a \mathcal{Q} -submartingale, not a \mathcal{Q} -martingale. In fact, for such a uniform decomposition, A is usually only adapted, not predictable.*

The second fundamental theorem concerns the preservation of the (super)–martingale property under optimal stopping.

Theorem C.5 (Optional Sampling Theorem) *Let Z be a bounded minimax supermartingale with respect to \mathcal{Q} . Let $\sigma \leq \tau$ be stopping times. Assume that τ is universally finite in the sense that $P[\tau < \infty] = 1$ for all $P \in \mathcal{Q}$. Then*

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [Z_\tau | \mathcal{F}_\sigma] \leq Z_\sigma.$$

PROOF: The Doob decomposition allows to write $Z = M - A$ for a minimax martingale M and a nondecreasing predictable process A . By the above Lemma 3.4, there exists $P^* \in \mathcal{Q}$ such that M is a supermartingale under P^* . The standard Optional Sampling Theorem states that

$$\mathbb{E}^{P^*} [Z_\tau | \mathcal{F}_\sigma] \leq Z_\sigma.$$

As a consequence,

$$\operatorname{ess\,inf}_{P \in \mathcal{Q}} \mathbb{E}^P [Z_\tau | \mathcal{F}_\sigma] \leq \mathbb{E}^{P^*} [Z_\tau | \mathcal{F}_\sigma] \leq Z_\sigma.$$

□

Remark C.6 *As minimax submartingales are \mathcal{Q} -submartingales, the above theorem holds also true for minimax submartingales.*

D Proof of Theorem 5.1

Note that every probability measure that is locally equivalent to the uniform probability has a density process of the form (5), where $\alpha_t = P[X_t = 1 | \mathcal{F}_{t-1}]$.

We start by showing that every class of priors \mathcal{Q} that satisfies the Assumptions 2.2, 2.3, and 3.2 can be represented as in (5). Let us fix $t > 0$ and a prior $P \in \mathcal{Q}$. As \mathcal{Q} is compact by Assumption 2.3, the lower bound $\underline{p}_t := \inf_{P \in \mathcal{Q}} P[X_t = 1 | \mathcal{F}_{t-1}]$ is actually attained, and thus strictly positive by Assumption 2.2. By the same argument, one obtains $\bar{p}_t := \sup_{P \in \mathcal{Q}} P[X_t = 1 | \mathcal{F}_{t-1}] < 1$.

Let \underline{P}, \bar{P} be the measures in \mathcal{Q} whose density processes are represented by the processes \underline{p} and \bar{p} resp. Due to time-consistency, $\underline{P}, \bar{P} \in \mathcal{Q}$.

Now let $\hat{\mathcal{Q}}$ be the class of all priors that satisfy the requirement (5) with $\underline{p} \leq a \leq \bar{p}$. From the preceding, we know that $\mathcal{Q} \subseteq \hat{\mathcal{Q}}$. We show by induction

over t that

$$\left\{ \frac{dP}{dP_0} \Big|_{\mathcal{F}_t} \mid P \in \mathcal{Q} \right\} = \left\{ 2^t \prod_{s=1}^t \alpha_s^{X_s} (1 - \alpha_s)^{1-X_s} \mid \underline{p}_s \leq \alpha_s \leq \bar{p}_s, (s = 1, \dots, t), \right. \\ \left. (\alpha_s) \text{ predictable} \right\}.$$

For $t = 1$, this follows from convexity of \mathcal{Q} . For the induction step, suppose that we have proved the claim for t . Take some predictable process $(\alpha_s)_{s=1, \dots, t+1}$ with $\underline{p}_s \leq \alpha_s \leq \bar{p}_s$ for $s = 1, \dots, t+1$. We have to show that

$$Z = 2^{t+1} \prod_{s=1}^{t+1} \alpha_s^{X_s} (1 - \alpha_s)^{1-X_s}$$

belongs to \mathcal{D}_{t+1} . By time-consistency, the measures \underline{R} and \bar{R} with densities

$$\frac{d\underline{R}}{dP_0} \Big|_{\mathcal{F}_{t+1}} = 2\underline{p}^{X_{t+1}} (1 - \underline{p})^{1-X_{t+1}}$$

and

$$\frac{d\bar{R}}{dP_0} \Big|_{\mathcal{F}_{t+1}} = 2\bar{p}^{X_{t+1}} (1 - \bar{p})^{1-X_{t+1}}$$

are in \mathcal{D}_{t+1} . As \mathcal{Q} is convex and α_{t+1} assumes at most 2^t many values, also the measure R with density

$$\frac{dR}{dP_0} \Big|_{\mathcal{F}_{t+1}} = 2\alpha_{t+1}^{X_{t+1}} (1 - \alpha_{t+1})^{1-X_{t+1}}$$

is in \mathcal{D}_{t+1} . Now use time-consistency again to see that $Z \in \mathcal{D}_{t+1}$.

For the converse, let \mathcal{Q} be a class of priors that satisfy the requirement (5) with $\underline{p} \leq a \leq \bar{p}$ for two predictable processes (\underline{p}_t) and (\bar{p}_t) with $0 < \underline{a}_t \leq \bar{a}_t < 1$. Then local equivalence, i.e. Assumption 2.2 is obvious. (Weak) compactness, i.e. Assumption 2.3 follows from $0 < \underline{p}_t \leq \bar{p}_t < 1$. To show time consistency, let P and Q be represented by the processes (α_t) and (β_t) resp. Set $\gamma_t = \alpha_t$ on $\{t \leq \tau\}$ and $\gamma_t = \beta_t$ on $\{t > \tau\}$. Let R be the measure that is represented by (γ_t) . R belongs to \mathcal{Q} and its density process satisfies (2). This proves time-consistency.

References

- ARROW, K. (1971): *Essays in the Theory of Risk Bearing*. Markham Publishing Co., Chicago.
- ARTZNER, P., F. DELBAEN, J.-M. EBER, D. HEATH, AND H. KU (2002): “Multiperiod Risk and Coherent Multiperiod Risk Measurement,” Mimeo, ETH Zürich.
- CHATEAUNEUF, A., F. MACCHERONI, M. MARINACCI, AND J.-M. TALLON (2005): “Monotone Continuous Multiple Priors,” *Economic Theory*, 26(4), 973–982.
- CHOW, Y., H. ROBBINS, AND D. SIEGMUND (1971): *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin Comp., Boston.
- COQUET, F., Y. HU, J. MÉMIN, AND S. PENG (2002): “Filtration-consistent Nonlinear Expectations and Related g-Expectations,” *Probability Theory and Related Fields*, 123, 1–27.
- COX, J. C., S. A. ROSS, AND M. RUBINSTEIN (1979): “Option Pricing: A Simplified Approach,” *Journal of Financial Economics*, 7, 229–263.
- DELBAEN, F. (2002a): “Coherent Risk Measures on General Probability Spaces,” in *Essays in Honour of Dieter Sondermann*, ed. by K. Sandmann, and P. Schönbucher, pp. 1–37. Springer-Verlag, New York.
- (2002b): “The Structure of m-Stable Sets and in Particular of the Set of Risk neutral Measures,” mimeo, ETH Zürich.
- DETLEFSEN, K., AND G. SCANDOLO (2005): “Conditional and Dynamic Convex Risk Measures,” *Finance and Stochastics*, 9(5), 539–561.
- DIXIT, A., AND R. PINDYCK (1994): *Investment under Uncertainty*. Princeton University Press, Princeton.
- EICHBERGER, J., AND D. KELSEY (1996): “Uncertainty Aversion and Dynamic Consistency,” *International Economic Review*, 37, 625–640.
- EL KAROUI, N., AND M.-C. QUENEZ (1995): “Dynamic Programming and Pricing of Contingent Claims in an Incomplete Market,” *SIAM Journal of Control and Optimization*, 33, 29–66.

- EPSTEIN, L., AND M. MARINACCI (2006): “Mutual Absolute Continuity of Multiple Priors,” Working Paper.
- EPSTEIN, L., AND M. SCHNEIDER (2003a): “IID: Independently and Indistinguishably Distributed,” *Journal of Economic Theory*, 113(1), 32–50.
- (2003b): “Recursive Multiple Priors,” *Journal of Economic Theory*, 113, 1–31.
- FERGUSON, T. (2006): “Optimal Stopping and Applications,” Electronic Text, <http://www.math.ucla.edu/~tom/Stopping/Contents.html>, University of California, Los Angeles.
- FÖLLMER, H., AND Y. KABANOV (1998): “Optional Decomposition and Lagrange Multipliers,” *Finance and Stochastics*, 2, 1–25.
- FÖLLMER, H., AND I. PENNER (2007): “Convex Risk Measures and The Dynamics of Their Penalty Functions,” *Statistics and Decisions*, p. to appear.
- FÖLLMER, H., AND A. SCHIED (2002): *Stochastic Finance*, Studies in Mathematics. de Gruyter, Berlin, New York.
- FÖLLMER, H., AND A. SCHIED (2004): *Stochastic Finance, 2nd edition*. de Gruyter, Berlin, New York.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- KARATZAS, I., AND S. KOU (1998): “Hedging American Contingent Claims with Constrained Portfolios,” *Finance and Stochastics*, 2, 215–258.
- KARATZAS, I., AND I. ZAMFIRESCU (2003): “Game Approach to the Optimal Stopping Problem,” Working Paper.
- KRAMKOV, D. (1996): “Optional Decomposition of Supermartingales and Hedging Contingent Claims in Incomplete Security Markets,” *Probability Theory Related Fields*, 105, 459–479.
- LERCHE, R., R. KEENER, AND M. WOODROOFE (1994): “A Generalized Parking Problem,” in *Statistical Decision Theory and Related Topics V*, ed. by J. B. S.S. Gupta, pp. 523–532. Springer Verlag.

- MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): “Dynamic Variational Preferences,” *Journal of Economic Theory*, 128, 4–44.
- MACHINA, M. (1989): “Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty,” *Journal of Economic Literature*, 27(4), 1622–1668.
- NISHIMURA, K. G., AND H. OZAKI (2007): “Irreversible Investment and Knightian Uncertainty,” *Journal of Economic Theory*, p. to appear.
- PENG, S. (1997): “BSDE and Related g-Expectation,” in *Backward Stochastic Differential Equation*, ed. by N. El Karoui, and L. Mazliak, no. 364 in Pitman Research Notes in Mathematics Series. Addison Wesley Longman Ltd.
- RIEDEL, F. (2004): “Dynamic Coherent Risk Measures,” *Stochastic Processes and Their Applications*, 112, 185–200.
- SARIN, R., AND P. WAKKER (1998): “Dynamic Choice and Nonexpected Utility,” *Journal of Risk and Uncertainty*, 17, 87–119.
- SNELL, L. (1952): “Applications of the Martingale Systems Theorem,” *Transactions of the American Mathematical Society*, 73, 293–312.
- TUTSCH, S. (2006): “Konsistente und konsequente dynamische Risikomaße und das Problem der Aktualisierung,” Ph.D. thesis, Humboldt University at Berlin.
- YOO, K.-R. (1991): “The Iterative Law of Expectation and Non-Additive Probability Measures,” *Economics Letters*, 37, 145–149.
- ZAMFIRESCU, I.-M. (2003): “Optimal Stopping under Model Uncertainty,” Ph.D. thesis, Columbia University.

Figure 1: Tree for Example 3.1. The decision maker uses two probabilistic models. In Model 1, the (conditional) probability of moving up is $1/3$ in every node; in Model 2, this probability is $2/3$. The payoff from stopping is indicated by the bold numbers at the nodes. Time consistency fails because the decision maker does not take the *worst case* probability measure into account. Here, the worst case probability measure would have probability $2/3$ of moving up in the upper node and probability $1/3$ of moving up in the lower node.