On Dynamic Knightian Uncertainty Models: Time-Consistency and Optimal Behavior

Inaugural-Dissertation

zur Erlangung des Grades eines Doktors

der Wirtschaftswissenschaften

durch die

Fakultät für Wirtschaftswissenschaften

der Universität Bielefeld

vorgelegt von

Monika Bier

aus Wheeling, West Virginia, USA

Bielefeld, 2009

Erstgutachter:
Prof. Dr. Frank Riedel
Institut für Mathematische
Wirtschaftsforschung (IMW)
Universität Bielefeld

Zweitgutachter:
JProf. Dr. Dr. Frederik Herzberg
Institut für Mathematische
Wirtschaftsforschung (IMW)
Universität Bielefeld

Gedruckt auf alterungsbeständigem Papier nach DIN-ISO 9706

to my grandma

Acknowledgments

In writing this thesis I have greatly benefited from the help and encouragement of numerous people.

First and foremost, I appreciate the guidance and support of my advisor Professor Dr. Frank Riedel. This doctoral thesis would not have been possible without his motivation, advice and ideas for further investigation. In many fruitful discussions he helped me with valuable comments, and suggestions and prevented me from getting lost in detours making my research aims reachable again. He encouraged me to present my work on several opportunities for which I am thankful. In addition, I am indebted to JProfessor Dr. Dr. Frederik Herzberg for his efforts in surveying this thesis.

I acknowledge financial support from the German Research Foundation (DFG) and from the Institute of Mathematical Economics (IMW) at Bielefeld University.

Thanks also to various seminar and conference participants who gave me helpful suggestions and encouragement.

A huge contribution to this work was also made by several colleagues and friends: Most important my former office mate, friend, and coauthor Daniel Engelage with whom I enjoyed many important discussions in- and outside of the office. In my opinion we complement each other perfectly as researchers but he has also become a valuable friend. At this point I also want to mention his successor as office mate Tatjana Chudjakow who was never short of an economic interpretation or a motivational word, whatever was needed more. On top of that she is an excellent travel companion and guide. I also want to thank Simon, Jan, and Jörg for numerous helpful mathematical remarks and inspiring coffee breaks. Special thanks to Simon who never tired of trying to convince me of Bielefeld's positive sides and making me feel welcome.

Furthermore I would like to thank my colleagues and friends at the Bonn Graduate School of Economics (BGSE) for making me feel at home. Here I want to especially mention Almira Buzaushina, Marcelo Cadena, Jördis Hengelbrock, Stefanie Lehmann, Christina Matzke, Klaas Schulze, and Bernd Schlusche for helpful suggestions and enjoyable times.

Huge thanks also to the members of the Institute of Mathematical Economics (IMW) for giving me such a hearty welcome and making my short stay as pleasant as it was. Thanks also for the many amusing evenings in the Lounge, especially our foosball tournaments. Special thanks to Bettina Buiwitt-Robson for just being who she is, to Matthias Schleef for solving all my IT and bicycle problems, and Sonja Brangewitz for all her special care.

Last but not least, special thanks go to Marcus, my family and my friends for their patience, their unconditional emotional support, and encouragement as well as their enduring belief in me.

Contents

Chapter 1

General Introduction

As can be well observed in the current financial crisis choosing the wrong model specifications for making decisions or assessing risk can have very severe consequences if the underlying theory is not robust to model uncertainty. This seems to have been one of the major shortcomings which lead to the present economic situation. Therefore a major aim in economics is to reduce model risk by developing robust approaches to decision making.

A link between robust control theory and decision theory was shown in [Hansen & Sargent, 01]. In robust control theory model uncertainty arises by the perturbation of a unique approximating model. This corresponds to uncertainty about the true distribution in a decision theoretical ansatz which will be the focus of the following.

The well known works of [von Neumann & Morgenstern, 44], [Savage, 54] and [Anscombe & Aumann, 63] are among the first theoretical models on decision making. One common property of all these models however is that they incorporate only one single known distribution of the outcomes, i.e. what is often called a purely risky setting, which in reality is a situation seldom found. The first comment on the fact that there is more than pure risk can be found in [Knight, 21] which is why uncertainty about the true distribution is often referred to as *Knightian uncertainty* but sometimes also as *ambiguity*.

1.1 Knightian Uncertainty

In his seminal paper [Knight, 21] suggested that there exist random outcomes which cannot be represented by numerical probabilities, i.e. he establishes a clear distinction between *measurable uncertainty* he calls risk and *unmeasurable uncertainty*. This unmeasurable uncertainty can among other reasons arise when the decision-maker is ignorant of statistical frequencies relevant to his decision or when a priori estimations are impossible to obtain or the decision is unique in the sense that there is no information to build an approximation for numerical probabilities. To make this differentiation a bit more explicit think of following examples. If one has to bet heads or tails in a coin toss one would assume that both are equally likely based on one's experience, i.e. one can somehow measure the probability and finds oneself in a pure risky setting. However if one is asked to bet on e.g. the outcome of a tennis match without possessing any information about the players, it is not clear how to assign a unique probability to the outcomes and one is in an uncertain setting.

Based on this theoretical approach is the empirical work of [Ellsberg, 61]. Following the aim of giving evidence for Knight's theory, he constructed the following urn-experiment with two urns containing 100 black and red balls. In the first urn the ratio of black to red is unknown, it can be anything between 0 and 100. In the second urn there are 50 red and 50 black balls. If one denotes the bet of getting \$100 if a red ball is pulled from urn i and \$0 else by RED_i, and $\mathrm{BLACK}_\mathrm{i}$ for the respective bet on a black ball Ellsberg claims that a majority of people show following preferences: They are indifferent between RED_i and $\mathrm{BLACK}_\mathrm{i}$ for $i = 1, 2$ but prefer RED_2 to RED_1 and BLACK_2 to BLACK_1.

He shows that there is absolutely no way to assign probabilities to the event of red being pulled from the first urn, to explain these preferences implying that the classical decision theory dealing with one unique distribution, as e.g. introduced in [Savage, 54] cannot contain the whole truth and that there do exist the *unmeasurable uncertainties* proposed by Knight.

An attempt to underpin these findings with a theoretical model is given in [Gilboa & Schmeidler, 89]. They set up axioms for preferences which lead to Maxmin Expected Utility or the Multiple Priors Model. They weakened the Independence Axiom of Anscombe and Aumann and introduced an axiom formalizing Uncertainty Aversion to the model. By this slight change they accomplished a representation of their preferences which instead of one unique distribution contained a whole set of possible distributions. It leads the decision maker to maximize $\inf_{\mathbb{P}\in\mathcal{C}} \mathbb{E}^{\mathbb{P}}[u \circ f]$ among all possible acts f, where $\mathcal{C}$ is a non-empty set of distributions. Heuristically one can interpret this as having the decision maker think all distributions in $\mathcal{C}$ possible and in order to be on the safe side he always looks at the one which gives him the lowest utility. If the utility under the worst possible distribution is enough to make him choose act f surely it is high enough under all other distributions as well. Seen this way these preferences correspond to an extremely conservative decision maker.

1.2 Time-Consistency

Since up to now the models were purely atemporal [Epstein & Schneider, 03] generalized the above approach to a dynamic setting. They appropriately modified the axioms of [Gilboa & Schmeidler, 89] to be not only state but also time dependent and additionally asked for *Dynamic Consistency* in the preferences. With this assumption they meant if two acts are identical up to some time t but one is preferred over the other in $t+1$ then this should already be the case at time t implying that a decision maker will never regret his earlier choices. This restriction on the preferences yields a very special property for the set of distributions in the Utility functional. They showed that preferences are dynamically consistent if and only if the corresponding set of distributions arising in their *Recursive Multiple Priors Representation* is rectangular. Rectangularity is a restriction on the whole set of measures.

Equivalent definitions were formulated by various authors. A survey of the different concepts and a proof of their equivalence can be found in [Riedel, 09] and [Delbaen, 03]. Among these concepts is the above mentioned rectangularity introduced in [Epstein & Schneider, 03] which is a property concerning the one-step-ahead measures. They asked that at every point in time all possible one-step-ahead measures can be added. Another concept is stability introduced in [Föllmer & Schied, 04]. Here for two measures $\mathbb{P}$ and $\mathbb{Q}$ in the set of measures and every stopping time τ the measure that takes $\mathbb{P}$ up to τ and $\mathbb{Q}$ afterwards is also contained in the set. The last concept is time-consistency which was introduced in [Delbaen, 03]. This property demands that at every stopping time density processes can be consistently pasted together. It is also equivalent to a Law of Iterated Expectations meaning that at time s my expected future payoff is the same as the expectation in s of my future expected payoff in $t \geq s$.

These may seem as rather technical assumptions but they also have some intuitive consequences for the decision maker. He can for instance change his mind in every time period about which measure he thinks is the true one or the worst one and time-consistency guarantees that this measure is contained in the set of his possible measures. This implies that as time passes he will never regret his previous decisions since at every point in time he can decide optimally. Another implication is that he can use backward induction for solving problems which makes large classes of problem a lot more tractable.

One further generalization of these utility models is the variational preference model introduced in a static set up in [Maccheroni et al., 06a] and extended to a dynamic framework in [Maccheroni et al., 06b]. Since up to now all distributions in the set were conceived as equally likely they introduced penalty functions allowing to differentiate the different distributions according to their likelihood.

1.3 Risk Measures

Up to now we have focused on utility models that arise from preferences but analogous models can also be found by axiomatizing risk measures.

Static coherent risk measures which correspond to the multiple prior model under the assumption of risk neutrality, a discount factor of one and no intermediate payoffs were first axiomatized in [Artzner et al., 99] and their dynamic generalizations can inter alia be found in [Riedel, 04] or [Artzner et al., 07]. The robust representation of a coherent risk measure is identical to the one for multiple prior preferences up to a minus sign, having the decision maker look at the largest expected loss as a basis for his decision.

Since for risk measures it is important to incorporate liquidity risk and to give less conservative assessments which coherent risk measures cannot provide convex risk measures were introduced and can e.g. be found in [Föllmer & Schied, 04] for a static setting. For a dynamic setting we refer to [Föllmer & Penner, 06] or [Föllmer et al., 07] for risky projects seen as payoffs in the last period while risky projects seen as stochastic processes are studied in [Cheridito et al., 06]. Here again the equivalence to variational preferences is given up to a minus sign.

Time-consistency concepts are the same in both approaches although in terms of risk measures one usually works with the one corresponding to the law of iterated expectations.

Considering this it makes no difference if we make the following observations in terms of utility functionals or risk measures. Each chapter may be reformulated in terms of the other approach resrectively since for simplicity we assume risk neutral decision makers without discounting in all chapters. However in Chapter 2 and 3 we will look at utility functionals while Chapter

4 will be mainly expressed in terms of risk measures where we will give a further short overview of the theory of risk measures and their robust representation.

1.4 Particular Considerations

The main chapters of this thesis, each of which is self-contained, are based on three articles. The first deals with the construction or an alternative characterization of time-consistent sets of measures in the special framework of finite trees, the second shows as a main result a duality theorem this however in a continuous time setting while the topic of the third one, coauthored by Daniel Engelage is concerned with the merging of convex risk measures as information is gained in the course of time.

In **Chapter 2** we solve the question what time-consistent sets of measures look like in finite trees. In [Riedel, 09] time-consistent sets of measures in a discrete setting are constructed via their density processes which gave rise to the question if all time-consistent sets can be constructed in this way. When restricting the framework to finite trees with a constant and finite splitting function we show that every time-consistent set of measures can be described via predictable processes. For each measure $\mathbb{P}$ in a time-consistent set $\mathcal{P}$ we get a predictable process $\alpha^{\mathbb{P}}$ whose dimension is one less than the splitting value of the tree. The set of predictable processes $\mathcal{A} = \{\alpha^{\mathbb{P}} \mid \mathbb{P} \in \mathcal{P}\}$ that arises via this identification has specific features. These features in return guarantee that a time-consistent set of measures can be created from a set of predictable processes with these properties. After this characterization we show some examples or applications for the use of this theorem. Additionally we show that standard generalizations of this theorem fail to go through, showing that our characterization is a universal one.

Chapter 3 contains a Duality theorem which allows to switch the order of minimization and maximization in order to solve optimal stopping problems under ambiguity. This theorem is again set in the recursive multiple priors model of [Epstein & Schneider, 03] and works for fairly general assumptions on our payoff process X and rather standard assumptions on our set of measures $\mathcal{P}$. We make strong use of an explicit but general construction for time-consistent sets of measures given in [Delbaen, 03]. We also apply this theorem to specific classes of payoff processes. It allows to determine an optimal stopping time for multiple prior super- and submartingales as payoff processes and in the case of κ-ambiguity adapted to our framework and a Brownian motion with drift as payoff process we are able to identify the worst case distribution and hence our ambiguous stopping problems shrinks to a classical problem.

In **Chapter 4**, coauthored by Daniel Engelage, we tackle the question of how in an ambiguous environment the assessment of risk and with this optimal behavior changes as time passes and information increases. The vehicle for this analysis will be convex risk measures or dynamic variational preferences equivalently. Our first approach to incorporate a kind of learning mechanism into convex risk measures is a constructive one via the minimal penalty function in the robust representation. However in our explicit approach to construct the penalty via the likelihoods of the distributions time-consistency seems to be a major problem. Therefore in our second approach we take a dynamically consistent set of risk measures as given and show that in the long run all uncertainty is revealed, leaving the decision maker behave as a utility maximizer under the true distribution. Formulating this result closer to the fundamental Blackwell-Dubins theorem of which this is a generalization: two decision makers who agree on sure and impossible events but with different opinions of the risk they face, modeled here via different penalty functions, tend towards agreeing on the true underlying distribution in the end. As mentioned we extend the main result in [Blackwell & Dubins, 62] which holds for probability measures to convex risk measures. An important step in this generalization is also the extension to not necessarily time-consistent convex risk measures. With this our existence result for a limiting distribution becomes more general than the one found in [Föllmer & Penner, 06]. To make things clearer we study entropic risk measures as an application.

To this point we have given a brief outline of the general context and developments which lead to this work. Since the questions and topics treated in the following chapters differ a more detailed scientific placement of this work will be discussed in each chapter separately.

Chapter 2

Time-Consistent Sets of Measures on Finite Trees

2.1 Introduction

In 1944 von Neumann and Morgenstern formulated their famous axioms for preferences over random payoffs (see [von Neumann & Morgenstern, 44]) and showed that these preferences are equivalent to an *Expected Utility Representation* of preferences. After some time their model was criticized because the distributions of their payoffs were exogenously given and purely objective. Since this is a very restrictive assumption their model was extended in [Savage, 54] and in [Anscombe & Aumann, 63]. In contrast to the von Neumann and Morgenstern model Savage regarded the distributions of the payoffs to be purely subjective and endogenous. Anscombe and Aumann then combined both models taking some objective distributions as given and having others arising purely out of the model.

At some point criticism also arose against these models. One of the most mentioned objections can be found in [Ellsberg, 61]. He conducted experiments and empirically showed that Expected Utility models do not always mirror reality. One way of explaining these findings is that people behave only boundedly rational. Another way is to distinguish between *uncertainty* and *risk*. While in a risky setting the decision maker is sure of the distributions of the outcomes in an uncertain setting he is unsure of the right distribution and thinks more than one possible. Following this idea Gilboa and Schmeidler developed their Multiple Priors Model in [Gilboa & Schmeidler, 89] using Anscombe's and Aumann's model as a basis. They weakened the Independence Axiom and added an additional axiom formalizing Uncertainty Aversion. This lead the decision maker to maximize $\inf_{\mathbb{P}\in\mathcal{C}} \mathbb{E}^{\mathbb{P}}[u \circ f]$ among all possible acts f, where $\mathcal{C}$ is a non-empty, closed and convex set of probability measures.

Since this is a purely atemporal model in [Epstein & Schneider, 03] the Multiple Priors Model was expanded to incorporate the factor time. They modified preferences to be not only state but also time-dependent, adjusted the G-S-axioms appropriately and asked for *Dynamic Consistency* as an additional axiom. This restriction on preferences yields a very specific property of the set of measures in their Utility Representation. They found out that preferences are dynamically consistent if and only if the set of measures in their Recursive Multiple Priors Representation is rectangular. Rectangularity is a restriction on the whole set of measures. It demands that it is possible for the one-step-ahead measures to be mixed arbitrarily. Since for some purposes (e.g. solving concrete optimal stopping problems) this is not a very easy definition but never the less an important one it is very natural to try and find equivalent definitions.

This was done by various authors. In [Riedel, 09] one can find a survey of the different concepts and a proof of their equivalence. Among these concepts is rectangularity which was introduced in [Epstein & Schneider, 03] and is a property concerning the one-step-ahead measures. They asked that at every point in time all possible one step ahead measures can be added. Another concept is stability. It was introduced in [Föllmer & Schied, 04]. Here for two measures $\mathbb{P}$ and $\mathbb{Q}$ in the set of measures and every stopping time τ the measure that takes $\mathbb{P}$ up to τ and $\mathbb{Q}$ afterwards also lies in the set. The last concept is time-consistency which was introduced in [Delbaen, 03]. This property demands that at every stopping time density processes can be consistently pasted together. A more formal definition of this specific property will be given in the next section.

In the above cited paper Riedel among other things constructed time-consistent sets of measures via their density processes. Consequently the question arose if in this special setting all time-consistent sets of measures can be constructed in this way. That is why we took a closer look at time-consistent sets of measures and found out that not quite all sets are of this kind. However a slight modification of his construction does the trick.

The main content of this paper is this alternative characterization of time-consistent sets. They are described via a set of predictable processes with specific properties. This will be our first and main theorem. In addition to showing how the set of measures can be related to this set of processes we will also show that sets of processes with the assumed properties define sets of time-consistent measures. This will be the content of our second theorem. So altogether we will provide an equivalent formulation for time-consistent sets of measures.

The build-up of this paper will be the following. After pinning down the model framework and specifying the attributes of our sets more precisely in Section 2.2 we will deduct the first theorem in the succeeding Section 2.3. Then in Section 2.4 we will commit ourselves to proving the second theorem. In the following fifth section we will introduce some example setting where our results are applicable an might simplify calculations. After that we discuss possible extensions in Section 2.6 and then conclude in the last and seventh section.

2.2 Model

To specify the setting we start with a discrete set $\Omega = \{\omega_1, ..., \omega_k\}$. On this state space we have an information structure $\{\mathbf{F}_t\}_{t=0,...,T}$ with $\mathbf{F}_0 = \Omega$ and $\mathbf{F}_T = \{\{\omega_1\}, ..., \{\omega_k\}\}$. This is a sequence of partitions of Ω, which become finer as time progresses, i.e. every set of $\mathbf{F}_{t+1}$ is a subset of some set of $\mathbf{F}_t$ for all t.

Heuristically this concept describes the information of the prevailing state available at a certain time t. This means for a fixed time t the decision maker will not necessarily be able to observe the exact state which occurs but merely which subset of $\mathbf{F}_t$ is realized. If the observed subset consists of only a single state then of course the decision maker has full knowledge of the realization.

If you want to express this in terms of σ-fields and filtrations you just take the power set $Pot(\Omega)$ for the filtration $\mathcal{F}$ and define the filtration $\{\mathcal{F}_t\}_t$ by setting $\mathcal{F}_t := \sigma(\mathbf{F}_t)$ i.e. $\mathbf{F}_t$ is the set of atoms generating $\mathcal{F}_t$.

For our considerations we assume our information structure to have a constant and finite splitting function with splitting value ν. This implies that if you draw the filtration as an information tree it will have the same finite number of branches at every vertex. Formally the splitting function f of an information structure $\{\mathbf{F}_t\}_t$ is defined in the following way

$$f : \Omega \times [0, \infty) \to \mathbb{N} \quad , \quad f(\omega, t) = \sharp\{A \in \mathbf{F}_{t+1} \mid A \subseteq \mathbf{F}_t(\omega)\}$$

where $\mathbf{F}_t(\omega)$ is the set $B \in \mathbf{F}_t$ with $\omega \in B$. The finiteness of this index will allow us to apply the martingale representation given in Theorem 5.15 in [Dothan, 90] and the constancy will result in unique processes in the representation. We will make these two things more precise in the following section.

For now we will also restrict this model to a finite time horizon $[0, T]$. The finite splitting index and the finite time horizon result in a finite Ω.

To complete our probability space we still need to fix a probability measure $\mathbb{P}_0$ as a reference measure which pins down the sets of measure zero. Since we are on a tree like structure any measure which assigns non-zero probability to each branch will do, for simplicity let us choose the uniform distribution.

The set of measures we want to characterize will be denoted by $\mathcal{P}$. In the following we will make some assumptions on this set and justify their plausibility.

Our first assumption will be

Assumption 2.2.1. *We assume $\mathbb{P}_0 \in \mathcal{P}$ and for all other measures $\mathbb{P} \in \mathcal{P}$ $\mathbb{P}(A) > 0$ for all $A \in \mathbf{F}_T$*

In this assumption $\mathbb{P}_0$'s function as a reference measure becomes clear. One can see that it has no influence on the stochastic structure of the other measures. It simply implies that all measures contained in $\mathcal{P}$ have the same null sets which means that we know what sure and impossible events are.

In [Epstein & Marinacci, 06] an economic interpretation of this assumption was given. They related it to an axiom on preferences first postulated in [Kreps, 79]. He claimed that if a decision maker is ambivalent between an act x and $x \cup x'$ then he should also be ambivalent between $x \cup x''$ and $x \cup x' \cup x''$. Meaning if the possibility of choosing x' in addition to x brings no extra utility compared to just being able to choose x, then also no additional utility should arise from being able to choose x' supplementary to $x \cup x''$.

In our second assumption we claim

Assumption 2.2.2. *$\mathcal{P}$ is time-consistent. This means for a stopping time τ and densities $p_t := \left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t$ and $q_t := \left(\frac{d\mathbb{Q}}{d\mathbb{P}_0}\right)_t$ belonging to $\mathbb{P}, \mathbb{Q} \in \mathcal{P}$ that the measure $\tilde{\mathbb{P}}$ defined by the density*

$$\left(\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}_0}\right)_t = \begin{cases} p_t & \text{if } t \leq \tau \\ \frac{p_\tau q_t}{q_\tau} & \text{else} \end{cases}$$

belongs to $\mathcal{P}$ as well.

As mentioned in the introduction this assumption also originates from a feature claimed for preferences introduced in [Epstein & Schneider, 03]. They expanded the Multiple Priors Model (cp [Gilboa & Schmeidler, 89]) to a dynamic setting and asked the decision maker to be dynamically consistent in his decisions. With this they meant that if two acts are identical up to some time t but in $t+1$ the one is preferred over the other, then this should already be the case at time t. This implies that a decision maker will never regret his earlier decisions. In their paper Epstein and Schneider then showed that preferences fulfill this requirement if and only if the utility functional one obtains contains a rectangular set of measures. Rectangularity is equivalent to time-consistency. Time-consistency was introduced in [Delbaen, 03] where he also showed the equivalence to rectangularity. These two features stand for being able to judge each period in time with a different measure. More technically they allow to consistently paste together different densities at different times and still stay in the set. They also make it possible to use backward induction in discrete settings and allow for a Law of Iterated Expectations.

The set used to characterize $\mathcal{P}$ will be denoted by $\mathcal{A}$. We will show that it consists of predictable processes, is compact and that the process constant to zero is contained in it. Furthermore we will see that it fulfills a property we call stable under pasting and define in the following way.

Definition 2.2.3. *A set of processes $\mathcal{A}$ is called stable under pasting if for every stopping time τ and all processes $(\alpha_t)_t, (\beta_t)_t \in \mathcal{A}$ the process defined by*

$$\gamma_t := \begin{cases} \alpha_t & \text{if } t \leq \tau \\ \beta_t & \text{else} \end{cases}$$

belongs to $\mathcal{A}$ as well.

Later on we will show if we assume these properties for a set $\mathcal{A}$ then we can derive a set of measures $\mathcal{P}$ that features our original characteristics.

2.3 From $\mathcal{P}$ to $\mathcal{A}$

The goal of this section is to prove the main theorem of this paper, which tells us, that every time-consistent set of measures in our setting can also be described via a set of predictable processes fulfilling certain properties.

Expressed more formally this results in

Theorem 2.3.1. *For every set of measures $\mathcal{P}$ satisfying Assumptions 2.2.1, 2.6.1 and 2.2.2 there is a set of predictable processes $\mathcal{A}$ such that*

$$\mathcal{P} = \left\{ \mathbb{P} \;\middle|\; \left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t = \tilde{\mathcal{E}}_t(\alpha)\ ,\ \alpha \in \mathcal{A}^t, t \in \{0, ..., T\} \right\} \qquad \textit{where}$$

$$\tilde{\mathcal{E}}_t(\alpha) = \exp\left(\sum_{s=1}^{t}\sum_{h=1}^{\nu-1} \alpha_{hs}\Delta m_{hs} - \sum_{s=1}^{t} \ln \mathbb{E}\left[\exp\left(\sum_{h=1}^{\nu-1} \alpha_{hs}\Delta m_{hs}\right)\right]\right)$$

The $\mathcal{A}$ resulting from each $\mathcal{P}$ inhabits following features:

- $0 \in \mathcal{A}$
- $\mathcal{A}$ *is compact.*
- $\mathcal{A}$ *is stable under pasting.*

In order to prove this theorem we will derive a set of predictable processes $\mathcal{A}$ for every time-consistent set $\mathcal{P}$ and then show that it inhabits the requested features. One important step along this way will be a martingale representation theorem which we will explain more thoroughly in the next subsection. After that we will show the construction of the processes starting with an arbitrary time-consistent set of measures satisfying the above assumptions. Following this we will show that the constructed processes really are what we asked for.

2.3.1 Martingale Representation

This important tool which we will use in our proof tells us that in our setting we can find a set of martingales with which we can represent every other martingale in our setting with the help of predictable processes. A set of martingales which has this representation property is called a martingale basis. More formally we define

Definition 2.3.1. *A finite set of martingales* $\{m_{1t}\}, ..., \{m_{kt}\}$ *is called a basis iff for every martingale* $\{x_t\}$ *there are predictable processes* $\{\alpha_{1t}\}, ..., \{\alpha_{kt}\}$ *such that for every* $1 \leq t \leq T$

$$x_t = x_0 + \sum_{h=1}^{k} \sum_{s=1}^{T} \alpha_{hs} \Delta m_{hs} \qquad \textit{where } \Delta m_{hs} = m_{hs} - m_{h,s-1}$$

If the martingales $\{m_{1t}\}, ..., \{m_{kt}\}$ *are pairwise orthogonal, i.e. for every* $1 \leq j \leq k$, $1 \leq h \leq m$, $j \neq h$ *and every* $0 \leq t \leq T$, $\langle m_j, m_h \rangle_t = 0$*, then the basis* $\{m_{1t}\}, ..., \{m_{kt}\}$ *is called orthogonal.*

For our purposes it would be good to know in which cases such a basis exists especially with unique α's. An answer for this is provided by the following proposition. A slightly different version of this can be found in [Dothan, 90] but since we are looking for a unique representation we need to restrict the setting to a constant splitting function of our information structure. The proof is works along the same line as the one in [Dothan, 90].

Proposition 2.3.2. (Martingale Representation)
Given a discrete space $\Omega = \{\omega_1, ..., \omega_k\}$ *which is endowed with an information structure* $\{\mathbf{F}_t\}_{t=0,...,T}$ *with* $\mathbf{F}_0 = \Omega$ *and* $\mathbf{F}_T = \{\{\omega_1\}, ..., \{\omega_k\}\}$ *and a constant splitting function with value* ν*. Then there exists an orthogonal martingale basis* $m_{1t}, ..., m_{\nu-1,t}$ *for which the predictable processes* $\{\alpha^x_{1t}\}, ..., \{\alpha^x_{\nu-1,t}\}$ *in the representation of every* $\{x_t\}$ *are unique.*

Remark 2.3.3. *Since under the assumption of "no arbitrage" discounted assets are martingales for a martingale measure* $\mathbb{P}^*$ *this means for a binomial* *tree setting that there is one asset M_t with which every other asset X_t can be replicated and therefore hedged. More general in an n-nomial tree we can replicate every asset with a set of $n-1$ many assets.*

2.3.2 Exponential form of the densities

The next step we will take is to show that every measure $\mathbb{P} \in \mathcal{P}$ can be uniquely related to predictable processes $\left(\alpha^{\mathbb{P}}_{1s}\right)_s, ..., \left(\alpha^{\mathbb{P}}_{\nu-1,s}\right)_s$.

Remark that this is exactly one process less than our splitting value ν.

The equivalence of the measures in addition to $\mathbb{P}_0 \in \mathcal{P}$ (Ass.2.2.1) gives us the possibility to identify each $\mathbb{P} \in \mathcal{P}$ uniquely with its density with respect to $\mathbb{P}_0$.

If you define $\left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t := \mathbb{E}\left[\left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)\Big|_{\mathcal{F}_t}\right]$ for every $t \leq T$ and every $\mathbb{P} \in \mathcal{P}$ with the expectation taken under $\mathbb{P}_0$ you obtain density processes which are $\mathbb{P}_0$-martingales.

Using Jensen's inequality and Doob's decomposition theorem each of the above densities can be written in the following form where $(M_t)_t$ is also a $\mathbb{P}_0$-martingale and $(A_t)_t$ is a non-decreasing and predictable process with $A_0 = 0$

$$\left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t = \exp(M_t - A_t).$$

Now applying the martingale representation theorem to M_t we obtain an orthogonal martingale basis $(m_{1s})_s, ..., (m_{\nu-1,s})_s$. This implies that there are predictable processes $\left(\alpha^{\mathbb{P}}_{1s}\right)_s, ..., \left(\alpha^{\mathbb{P}}_{\nu-1,s}\right)_s$ such that our densities can now be written in the following manner where $\Delta m_{hs} = m_{hs} - m_{h,s-1}$

$$\left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t = \exp\left(\sum_{s=1}^{t}\sum_{h=1}^{\nu-1} \alpha^{\mathbb{P}}_{hs}\Delta m_{hs} - A_t\right).$$

Now we still have to determine the A_t's. Using the martingale property of the densities and the measurability of the A_t's we receive the following recursive relation

$$A_{t+1} - A_t = \ln \mathbb{E}\left[\exp\left(\sum_{h=1}^{\nu-1} \alpha^{\mathbb{P}}_{h,t+1} \Delta m_{h,t+1}\right) \;\middle|\; \mathcal{F}_t\right].$$

This results in

$$A_t = \sum_{s=1}^{t} (A_s - A_{s-1}) = \sum_{s=1}^{t} \ln \mathbb{E}\left[\exp\left(\sum_{h=1}^{\nu-1} \alpha^{\mathbb{P}}_{hs} \Delta m_{hs}\right) \;\middle|\; \mathcal{F}_{s-1}\right].$$

Additionally thanks to the assumptions on our information structure, we can show that our filtration is generated by our martingale basis and this in addition to the predictability of the α's allows us to drop the conditioning on $\mathcal{F}_{s-1}$.

So for our density $\left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t$ we now have following representation

$$\left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t = \exp\left(\sum_{s=1}^{t}\sum_{h=1}^{\nu-1} \alpha^{\mathbb{P}}_{hs} \Delta m_{hs} - \sum_{s=1}^{t} \ln \mathbb{E}\left[\exp\left(\sum_{h=1}^{\nu-1} \alpha^{\mathbb{P}}_{hs} \Delta m_{hs}\right)\right]\right). \tag{2.1}$$

This construction now allows us to not only identify a measure $\mathbb{P}$ with its density with respect to $\mathbb{P}_0$ and the associated density process but also with the predictable processes in the above representation $\left(\alpha^{\mathbb{P}}_{1s}\right)_s, ..., \left(\alpha^{\mathbb{P}}_{\nu-1,s}\right)_s$. Consequently it gives us a mapping from our density processes to sets of predictable processes.

For notational convenience and in resemblance to a stochastic exponential we will denote the right hand side of (2.1) as $\tilde{\mathcal{E}}_t(\alpha^{\mathbb{P}})$ seeing $\alpha^{\mathbb{P}} = (\alpha^{\mathbb{P}}_1, ..., \alpha^{\mathbb{P}}_{\nu-1})$ as a $\nu - 1$-dimensional process.

So now if we denote the set of processes generated via this construction and the densities up to time t by

$$\mathcal{A}^t := \left\{\left(\alpha^{\mathbb{P}}_{1,s}, ..., \alpha^{\mathbb{P}}_{\nu-1,s}\right)_{s\in\{0,...,t\}} \mid \mathbb{P} \in \mathcal{P}\right\} \qquad \text{and}$$

$$\mathcal{D}^t := \left\{\left(\left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_1, ..., \left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t\right) \mid \mathbb{P} \in \mathcal{P}\right\}$$

we have constructed a mapping $\tilde{\mathcal{E}}_t^{-1} : \mathcal{D}^t \to \mathcal{A}^t$.

From this construction and from the assumption that $\mathbb{P}_0 \in \mathcal{P}$ we directly conclude that the α's are predictable and that $0 \in \mathcal{A} := \mathcal{A}^T$.

2.3.3 Compact-valuedness of the α's

One further thing we want to show is that the compactness of the densities resulting from $\mathcal{P}$ implies compactness of $\mathcal{A}$. The compactness on $\mathcal{A}^t$ is defined via the norm $||\alpha||_{t,L^1} := \max\limits_{s\in\{0,\dots,t\}} ||\alpha_s||_{L^1}$.

This is a straight forward consequence of our assumptions and the preceding construction. In the construction of the α's every step was unique thanks to our assumptions. A density with respect to a designated measure uniquely characterizes a measure, the same is true for the construction of our density processes. Doob's decomposition is also unique and since we assumed a finite and constant splitting function the martingale representation also delivers unique predictable processes once the martingale basis is fixed. All in all the set of α's that belongs to one $\mathbb{P}$ is unique. Additionally a set of α's provides exactly one density and through that uniquely one measure. For this reason our $\tilde{\mathcal{E}}_t$ gives us a bijective mapping from the set of predictable processes $\mathcal{A}^t$ to our set of densities $\mathcal{D}^t$. This mapping is also continuous since the elements of our martingale basis are bounded thanks to the finite splitting index.

Since this also implies a continuous mapping between the densities and the predictable processes, the compactness on one side carries over to the other.

2.3.4 Stability under Pasting

The final property we claimed for our processes is stability under pasting. This property however follows directly from the assumption that $\mathcal{P}$ is time-consistent. To make this more clear for $(\alpha_t^{\mathbb{P}})_t, (\alpha_t^{\mathbb{Q}})_t \in \mathcal{A}$ and a stopping time $\tau \leq T$ define

$$\beta_t := \begin{cases} \alpha_t^{\mathbb{P}} & \text{if } t \leq \tau \\ \alpha_t^{\mathbb{Q}} & \text{else} \end{cases}$$

Our aim now is to show that this process lies in $\mathcal{A}$, i.e. that there exists a $\mathbb{P}^* \in \mathcal{P}$ such that $\left(\frac{d\mathbb{P}^*}{d\mathbb{P}_0}\right)_t = \tilde{\mathcal{E}}_t(\beta)$. If we plug β into Equation (2.1) and define $\mathbb{P}^*$ by

$$\left(\frac{d\mathbb{P}^*}{d\mathbb{P}_0}\right)_t := \begin{cases} \tilde{\mathcal{E}}_t(\alpha^{\mathbb{P}}) & \text{if } t \leq \tau \\ \frac{\tilde{\mathcal{E}}_t(\alpha^{\mathbb{Q}})\tilde{\mathcal{E}}_\tau(\alpha^{\mathbb{P}})}{\tilde{\mathcal{E}}_\tau(\alpha^{\mathbb{Q}})} & \text{else.} \end{cases}$$

we notice that $\beta \in \mathcal{A}$ is equivalent to $\mathbb{P}^* \in \mathcal{P}$. The fact that $\mathbb{P}^* \in \mathcal{P}$ however follows directly from our assumption of time-consistency.

If we now combine the above propositions we have shown Theorem 2.3.1.

2.4 Necessity

Now let us look at the conversion of the theorem above. The goal of this section will be to show that every $\mathcal{A}$ with the above properties defines a time-consistent set of measures. So we see that the properties of $\mathcal{A}$ are not only sufficient but also necessary. For this purpose we will derive a set of measures $\mathcal{P}$ from a given set $\mathcal{A}$ of predictable processes which are assumed to be compact-valued and stable under pasting. Additionally we claim that $\mathcal{A}$ contains the process constant to zero. Our goal will be to verify that the derived $\mathcal{P}$ satisfies the assumptions made in the model specifications.

Formally this will lead to following theorem

Theorem 2.4.1. *For every set of predictable processes $\mathcal{A}$ that satisfies the properties shown in Theorem 2.3.1 there exists a set of measures $\mathcal{P}$, such that*

$$\mathcal{A} = \left\{ \alpha \mid \left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t = \tilde{\mathcal{E}}_t(\alpha) \quad , \quad \mathbb{P} \in \mathcal{P} \right\}.$$

Every $\mathcal{P}$ constructed in this way has the following properties:

- $\mathbb{P}_0 \in \mathcal{P}$ *and* $\mathbb{P} \sim_{loc} \mathbb{P}_0$ *for all* $\mathbb{P} \in \mathcal{P}$
- $\mathcal{P}$ *is compact*
- $\mathcal{P}$ *is time-consistent.*

2.4.1 Construction of $\mathcal{P}$

If we use the same identification as in part 2.3.2 between the processes $(\alpha_t)_{t\in\{0,\dots,T\}}$ and the densities we are able to construct a density process $\left(\left(\frac{d\mathbb{P}^\alpha}{d\mathbb{P}_0}\right)_t\right)_t$ for every $\alpha \in \mathcal{A}$.

From the construction it follows immediately that the obtained processes are $\mathbb{P}_0$-martingales with expectation 1 and since the they are clearly strictly larger than zero they are indeed density processes.

Let us define our new set of measures by

$$\mathcal{P} := \left\{ \mathbb{P} \;\middle|\; \frac{d\mathbb{P}}{d\mathbb{P}_0}\Big|_{\mathcal{F}_t} = \tilde{\mathcal{E}}_t(\alpha) \text{ for } \alpha \in \mathcal{A} \right\}.$$

Since the process $\alpha \equiv 0$ is assumed to be an element of $\mathcal{A}$ we get that $\mathbb{P}_0 \in \mathcal{P}$. From the fact that all $\mathbb{P} \in \mathcal{P}$ are constructed via density processes with respect to $\mathbb{P}_0$ that are strict positive we can also directly conclude that our measures are all equivalent to our reference measure.

2.4.2 Time-Consistency

As when showing that we can derive $\mathcal{A}$ from $\mathcal{P}$ time-consistency in our set $\mathcal{P}$ is equivalent to stability under pasting in our set $\mathcal{A}$ and thus this property follows instantly from our assumptions.

2.4.3 Compactness of densities

Here again the fact that the $\tilde{\mathcal{E}}$ is a bijective and continuous mapping is the reason why the compactness of the α's implies compactness of the densities.

And again summarizing the above propositions leads us to the proof of Theorem 2.3.1.

2.5 Examples

In this section we introduce some examples for which this result is applicable and might simplify calculations.

2.5.1 Binomial Tree

The most basic example one can think of in this setting is a binomial tree. It has a constant and finite splitting index of two. Here things are still very basic to calculate. One can for instance show that a convex set of priors results in a convex set of processes and vice versa which is in general not true for a higher splitting index. Put more formally we have

Proposition 2.5.1. *On a binomial tree every convex set of measures fulfilling Assumptions 2.2.1 and 2.2.2, i.e.* $\mathcal{P} = \{(p_1, ..., p_T) \; | p_t \in [\underline{p_t}, \overline{p_t}]$ *for all* $t = \{0, ..., T\}\}$*, is equivalent to the respective processes lying in a predictable interval* $[a_t, b_t]$*, where* $p_t = P[X_t = up \; |\mathcal{F}_{t-1}]$.

Proof. For the proof we will work ourselves through the tree successively for every time period t.

Starting with $t = 1$ the density for a fixed $\mathbb{P}$ takes following form

$$\frac{d\mathbb{P}}{d\mathbb{P}_0}\bigg|_{\mathcal{F}_1}(\text{up}) = 2p = \frac{2\exp(\alpha\Delta m_1(\text{up}))}{\exp(\alpha\Delta m_1(\text{up})) + \exp(\alpha\Delta m_1(\text{down}))}$$

this can be transformed to

$$\alpha = \ln\left(\frac{1-p}{p}\right)(m_1(\text{down}) - m_1(\text{up}))^{-1}$$

which is a function that is monotone and continuous in p. So if $p \in [\underline{p}, \overline{p}]$ then this results in boundaries a, b which are $\mathcal{F}_0$-measurable s.t. $\alpha \in [a, b]$.

One can show the conversion by the same argumentation since the above formula can be converted to a function $p(\alpha)$ which is also monotone and continuous in α. Therefore a convex set of α's gives us a convex set of probabilities $[\underline{p_t}, \overline{p_t}]$ where $\underline{p_t} = \inf_{\mathbb{P}\in\mathcal{P}} \mathbb{P}[X_t = \text{up} \; |\mathcal{F}_{t-1}]$.

This can easily be extended to further time periods by just looking at the one step ahead measures or densities in an analogous way. □

[Chudjakow & Vorbrink, 09] present applications of this to american exotic options on a binomial tree.

2.5.2 Exponential Families

A further example for expressing time-consistent sets of measures via predictable processes was given in [Riedel, 09]. He introduced what he calls dynamic exponential families which is the discrete version of κ-ambiguity in [Chen & Epstein, 02] but with predictable bounds.

He starts with a probability state space $(S, \mathcal{S}, \nu_0)$ with $S \subset \mathbb{R}^d$. With this he constructs a probability space with $(\Omega, \mathcal{B}, (\mathcal{F}_t)_{t=1,\dots,T}, P_0)$, where

- $\Omega = S^T$
- $\mathcal{B} = \bigotimes_{t=1}^T \mathcal{S}$ σ-field generated by all projections $\epsilon_t : \Omega \to S$
- $(\mathcal{F}_t)$ generated by the sequence (ϵ_t)
- $\mathbb{P}_0 = \bigotimes_{t=1}^T \nu_0$ probability s.t. ϵ_t iid with distribution ν_0

Then by assuming that $\int_S e^{\lambda \cdot x} \nu_0(dx) < \infty$ the log-Laplace function $L(\lambda) = \log \int_S e^{\lambda \cdot x} \nu_0(dx)$ is well defined and with the help of predictable processes $(\alpha_t)_t$ he then defines densities on $(\Omega, \mathcal{B}, (\mathcal{F}_t)_t, \mathbb{P}_0)$ via

$$\mathcal{D}_t^\alpha := \exp\left(\sum_{s=1}^t \alpha_s \epsilon_s - \sum_{s=1}^t L(\alpha_s)\right).$$

Then for fixed predictable processes $a < b$ one gets a set of densities which defines a time-consistent set of measures by setting

$$\mathcal{P}^{a,b} = \left\{\mathbb{P} \mid \left(\frac{d\mathbb{P}}{d\mathbb{P}_0}\right)_t = \mathcal{D}_t^\alpha \ , \ \alpha \in [a, b]\right\}.$$

2.5.3 Trinomial Tree

The purpose of the following example is to show that switching between these two representations does not work too well in general. Starting with a two period trinomial tree which means we have a state space $\Omega = \{s_1, ..., s_9\}$ and the information structure $\mathbf{F}_0 = \Omega$, $\mathbf{F}_1 = \{\{s_1, s_2, s_3\}, \{s_4, s_5, s_6\}, \{s_7, s_8, s_9\}\}$ and $\mathbf{F}_2 = \{\{s_1\}, ..., \{s_9\}\}$ we define the rather simple time-consistent set

$$\mathcal{P} = \left\{ \left(\frac{1}{3} + \epsilon, \frac{1}{3} + \delta, \frac{1}{3} - \epsilon - \delta \right) \;\middle|\; \epsilon, \delta \in \left(-\frac{1}{3}, \frac{1}{3} \right) \text{ and } \epsilon + \delta \neq \frac{1}{3} \right\}.$$

We then construct a martingale basis in this tree with respect to the uniform distribution and then show what this set looks like expressed via predictable processes and our basis.

A martingale basis $\{m_t^1\}, \{m_t^2\}$ in this case is given by

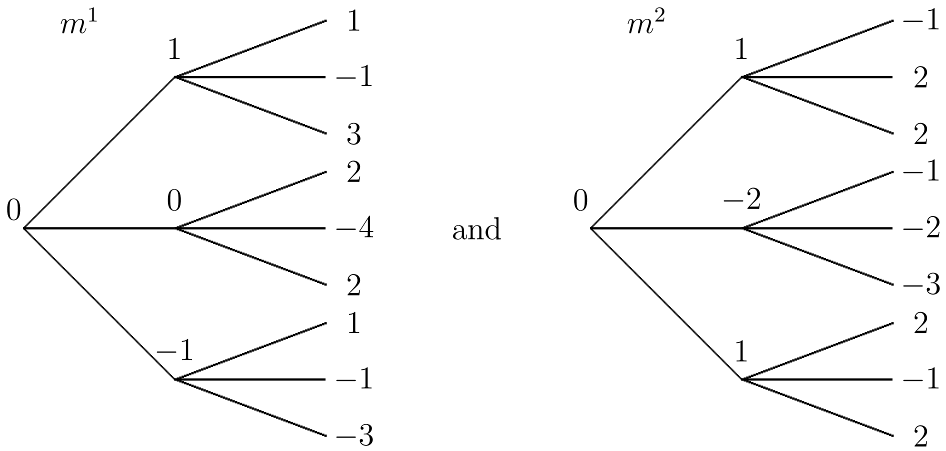

Figure 2.1: Martingale Basis

If we now calculate the processes that belong to each of the measures above we obtain
for $t = 1$ and $i = 1, ..., 9$

$$\alpha_1^1(s_i) = \frac{1}{2} \ln \frac{1 + 3\epsilon}{1 - 3\delta - 3\epsilon} \quad \text{and} \quad \alpha_1^2(s_i) = \frac{1}{3} \ln \frac{(1 + 3\epsilon)(1 - 3\epsilon - 3\delta)}{1 + 3\delta}$$

and for $t = 2$

$$\alpha_2^1(s_i) = \begin{cases} \frac{1}{4}\ln\frac{1-3\epsilon-3\delta}{1+3\delta} & \text{for } i = 1,2,3 \\ \frac{1}{6}\ln\frac{\sqrt{(1-3\epsilon-3\delta)(1+3\epsilon)}}{1+3\delta} & \text{for } i = 4,5,6 \\ \frac{1}{4}\ln\frac{1+3\epsilon}{1-3\epsilon-3\delta} & \text{for } i = 7,8,9 \end{cases}$$

$$\alpha_2^2(s_i) = \begin{cases} \frac{1}{3}\ln\frac{\sqrt{(1-3\epsilon-3\delta)(1+3\delta)}}{1+3\epsilon} & \text{for } i = 1,2,3 \\ \frac{1}{2}\ln\frac{1+3\epsilon}{1-3\epsilon-3\delta}1+3\delta & \text{for } i = 4,5,6 \\ \frac{1}{3}\ln\frac{\sqrt{(1+3\epsilon)(1-3\epsilon-3\delta)}}{1+3\delta} & \text{for } i = 7,8,9 \end{cases} .$$

As one can see a comparably simple set in the one representation can become relatively complicated in the other.

2.5.4 DTV@R

Another important area in which time-consistent sets of measures have been studied are risk measures. In [Artzner et al., 99] it is shown that every coherent risk measure ρ_t has a robust representation involving a set of measures $\mathcal{P}$, i.e.

$$\rho_t(X) = \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[X \mid \mathcal{F}_t\right].$$

Then in [Artzner et al., 07] it was shown that the family of dynamic risk measures $\rho = (\rho_t)_t$ is dynamically consistent iff the set $\mathcal{P}$ is time-consistent. [Roorda & Schumacher, 05] introduce dynamically consistent tail value at risk (DTV@R) as one of these time-consistent risk measures.

As the set $\mathcal{P}$ they take all measures $\mathbb{P}$ for which the one step ahead densities with respect to the reference measure $\mathbb{P}_0$ are bounded by $\frac{1}{\lambda}$ where $\lambda \in (0, 1]$ is the usual risk level. If we want to describe this in our characterization it gives us

$$\frac{\tilde{\mathcal{E}}_t(\alpha)}{\tilde{\mathcal{E}}_{t-1}(\alpha)} = \exp\left(\alpha_t \cdot \Delta m_t - \ln \mathbb{E}\left[\exp\left(\alpha_t \cdot \Delta m_t\right)\right]\right) = \frac{\exp\left(\alpha_t \cdot \Delta m_t\right)}{\mathbb{E}\left[\exp\left(\alpha_t \cdot \Delta m_t\right)\right]} \leq \frac{1}{\lambda}$$

for all $t = 1, ...T$ and all $\alpha \in \mathcal{A}$.

This allows to characterize the set $\mathcal{A}$ as soon as the martingale basis is fixed.

2.6 Possible Extensions

In this section we will discuss poaaible extensions which arise quite naturally.

2.6.1 Convexity

Since time-consistent sets are often used in optimization problems convexity of the sets is often assumed. It would be nice if this feature would carry over to the processes. Unfortunately this is not the case in general, as can be seen in the following counterexample.

Take for example a trinomial tree with states s_1, s_2 and s_3 and just one time period. As a reference measure we will fix

$$\mathbb{P}_0(s_1) = \frac{1}{2} \ , \ \mathbb{P}_0(s_2) = \frac{1}{4} \quad \text{and} \quad \mathbb{P}_0(s_3) = \frac{1}{4}.$$

A second measure will be given by

$$\mathbb{Q}(s_1) = \frac{1}{2} \ , \ \mathbb{Q}(s_2) = \frac{1}{8} \quad \text{and} \quad \mathbb{Q}(s_3) = \frac{3}{8}.$$

The density of $\mathbb{Q}$ with respect to $\mathbb{P}_0$ will then be

$$\frac{d\mathbb{Q}}{d\mathbb{P}_0}(s_1) = 1 \ , \ \frac{d\mathbb{Q}}{d\mathbb{P}_0}(s_2) = \frac{1}{2} \quad \text{and} \quad \frac{d\mathbb{Q}}{d\mathbb{P}_0}(s_3) = \frac{3}{2}.$$

Since we want to show that from a convex set of measures a non-convex set of processes can arise, let us define our set of measures via

$$\mathcal{P} := \text{convH}\left\{\mathbb{P}, \mathbb{Q}\right\}.$$

Then let us look at the set of processes $\mathcal{A}$ arising from this convex set, especially $\alpha^{\mathbb{P}_0}$ and $\alpha^{\mathbb{Q}}$. Now if $\mathcal{A}$ were a convex set, then every convex combination of $\alpha^{\mathbb{P}_0}$ and $\alpha^{\mathbb{Q}}$ has to be an element of $\mathcal{A}$. Since $\alpha^{\mathbb{P}_0}$ is zero, because we chose $\mathbb{P}_0$ as our reference measure we look at $\frac{1}{2}\alpha^{\mathbb{Q}}$. If we now calculate the associated density to this process, we see that it can never originate from a convex combination of our original measures and therefore $\frac{1}{2}\alpha^{\mathbb{Q}} \notin \mathcal{A}$ and hence $\mathcal{A}$ is not convex.

2.6.2 Infinite Horizon

When extending our statements to an infinite time horizon let us first remark that our model assumptions can all be transferred without complications. We will however need a further assumption on our set of measures. This assumption will be

Assumption 2.6.1. *The family of densities for a fixed t*

$$\mathcal{D}_t := \left\{ \frac{d\mathbb{P}}{d\mathbb{P}_0}\Big|_{\mathcal{F}_t} \mid \mathbb{P} \in \mathcal{P} \right\}$$

is weakly compact in $L^1(\Omega, \mathcal{F}, \mathbb{P}_0)$.

Technically this assumption ensures that when looking at expressions of the following kind $\inf\limits_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_\tau]$ the infimum is always attained for bounded stopping times τ. (cp. [Riedel, 09])

[Arrow, 71] already gives an economic interpretation of this property by claiming a feature of preferences which is related to this assumption in [Chateauneuf et al., 05]. The condition we need to ask of the preferences to obtain this feature is called Monotone Continuity. It means that if an act f is preferred over an act g then a consequence x is never that bad that there is no small p such that x with probability p and f with probability $(1-p)$ is still preferred over g. The same is true for good consequences mixed with g.

Critics tend to object to this assumption by saying that if the probability of dying is added to the better act f then surely the preferences have to be reversed. However if we take f for getting 100 dollars and g for getting nothing then having to drive 60 miles to get the 100 dollars and so adding a small probability of getting killed will normally not reverse the preferences.

Expressed formally this means for acts $f \succ g$, a consequence x and a sequence of events $\{E_n\}_{n\in\mathbb{N}}$ with $E_1 \supseteq E_2 \supseteq ...$ and $\cap_{n\in\mathbb{N}} E_n = \emptyset$ there exists an $\bar{n} \in \mathbb{N}$ such that

$$\begin{bmatrix} x \text{ if } s \in E_{\bar{n}} \\ f(s) \text{ if } s \notin E_{\bar{n}} \end{bmatrix} \succ g \quad \text{and} \quad f \succ \begin{bmatrix} x \text{ if } s \in E_{\bar{n}} \\ g(s) \text{ if } s \notin E_{\bar{n}} \end{bmatrix}.$$

The construction of the processes can also be maintained, since they are always constructed for a fixed time horizon up to a time t. That is also the reason why the mapping from our densities to our processes still inhabits the same features, i.e. it is continuous and bijective. Therefore in this case the compactness also carries over from one side to the other. It is also clear that stability under pasting is equivalent to time-consistency for an infinite horizon as well. So altogether our statements can smoothly be converted from a finite to an infinite time horizon.

2.6.3 Looser Assumptions on Splitting Function

Since our assumptions on the filtration are very restrictive, it would be nice if they could be relaxed in one way or another.

One way would be to give up the assumption of a constant splitting function. In this case however you run into the problem that the α's that arise from the martingale representation are no longer unique and with that the mapping no longer distinct and bijective.

A second way is allowing for the splitting value to become infinite. This however has the consequence that the martingale representation will not necessarily exist anymore.

2.7 Conclusions

For our special setting, i.e. discrete and with special assumptions on the information structure, we have constructed an alternative characterization for time-consistent sets of measures. We have shown that all sets of time-consistent sets of measures can be expressed by predictable processes and vice versa.

As can be seen in the extensions standard generalizations fail to work. So as far as I am concerned this is the most generalization that can be formulated in this setting.

For practical applications we have shown that for problems which can be modeled in the form of decision trees (with a constant number of branches e.g. trinomial trees) we now know what a time-consistent set of measures must look like expressed via predictable processes which might simplify calculations. So hopefully our construction will be helpful in the future e.g. for solving Optimal Stopping Problems which can be modeled in this framework.

Chapter 3

A Duality Theorem for Optimal Stopping Problems under Uncertainty

3.1 Introduction

Many investment problems arising in Economics, Finance or general Decision Theory do not have a fixed point in time when the decision must be made but a time period in which the choice of investment is possible. Especially if the decision is irreversible which means the investment cannot be recovered without considerable losses optimal timing is crucial. Important examples for these kind of investments are the market entry time of a firm, the optimal time to install a new technology or the exercise strategy of an American option but also at what bid to sell a house or which job offer to accept. This class of problems is called optimal stopping problems.

Each of these problems can be modeled by optimally trying to stop a stochastic process $(X_t)_t$ which describes the future random payoffs. Classical decision theory (e.g. [von Neumann & Morgenstern, 44]) proposes to maximize the expected payoff under a given distribution, i.e.

$$\text{maximize} \qquad \mathbb{E}[X_\tau] \qquad \text{among all stopping times } \tau.$$

But what if the distribution of $(X_t)_t$ is not (exactly) known or one is unsure about its true form. In this case the current literature on what is called ambiguity or uncertainty aversion often reverts to the multiple prior model introduced in [Gilboa & Schmeidler, 89] for the static case and in [Epstein & Schneider, 03] for a dynamic setting. They propose to look at a whole set of possible distributions and to take the worst expected value as a foundation of decision-making, i.e.

$$\text{maximize} \qquad \inf_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_\tau] \qquad \text{among all stopping times } \tau$$

where $\mathcal{P}$ is a set of measures with specific properties.

One reason this way of modeling decisions emerged was that empirical studies e.g. [Ellsberg, 61] gave substantial evidence that decision makers not only inhabit risk aversion but also uncertainty aversion.

What is the difference between risk and uncertainty? When talking about risk one means the randomness that is inherent in a given and fixed distribution while uncertainty or ambiguity describes a further source of randomness which springs from lacking knowledge of the correct distribution. This notion was first introduced in [Knight, 21], and hence is also often referred to as *Knightian Uncertainty.*

This ansatz can further be motivated using the framework of incomplete markets where the equivalent martingale measure is no longer unique and therefore one obtains a whole set of measures. Alternatively one can think of the set of measures being slight variations of the measure one thinks the right one, this corresponds to testing the robustness of a model (cp. [Hansen & Sargent, 01]).

The question we now want to study is, can the order of minimizing over the distribution and maximizing via the stopping time be switched. Since for several optimal stopping problems the classical solution, i.e. with one fixed and known distribution, is well studied the duality can be helpful in solving these problems under uncertainty. Therefore the main objective of this work is to study under what conditions it is possible to first maximize over all stopping times and then minimize over the distributions. Expressed more formally we will prove the following duality theorem

$$\sup_{\tau \in \mathcal{S}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_\tau] = \inf_{\mathbb{P} \in \mathcal{P}} \sup_{\tau \in \mathcal{S}} \mathbb{E}^{\mathbb{P}}[X_\tau].$$

Remark 3.1.1. *This problem can also be seen as a stochastic game between two players, while one player is the "maximizer" picking an optimal stopping time the second is the "minimizer" choosing the distribution. What our theorem then shows is that it is irrelevant in which order they make their decisions since the "value" process of this game is always the same.*

In this paper we prove in a rather general setting a minimax theorem for optimal stopping problems in continuous time under Knightian uncertainty and deduce an optimal stopping rule. More precisely for mild assumptions on the payoff process X (i.e. right continuous, class of $\mathcal{P}$-D, upper semicontinuous, adapted and an a.s. finite optimal stopping time) and rather standard assumptions on $\mathcal{P}$ (i.e. absolute continuity, weak compactness of densities and time-consistency) we obtain that

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

and that an optimal stopping strategy is given by

$$\rho_0^* = \inf\{ s \geq 0 | \operatorname*{ess\,sup}_{\tau \geq s} \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau | \mathcal{F}_s \right] = X_s \}.$$

We also show that this stopping time is the minimum of all optimal stopping rules for the classical solutions, i.e.

$$\rho_0^* = \inf_{\mathbb{P} \in \mathcal{P}} \rho_0^{\mathbb{P}} \qquad \text{, where} \qquad \rho_0^{\mathbb{P}} = \inf\{ s \geq 0 | \operatorname*{ess\,sup}_{\tau \geq s} \mathbb{E}^{\mathbb{P}}[X_\tau | \mathcal{F}_s] = X_s \}.$$

A great help in the proof of the main theorem is an explicit but general construction for time-consistent sets of measures introduced in [Delbaen, 03] and briefly reviewed in Section 3.4. He shows that every convex and time-consistent set of measures can be expressed with the help of a convex valued correspondence. He also shows that starting with a martingale and a convex valued corespondence, one can construct a convex and time-consistent set of measures.

In addition to the theorem with its implications we apply the results to special classes of payoff processes. First we look at the case where $(X_t)_t$ is either a multiple prior sub- or supermartingale, which leads to stopping at the last possible period if there is a finite time horizon for the submartingale or to stopping immediately in case of the supermartingale. After that we show how the theorem helps identify the worst case distribution in an adaption of κ-ambiguity[1] introduced in [Chen & Epstein, 02]. This allows for the ambiguous stopping problem to be transformed into a classical one.

Discrete versions of the theorem can be found in [Föllmer & Schied, 04] and [Riedel, 09] who also presents applications. In [Karatzas & Kou, 98] one can find the continuous time case for a finite time horizon and a strong focus on trading constraints. Unlike their paper we explicitly include the infinite time horizon and embed the explicit construction for all sets of time-consistent sets of measures introduced in [Delbaen, 03] into the proof.

The paper is organized in the following way. In the next section the model will be discussed in more detail and the assumptions we make justified. After that the problem that is to be solved is elaborated more thoroughly in Section 3.3. Section 3.4 contains a constructive description of the time-consistent sets we look at including the adapted version of κ-ambiguity. The succeeding Section 3.5 contains the proof of our main theorem which we apply

[1] Since we are going to look at problems with an infinite time horizon and the classical version of κ-ambiguity fails to satisfy the Novikov condition in $T = \infty$ we need to adapt this concept a bit.

to different stopping problems in Section 3.6 and Section 3.7 concludes.

3.2 Model

As a foundation for our model we begin with the filtered probability space $\left(\Omega, \mathcal{F}, \mathbb{P}_0, (\mathcal{F}_t)_{t\in[0,\infty]}\right)$, where the filtration satisfies the usual assumptions, also $\mathcal{F}_0$ is trivialand $\mathcal{F}$ the σ-field generated by the union of all $\mathcal{F}_t$. We will denote the class of all stopping times τ of the filtration $(\mathcal{F}_t)_t$ which satisfy $\mathbb{P}_0(\tau < \infty) = 1$ by $\mathcal{S}$ and those that are larger than or equal to a $t \in [0, \infty)$ by $\mathcal{S}_t := \{\tau \in \mathcal{S} \mid \tau \geq t\}$.

Further let $(X_t)_{t\in[0,\infty]}$ be a right continuous and adapted process describing the payoff from stopping. Our decision maker's task is to choose a stopping time τ of the filtration $(\mathcal{F}_t)_t$. If he chooses the stopping rule τ he gains the payoff $X_\tau(\omega) = X_{\tau(\omega)}(\omega)$ for $\omega \in \Omega$. His goal is to maximize his expected reward. Since our model is placed in an ambiguous setting our decision maker is uncertain about the true distribution of X. In order to capture the decision maker's uncertainty aversion we will use the Recursive Multiple Prior Model introduced in [Epstein & Schneider, 03]. As a consequence he considers a set of probability distributions $\mathcal{P}$ on $(\Omega, \mathcal{F})$ which he all assumes possible and his (minimax) expected reward for stopping in τ is given by

$$\inf_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau\right]. \tag{3.1}$$

Remark 3.2.1. *For simplicity we will only look at risk neutral decision makers. Plus we will not explicitly mention discounting or a special utility function.*

Since the expected payoff should be well-defined for all possible stopping times we introduce following notion of class $\mathcal{P} - D$ and assume this property for X.

Definition 3.2.2. *We say a right-continuous process* $\{X_t, \mathcal{F}_t; 0 \leq t < \infty\}$ *is* *of class $\mathcal{P}-D$ if*

$$\sup_{\mathbb{P}\in\mathcal{P}} \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}}\left[|X_\tau|\right] < \infty \qquad \textit{and}$$

$$\lim_{K\to\infty} \sup_{\mathbb{P}\in\mathcal{P}} \sup_{\tau\in\mathcal{S}} \int_{|X_\tau|>K} |X_\tau| \; d\mathbb{P} = 0.$$

Remark 3.2.3. *This property is a uniform integrability condition for ambiguous settings. We not only ask for uniform integrability under a fixed distribution $\mathbb{P}_0$ but under a whole set $\mathcal{P}$. However we only look at stopping times τ and not at the whole time set.*

A further property we want to assume for the payoff process X is upper semicontinuity.

Definition 3.2.4. *A stochastic process X is upper semicontinuous in expectation from the left with respect to the probability measure $\mathbb{P}_0$ if for any increasing sequence of stopping times $\{\tau_i\}_{i=1}^{\infty}$ converging to τ, we have*

$$\limsup_{i\to\infty} \mathbb{E}^{\mathbb{P}_0}[X_{\tau_i}] \leq \mathbb{E}^{\mathbb{P}_0}[X_\tau].$$

This ensures that the lower Snell envelope

$$V_\cdot = \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \operatorname*{ess\,sup}_{\tau\in\mathcal{S}_\cdot} \mathbb{E}^{\mathbb{P}}[X_\tau \mid \mathcal{F}_\cdot]$$

has a cadlag modification with the important consequence that for the stopping times $\rho_t := \inf\{s \geq t \mid X_s = V_s\}$ one actually obtains $V_{\rho_t} = X_{\rho_t}$. A thorough discussion of these results along with a different approach to Theorem 3.5.1 can be found in [Treviño, 09].

For the set $\mathcal{P}$ we will also make some assumptions. First of all, for mainly technical reasons we assume

Assumption 3.2.5. $\mathbb{P}_0 \in \mathcal{P}$ *and all other measures* $\mathbb{P} \in \mathcal{P}$ *are absolutely continuous with respect to* $\mathbb{P}_0$*, i.e.* $\mathbb{P}(A) = 0$ *if* $\mathbb{P}_0(A) = 0$ *for all* $A \in \mathcal{F}$*. Additionally we will ask for* $\mathcal{P}$ *to be convex, i.e. for* $\lambda \in (0,1)$ *and* $\mathbb{Q}, \mathbb{P} \in \mathcal{P}$ *we have* $\lambda\mathbb{Q} + (1-\lambda)\mathbb{P} \in \mathcal{P}$.

This assumption merely lets $\mathbb{P}_0$ fix some sets of measure zero and serve as a reference measure. This has no influence on the stochastic structure of the other measures. It simply implies that all measures contained in $\mathcal{P}$ have at least the same null sets as $\mathbb{P}_0$ which economically translates to the decision maker knowing some sure and impossible events. Technically it allows us to identify each measure $\mathbb{P} \in \mathcal{P}$ with its Radon-Nikodym density $\frac{d\mathbb{P}}{d\mathbb{P}_0}$ with respect to $\mathbb{P}_0$.

The second part of the assumption assures that the set satisfying the more stringent constraint of mutual continuity, i.e. $\mathcal{P}^e := \{\mathbb{Q} \in \mathcal{P} \mid \mathbb{Q} \sim \mathbb{P}_0\}$ lies dense in $\mathcal{P}$ since each $\mathbb{Q} \in \mathcal{P}$ can be approximated by elements in $\lambda\mathbb{Q} + (1-\lambda)\mathbb{P}_0 \in \mathcal{P}^e$. Therefore we achieve the same behavioral implications for our optimal stopping problem.

For the assumption of mutual continuity an interpretation was given in [Epstein & Marinacci, 06]. They related it to an axiom on preferences first postulated in [Kreps, 79]. He claimed that if a decision maker is ambivalent between an act x and $x \cup x'$ then he should also be ambivalent between $x \cup x''$ and $x \cup x' \cup x''$. Meaning if the possibility of choosing x' in addition to x brings no extra utility compared to just being able to choose x, then also no additional utility should arise from being able to choose x' supplementary to $x \cup x''$.

The second assumption for our set $\mathcal{P}$ will ensure that the infimum in (3.1) is always attained for bounded stopping times τ (cp. [Riedel, 09]). We assume

Assumption 3.2.6. *The family of densities*

$$\mathcal{D} := \left\{ \frac{d\mathbb{P}}{d\mathbb{P}_0} \mid \mathbb{P} \in \mathcal{P} \right\}$$

is weakly compact in $L^1(\Omega, \mathcal{F}, \mathbb{P}_0)$.

An economic interpretation of this property was given by [Arrow, 71] in claiming a feature of preferences which was related to this assumption in

[Chateauneuf et al., 05]. The condition we need to ask of the preferences to obtain this feature is called Monotone Continuity. It means that if an act f is preferred over an act g then a consequence x is never that bad that there is no small p such that x with probability p and f with probability $(1-p)$ is still preferred over g. The same is true for good consequences mixed with g.

Critics tend to object to this assumption by saying that if the probability of dying is added to the better act f then surely the preferences have to be reversed. However if we take f for getting 100 dollars and g for getting nothing then having to drive 60 miles to get the 100 dollars and so adding a small probability of getting killed will normally not reverse the preferences.

Expressed formally this means for acts $f \succ g$, a consequence x and a sequence of events $\{E_n\}_{n\in\mathbb{N}}$ with $E_1 \supseteq E_2 \supseteq ...$ and $\cap_{n\in\mathbb{N}} E_n = \emptyset$ there exists an $\bar{n} \in \mathbb{N}$ such that

$$\begin{bmatrix} x \text{ if } s \in E_{\bar{n}} \\ f(s) \text{ if } s \notin E_{\bar{n}} \end{bmatrix} \succ g \quad \text{and} \quad f \succ \begin{bmatrix} x \text{ if } s \in E_{\bar{n}} \\ g(s) \text{ if } s \notin E_{\bar{n}} \end{bmatrix}.$$

3.3 Problem

The question we want to study in the above setting is how to solve optimal stopping problems of the following form

$$\text{maximize } \inf_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_\tau] \text{ over all stopping times } \tau \in \mathcal{S}.$$

If $\mathcal{P}$ is singleton the problem reduces to subjective expected utility and the solution is well-known. We know an optimal solution is given by

$$\tau^* = \inf\left\{ s \geq 0 \mid X_s = \sup_{\tau\in\mathcal{S}_s} \mathbb{E}[X_\tau \mid \mathcal{F}_s] \right\}$$

see for example [El Karoui, 81]. As we will see later on the solution to our problem including uncertainty is very similar to this one.

In a discrete setting the problem was studied in [Riedel, 09]. He shows that with an added condition on $\mathcal{P}$ he attains following optimal stopping time as a result for the problem

$$\tau^* = \inf\{s \geq 0 \mid X_s = U_s\} \ \text{, where } U_s = \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_s} \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}[X_\tau \mid \mathcal{F}_s].$$

He also shows that under this extra assumption and with a finite time horizon U_s can be obtained recursively by setting

$$U_T = X_T \quad \text{and} \quad U_s = \max\left\{X_s,\ \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_{s+1} \mid \mathcal{F}_s]\right\} \quad \forall s = 0, ..., T-1.$$

This important condition on $\mathcal{P}$ is called time-consistency. It is a crucial assumption for making dynamically consistent decisions. In the discrete setting for instance it allows the use of backward induction. More general it implies following version of the law of iterated expectations:

$$\operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}\left[\operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}\left[X_\tau \mid \mathcal{F}_t\right] \mid \mathcal{F}_s\right] = \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}\left[X_\tau \mid \mathcal{F}_s\right] \qquad \text{for } t \geq s.$$

This can be interpreted in the following way, if the decision maker settles for the stopping rule τ then his expected return at time s is the same as the expectation in s of what he would get if he chose τ in t.

A technical formulation of this property can be found in the following final assumption on our set $\mathcal{P}$.

Assumption 3.3.1. $\mathcal{P}$ *is time-consistent. This means that for every stopping time* τ *and every pair* $\mathbb{P}^1, \mathbb{P}^2 \in \mathcal{P}$ *with density processes* $p^1_t = \frac{d\mathbb{P}^1}{d\mathbb{P}_0}$ *and* p^2_t*, respectively, the measure* $\mathbb{Q}$ *defined by the density process*

$$\left.\frac{d\mathbb{Q}}{d\mathbb{P}_0}\right|_{\mathcal{F}_t} = \begin{cases} p^1_t & \text{if } t \leq \tau \\ \frac{p^1_\tau p^2_t}{p^2_\tau} & \text{else} \end{cases}$$

belongs to $\mathcal{P}$ *as well.*

Further implications or equivalent definitions can inter alia be found in [Delbaen, 03] or [Riedel, 09].

3.4 Construction of $\mathcal{P}$

As is typical for these models we will use our first assumption on $\mathcal{P}$, the absolute continuity, to describe our set of measures since it allows us to identify each measure with its density function with respect to the reference measure $\mathbb{P}_0$. A further useful and well-known fact is that for a given martingale $(M_t)_t$ with respect to $\mathbb{P}_0$ and $(\mathcal{F}_t)_t$ densities can be generated by predictable processes since the stochastic exponential $\mathcal{E}(\theta \cdot M)$ describes a density if it is a non-negative martingale.

Remark 3.4.1. *As a reminder for a semimartingale X with $X_0 = 0$ the stochastic exponential is defined in the following way:*

$$\mathcal{E}_t(X) = \exp\left(X_t - \frac{1}{2}[X,X]_t\right) \prod_{0<s\leq t} (1 + \Delta X_s) \exp\left(-\Delta X_s + \frac{1}{2}(\Delta X_s)^2\right)$$

where $\Delta X = X_s - X_{s_-}$ and the infinite product converges.

In [Delbaen, 03] one can find a thorough study of such constructions, he shows that under certain assumptions time-consistent sets can always be represented in this way and gives a description of what the set of densities will look like. More explicit he proves for every convex and time-consistent set of densities $\mathcal{P}$ containing the reference measure $\mathbb{P}_0$, that if there exists a continuous martingale $(M_t)_t$ such that for all measures $\mathbb{P} \in \mathcal{P}^e$ there exists a predictable process θ such that $\mathbb{E}^{\mathbb{P}_0}\left[\frac{d\mathbb{P}}{d\mathbb{P}_0} \middle| \mathcal{F}_t\right] = \mathcal{E}_t(\theta \cdot M)$, then there exists a predictable, convex correspondence $C : \mathbb{R}_+ \times \Omega \to \mathcal{B}(\mathbb{R}^d)$ such that $0 \in C(t,\omega)$ for all (t,ω) and such that

$$\mathcal{P} = \text{cl } \mathcal{P}^e \quad \text{with} \quad \mathcal{P}^e = \left\{\mathbb{P}^\theta \mid \frac{d\mathbb{P}^\theta}{d\mathbb{P}_0} = \mathcal{E}(\theta \cdot M) \text{ with } \theta \in \Theta\right\}$$

where

$$\begin{aligned} \Theta &= \{\, \theta \mid \theta \text{ is a predictable process with } \theta(t,\omega) \in C(t,\omega), \\ &\qquad \text{s.t. } \mathcal{E}(\theta \cdot M) \text{ is a positive uniformly integrable martingale} \,\} \end{aligned}$$

and the closure is taken with respect to the L^1 norm on the space of the density processes.

Remark 3.4.2. *The uniform integrability of the stochastic exponential guarantees that the arising distributions are absolutely continuous with respect to $\mathbb{P}_0$ and the positivity, i.e. $\mathcal{E}_\infty(\theta \cdot M) > 0$, even guarantees equivalence.*

Remark 3.4.3. *Two important examples for these kind of sets are the extreme constructions where $C(t,\omega) = \{0\}$ and $C(t,\omega) = \mathbb{R}^d$. In the first case $\mathcal{P} = \{\mathbb{P}_0\}$, i.e. singleton, and in the second case $\mathcal{P}$ is the set of all absolutely continuous probability measures whose densities have the appropriate form.*

Remark 3.4.4. *Delbaen also shows the conversion of the theorem used above to describe our set. He shows that for a martingale M and a predictable convex correspondence $C : \mathbb{R}_+ \times \Omega \to \mathcal{B}(\mathbb{R}^d)$ satisfying $0 \in C(t,\omega)$ for all (t,ω) and that the projection of C onto the predictable range of M is closed*[2] *we can construct a time-consistent convex set of measures.*

3.4.1 κ-Ambiguity

An explicit example for the construction of such a set on the probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,\infty]}, \mathbb{P}_0)$ is given in the following. It is strongly related to κ-ambiguity introduced in [Chen & Epstein, 02], but since the classical definition does not fulfill the Novikov condition in $T = \infty$ we need to slightly adapt it.

Fix a $\mathbb{P}_0$-martingale $(M_t)_{t\in[0,\infty]}$ and the set of predictable processes

$$\Theta := \left\{(\theta_t)_{t\in[0,\infty]} \mid \theta \text{ predictable process with } |\theta_t| \le \kappa_t\right\}$$
$$\text{where} \quad \kappa_t = \begin{cases} 1 & t \le 1 \\ \frac{1}{t^2} & \text{else.} \end{cases}$$

[2]This assumption is made in order to deal with those density generators θ that are not identically zero but are such that $\theta \cdot M$ is zero. It guarantees that the elements lying in the closure of the constructed set that are equivalent to $\mathbb{P}_0$ also have the form of a stochastic exponential. For more details we again refer to [Delbaen, 03].

Here $C(t,\omega) = [-\kappa_t, \kappa_t]$ is a predictable and convex correspondence with $0 \in C(t,\omega)$ for all (t,ω).

Using the analysis found in [Czichowsky & Schweizer, 09] it follows that the projection of C onto the predictable range of M is closed and since κ_t tends to zero fast enough to fulfill the Novikov condition we get the positivity and the uniform integrability of the stochastic exponential. Therefore Theorem 1 in [Delbaen, 03] tells us that the L^1-closure of

$$\mathcal{P}^e = \left\{ \mathbb{P}^\theta \;\middle|\; \frac{d\mathbb{P}}{d\mathbb{P}_0} = \mathcal{E}(\theta \cdot M), \quad \theta \in \Theta \right\}$$

is a time-consistent and convex set.

The set $\mathcal{P} = \mathrm{cl}\ \mathcal{P}^{\mathrm{e}}$ is also a weakly compact set, since it is weakly closed, uniformly integrable, and bounded hence it is a set fulfilling all our assumptions.

Remark 3.4.5. *When constructing time-consistent sets in a setting with an infinite time horizon in this fashion one needs to consider that many martingales known to us when $t \in [0,\infty[$ no longer satisfy the martingale condition when $t = \infty$ is included. A prominent example for this is the geometric Brownian motion.*

3.5 Main Part

Before we come to the main theorem of the paper let us make sure the conditional expectations we will speak about are properly, i.e. $\mathbb{P}_0$-a.s., defined. Since our sets of measures constructed with the help of [Delbaen, 03] are convex we know that every measure $\mathbb{Q} << \mathbb{P}_0$ can be approximated by measures $\mathbb{Q}_n \sim \mathbb{P}_0$. Hence $\mathcal{P}^e$ lies dense in $\mathcal{P}$ and we define

$$\operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[X_\tau|\mathcal{F}_t\right] := \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}^e} \mathbb{E}^{\mathbb{P}}\left[X_\tau|\mathcal{F}_t\right].$$

To simplify notations we will define the value function of stopping under the worst case measure by

$$U_t := \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t} \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau | \mathcal{F}_t \right],$$

the value function of the interchanged problem by

$$V_t := \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t} \mathbb{E}^{\mathbb{P}} \left[X_\tau | \mathcal{F}_t \right]$$

and for fixed $\mathbb{P}$ we define

$$U_t^{\mathbb{P}} := \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t} \mathbb{E}^{\mathbb{P}} \left[X_\tau \mid \mathcal{F}_t \right].$$

The equations are all to be understood $\mathbb{P}_0$-almost surely. Additionally we will ask for the stopping times $\rho_t := \inf\{s \geq t \mid X_s = U_s\}$ to be $\mathbb{P}_0$-almost surely finite for all $t < \infty$.

The main statement of this paper is the following Duality Theorem which can help solve optimal stopping problems in the above model since it allows to interchange the infimum and supremum.

Theorem 3.5.1. *For $X, \mathcal{F}$ and $\mathcal{P}$ satisfying the above claims we get*

$$U_t := \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t} \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau | \mathcal{F}_t \right] = \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t} \mathbb{E}^{\mathbb{P}} \left[X_\tau | \mathcal{F}_t \right] =: V_t$$

Remark 3.5.1. *As is always the case with statements of duality one inequality is trivial. So in our case we only have to show that the left hand side is greater than or equal to the right hand side.*

Proof of the theorem. The proof will consist of two claims leading to the main statement. The first claim is
Claim 1:

$$V_t = \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_{\rho_t} | \mathcal{F}_t \right] \qquad \text{for all } 0 \leq t < \infty,$$
$$\text{where } \rho_t = \inf\{s \in [t, \infty) \mid X_s = V_s\}$$

and from this claim the theorem results at once.

In order to prove this claim first observe that $\rho_s \in \mathcal{S}$ and we immediately obtain that the l.h.s. is greater than or equal to the r.h.s. The opposite inequality remains to be shown.

To do so, we propose

Claim 2:

$$V_s \leq \mathbb{E}^{\mathbb{Q}}\left[V_\rho | \mathcal{F}_s\right] \qquad \text{for all } \mathbb{Q} \in \mathcal{P} \text{ and stopping times } s \leq \rho \leq \rho_s$$

This immediately yields the first claim by setting $\rho = \rho_s$.

Remark 3.5.2. *In the following the explicit construction of the set of measures as described in Section 3.4 will play a very prominent role. Instead of directly working with the distributions, we can restrict ourselves to the density generators.*

We start the proof of the second claim by first fixing an arbitrary measure $\mathbb{Q} \in \mathcal{P}$ whose density process regarding $\mathbb{P}_0$ is $\mathcal{E}(\theta \cdot M)_t$. Now look at a sequence $\{\theta_k\}_{k\in\mathbb{N}}$ where $\theta_k \in \Theta$ for all k and all θ_k coincide with θ on the stochastic interval $[\![s, \rho]\!]$ and which suffices following convergence

$$\lim_{k\to\infty} \operatorname*{ess\,sup}_{\tau\in\mathcal{S}_t} \mathbb{E}^{\mathbb{P}^k}\left[X_\tau \mid \mathcal{F}_t\right] = V_t,$$

where $\mathbb{P}^k$ denotes the measure with the density $\mathcal{E}(\theta_k \cdot M)$.

For such a sequence to exist we need to show that the set

$$\Phi := \{\operatorname*{ess\,sup}_{\tau\in\mathcal{S}_t} \mathbb{E}^{\mathbb{P}}\left[X_\tau \mid \mathcal{F}_t\right] \mid \mathbb{P} \in \mathcal{P}\}$$

is directed downwards, i.e. for all $\phi_1, \phi_2 \in \Phi$ we have that $\phi_1 \wedge \phi_2 \in \Phi$. This follows at once from Lemma 17 of [Delbaen, 03], where he shows that

$$\tilde{\Phi} := \{\mathbb{E}^{\mathbb{P}}\left[X_\tau \mid \mathcal{F}_t\right] \mid \tau \geq t \text{ is a stopping time and } \mathbb{P} \in \mathcal{P}\}$$

is a lattice.

Remark 3.5.3. *In the respective proof one can see that time-consistency is a crucial assumption here.*

This sequence is dominated by $\text{ess sup}_{\mathbb{P}\in\mathcal{P}} \text{ess sup}_{\tau\in\mathcal{S}_t} \mathbb{E}^{\mathbb{P}}[X_\tau \mid \mathcal{F}_t]$, which is $\mathbb{Q}$-integrable due to X being of class $\mathcal{P}-D$. Now applying the Dominated Convergence Theorem and using the facts that the stopped process $\left(U^{\mathbb{P}_{t\wedge\rho_s}}\right)_{t\geq s}$ is a $\mathbb{P}$-martingale for $\rho_s^{\mathbb{P}} = \inf\{u \geq s \mid X_u = U_u^{\mathbb{P}}\}$ (cp e.g. [Karatzas & Shreve, 98]) and $\rho_t \leq \text{ess inf}_{\mathbb{P}\in\mathcal{P}} \rho_t^{\mathbb{P}}$ (cp subsequent lemma) we get the second claim through

$$\begin{aligned}
\mathbb{E}^{\mathbb{Q}}[V_\rho \mid \mathcal{F}_s] &= \mathbb{E}^{\mathbb{Q}}\left[\lim_{k\to\infty} \operatorname*{ess\,sup}_{\tau\in\mathcal{S}_\rho} \mathbb{E}^{\mathbb{P}^k}[X_\tau \mid \mathcal{F}_\rho] \,\Big|\, \mathcal{F}_s\right] \\
&= \lim_{k\to\infty} \mathbb{E}^{\mathbb{Q}}\left[\operatorname*{ess\,sup}_{\tau\in\mathcal{S}_\rho} \mathbb{E}^{\mathbb{P}^k}[X_\tau \mid \mathcal{F}_\rho] \,\Big|\, \mathcal{F}_s\right] \\
&= \lim_{k\to\infty} \mathbb{E}^{\mathbb{P}^k}\left[\operatorname*{ess\,sup}_{\tau\in\mathcal{S}_\rho} \mathbb{E}^{\mathbb{P}^k}[X_\tau \mid \mathcal{F}_\rho] \,\Big|\, \mathcal{F}_s\right] \\
&\geq \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[U_\rho^{\mathbb{P}} \big| \mathcal{F}_s\right] = \mathbb{E}^{\mathbb{P}^*}[U^{\mathbb{P}^*}_{\rho\wedge\rho_s} \mid \mathcal{F}_s] \\
&= U^{\mathbb{P}^*}_{s\wedge\rho_s} \geq \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \operatorname*{ess\,sup}_{\tau\in\mathcal{S}_s} \mathbb{E}^{\mathbb{P}}[X_\tau \mid \mathcal{F}_s] \\
&= V_s
\end{aligned}$$

where $\mathbb{P}^*$ denotes the minimizing $\mathbb{P}$ in the foregoing equation.

□

To complete the proof we still need to show

Lemma 3.5.4. *Defining*

- $\rho_t := \inf\{s \geq t \mid X_s = V_s\}$ *and*
- $\rho_t^{\mathbb{P}} := \inf\{s \geq t \mid X_s = U_s^{\mathbb{P}}\}$

then it holds that $\rho_t \leq \rho_t^{\mathbb{P}}$ *for all* $\mathbb{P} \in \mathcal{P}$.

Proof. We know following equations hold almost surely for all $\mathbb{P} \in \mathcal{P}$

1. $X_s < V_s$ and $X_{\rho_t} = V_{\rho_t}$ for all $t \leq s < \rho_t$
2. $X_s < U_s^{\mathbb{P}}$ and $X_{\rho_t^{\mathbb{P}}} = U^{\mathbb{P}}_{\rho_t^{\mathbb{P}}}$ for all $t \leq s < \rho_t^{\mathbb{P}}$

3. $V_s \leq U_s^{\mathbb{P}}$ for all $s \in [0, \infty]$

Now if we assume $\rho_t > \rho_t^{\mathbb{P}}$ it follows from 1 and 2 that

$$V_{\rho_t^{\mathbb{P}}} > X_{\rho_t^{\mathbb{P}}} = U^{\mathbb{P}}_{\rho_t^{\mathbb{P}}}$$

which clearly contradicts 3 and therefore $\rho_t \leq \rho_t^{\mathbb{P}}$ for all $\mathbb{P}$. □

Intuitively this means our "worst case" decision maker has a more pessimistic apprehension of the future than the other investors and hence he values the expected payoff lower. So he is more likely to accept what he has earlier because he expects less in the future.

With the help of the theorem we can now immediately show that the ρ_t in the proof gives us a solution to our optimal stopping problem and we obtain.

Corollary 3.5.5. *(i) U is the smallest multiple prior supermartingale with respect to $\mathcal{P}$ that dominates X.*

(ii) An optimal stopping rule is given by $\rho_0^* = \inf \{s \geq 0 | U_s = X_s\}$.

Proof. We first show that $(U_t)_{t\in[0,\infty]}$ is a multiple prior supermartingale. For $t \geq s$ we have

$$\begin{aligned} \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[U_t \,|\mathcal{F}_s\right] &= \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[\operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[X_{\rho_t}|\mathcal{F}_t\right] \,|\mathcal{F}_s\right] \\ &= \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[X_{\rho_t}|\mathcal{F}_s\right] \leq \operatorname*{ess\,sup}_{\tau\in\mathcal{S}_s} \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[X_\tau|\mathcal{F}_s\right] = U_s. \end{aligned}$$

Remark 3.5.6. *The second equality is again due to our time-consistency assumption.*

Next we show that it is the smallest multiple prior supermartingale dominating X. For this assume that W is another multiple prior supermartingale dominating X, then it holds that

$$\operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[X_\tau \,|\mathcal{F}_t\right] \leq \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[W_\tau|\mathcal{F}_t\right] \leq W_t \qquad \forall \tau \in \mathcal{S}_t$$

and in particular

$$U_t = \operatorname*{ess\,sup}_{\tau \in \mathcal{S}_t} \operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau | \mathcal{F}_t \right] \leq W_t.$$

The last thing we want to show is the optimality of ρ_0^*, this follows directly with the help of the first claim in the proof, since

$$\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_{\rho_0^*} \right] = V_0 = U_0 = \sup_{\tau \in \mathcal{S}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau \right] \geq \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau \right] \qquad \forall \tau \in \mathcal{S}.$$

□

3.6 Applications

3.6.1 Sub- and Supermartingales

With the help of this theorem we want to solve optimal stopping problems. Two straightforward examples are if the payoff process $(X_t)_{t \in [0,T]}$ is either a multiple prior sub- or supermartingale.

In the case of the multiple prior submartingale which means

$$\operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_t \mid \mathcal{F}_s] \geq X_s \qquad \text{for all } t \geq s$$

we can show that

$$\sup_{\tau \in \mathcal{S}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}} \left[X_\tau \right] = \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_T].$$

This means it is optimal to wait until the very last period to stop and the expected payoff is the expected payoff in the last period under the worst case measure.

To see why this is true first of all remark that thanks to the optional sampling theorem

$$\mathbb{E}^{\mathbb{P}}[X_T] = \mathbb{E}^{\mathbb{P}}[\underbrace{\mathbb{E}^{\mathbb{P}}[X_T \mid \mathcal{F}_s]}_{\geq X_s}] \geq \mathbb{E}^{\mathbb{P}}[X_s]$$

for all stopping times $s \leq T$ and fixed $\mathbb{P}$ since multiple prior submartingales are submartingales for all $\mathbb{P} \in \mathcal{P}$.

Therefore $\sup_{\tau \in \mathcal{S}} \mathbb{E}^{\mathbb{P}}[X_\tau] = \mathbb{E}^{\mathbb{P}}[X_T]$ for all $\mathbb{P}$ and we get

$$\sup_{\tau \in \mathcal{S}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_\tau] = \inf_{\mathbb{P} \in \mathcal{P}} \sup_{\tau \in \mathcal{S}} \mathbb{E}^{\mathbb{P}}[X_\tau] = \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_T]$$

and with that the above statement.

In the case of a multiple prior supermartingale it turns out that stopping immediately is optimal. Since X being a multiple prior supermartingale means

$$\operatorname*{ess\,inf}_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_t \,|\mathcal{F}_s] \leq X_s \qquad \text{for all } t \geq s$$

we get

$$\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_\tau] \leq X_0 \qquad \forall \tau \in \mathcal{S}$$

and obtain that stopping immediately is optimal.

3.6.2 Exploiting Monotonicity in the Drift

In the following let $\mu, \sigma, \theta : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}$ be continuous, bounded and adapted functions. For convenience we abbreviate the functions $\mu(t, X_t), \sigma(t, X_t)$ and $\theta(t, X_t)$ by μ_t, σ_t and θ_t. Let $b^\theta = \mu + \theta\sigma$ fulfill a Lipschitz condition, i.e.

$$|b^\theta_t(x) - b^\theta_t(y)| \leq K|x - y| \qquad \text{for a positive constant } K.$$

On top of this let σ satisfy $|\sigma| \geq \epsilon > 0$ and

$$|\sigma_t(x) - \sigma_t(y)| \leq h(|x - y|)$$

where $h :]0, \infty[\to]0, \infty[$ is a strictly increasing function with $h(0) = 0$ and $\int_{(0,\epsilon)} h^{-2}(u)du = \infty \quad \forall \epsilon > 0$.

Falling back on the example for constructing time-consistent sets via our adaption of κ-ambiguity in Section 3.4.1 we will in this section show how the theorem helps identify the worst case measure in this set in the case of a

Brownian motion with drift as payoff process. With this we mean that our payoff process has following dynamics

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t^0$$

where W^0 is a Brownian Motion with respect to our underlying probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P}_0)$.

So the problem we want to study is

$$\sup_{\tau \in \mathcal{S}} \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}[X_\tau]$$

for the above X and the $\mathcal{P}$ from Section 3.4.1. For simplicity we will restrict ourselves to one dimensional processes giving us the advantage that we can use a comparison theorem for our drifts later on.

The first step in our analysis will be to transfer the ambiguity implemented by our set of measures onto the process, we want to stop. The main tool for this can be found in the theory of weak solutions for stochastic differential equations (SDEs) and Markov processes (e.g. cp [Revuz & Yor, 91] and [Shiryayev, 78]).

If a SDE has two weak solutions

$$(X^i, W^i), (\Omega^i, \mathcal{F}^i, (\mathcal{F}_t^i)_t, \mathbb{P}^i) \quad \text{for} \quad i = 1, 2 \quad \text{with} \quad X^1 = X^2 = x,$$

meaning in our case that

$$dX_t^i = \mu(t, X_t^i)dt + \sigma(t, X_t^i)dW_t^i$$

where W^i is a Brownian motion with respect to $(\Omega^i, \mathcal{F}^i, (\mathcal{F}_t^i)_t, \mathbb{P}^i)$ for $i = 1, 2$, then we know that X^1 and X^2 have the same law, i.e.

$$\mathbb{P}^1[X^1 \in \Gamma] = \mathbb{P}^2[X^2 \in \Gamma] \quad \text{for } \Gamma \in \mathcal{B}(C(\mathbb{R}_+)) \ .$$

In our case this means if we define a further auxiliary process $(X_t^\theta)_t$ via

$$dX_t^\theta = (\mu_t + \theta_t\sigma_t)dt + \sigma_t dW_t^0$$

then with the help of Girsanov's theorem and the construction of our measures we get that for W^θ being the $\mathbb{P}^\theta$-Brownian motion defined by

$$W_t^\theta = W_t^0 - \int_0^t \theta_s ds$$

where $\theta \in \Theta$ is one of the density generators from Section 3.4.1 that

$$(X, W^\theta), (\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P}^\theta) \quad \text{and} \quad (X^\theta, W^0), (\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P}_0)$$

are both weak solutions to the same SDE and therefore

$$\mathbb{P}^\theta[X \in \Gamma] = \mathbb{P}_0[X^\theta \in \Gamma] \quad \text{for } \Gamma \in \mathcal{B}(C(\mathbb{R}_+)) \ .$$

This equality now allows us to shift the ambiguity from the set of measures to the payoff process since $(X_t)_t$ together with $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P}^\theta)$ is a Markov process and hence

$$s(t,x) = \sup_{\tau \geq t} \mathbb{E}^{\mathbb{P}^\theta}[X_\tau \mid X_t = x]$$

is the smallest excessive majorant of the function $g(z) = z$ and with respect to the process X.

Remark 3.6.1. *Since we are looking for the optimal stopping time from the beginning, we set $t = 0$ and drop it in the following.*

We also know the smallest excessive majorant of $g(x)$ with respect to X can be approximated by

$$v(x) = \lim_n \lim_N Q_n^N g(x),$$

where $Q_n^N g(x)$ is the N^{th} power of of the operator Q_n defined via

$$Q_n g(x) = \max\{g(x), T_{2^{-n}} g(x)\}.$$

Remark 3.6.2. *The operator Q_n reminds of backward induction since $g(x)$ is the payoff of stopping and $T_{2^{-n}} g(x) = \int g(y) \mathbb{P}^\theta(2^{-n}, x, dy)$ is the expected payoff in $t = 2^{-n}$.*

This is now where the uniqueness in law from above comes in since it implies

$$\begin{aligned}\mathbb{P}^\theta(2^{-n}, x, dy) &:= \mathbb{P}^\theta(X_{2^{-n}} \in dy \mid X_0 = x)\\ &= \mathbb{P}_0(X^\theta_{2^{-n}} \in dy \mid X^\theta_0 = x) = \mathbb{P}_0(2^{-n}, x, dy).\end{aligned}$$

Implying since $(X^\theta_t)_t$ together with $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P}_0)$ is also a Markov process that the smallest excessive majorants in both settings are identical and we have

$$\sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}^\theta}[X_\tau \mid X_0 = x] = \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}_0}[X^\theta_\tau \mid X_0 = x]$$

which allows to shift the uncertainty of the distribution of X to uncertainty of the true drift of X via

$$\sup_{\tau\in\mathcal{S}} \inf_{\mathbb{P}^\theta\in\mathcal{P}} \mathbb{E}^{\mathbb{P}^\theta}[X_\tau \mid X_0 = x] = \sup_{\tau\in\mathcal{S}} \inf_{\theta\in\Theta} \mathbb{E}^{\mathbb{P}_0}[X^\theta_\tau \mid X^\theta_0 = x] \quad \text{for all } x \in \Omega.$$

The construction of Section 3.4.1 now tells us that

$$\mu_t + \theta_t\sigma_t \geq \mu_t - \kappa_t\sigma_t$$

for all $\theta \in \Theta$, meaning that $\mu_t - \kappa_t\sigma_t$ is the smallest possible drift our pay-off process can have and with a comparison result for stochastic differential equations we obtain that

$$\mathbb{P}_0\left[X^\theta_t \geq X^{-\kappa}_t \text{ for all } t \geq 0\right] = 1$$

and with that

$$\begin{aligned}&\sup_{\tau\in\mathcal{S}} \inf_{\theta\in\Theta} \mathbb{E}^{\mathbb{P}_0}\left[X^\theta_\tau\right] = \inf_{\theta\in\Theta} \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}_0}\left[X^\theta_\tau\right]\\ &\geq \inf_{\theta\in\Theta} \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}_0}\left[X^{-\kappa}_\tau\right] = \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}_0}\left[X^{-\kappa}_\tau\right].\end{aligned}$$

Since the converse inequality follows directly from the theorem

$$\sup_{\tau\in\mathcal{S}} \inf_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\left[X_\tau\right] = \inf_{\mathbb{P}\in\mathcal{P}} \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}}\left[X_\tau\right] = \inf_{\theta\in\Theta} \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}_0}\left[X^\theta_\tau\right] \leq \sup_{\tau\in\mathcal{S}} \mathbb{E}^{\mathbb{P}_0}\left[X^{-\kappa}_\tau\right]$$

we obtain that the theorem helps us identify the worst case distribution in this case and our ambiguous stopping problem is simplified into a classical stopping problem with a known distribution, i.e.

$$\sup_{\tau\in\mathcal{S}}\inf_{\mathbb{P}\in\mathcal{P}}\mathbb{E}^{\mathbb{P}}[X_\tau]=\sup_{\tau\in\mathcal{S}}\inf_{\theta\in\Theta}\mathbb{E}^{\mathbb{P}_0}[X_\tau^\theta]=\sup_{\tau\in\mathcal{S}}\mathbb{E}^{\mathbb{P}_0}\left[X_\tau^{-\kappa}\right]=\sup_{\tau\in\mathcal{S}}\mathbb{E}^{\mathbb{P}^{-\kappa}}\left[X_\tau\right].$$

3.7 Conclusion

Confronted with an optimal stopping problem and not (exactly) knowing or being unsure of the true distribution of the payoff process X we find ourselves in the framework of [Epstein & Schneider, 03] who propose to solve the following problem

$$\text{maximize}\qquad \inf_{\mathbb{P}\in\mathcal{P}}\mathbb{E}^{\mathbb{P}}\left[X_\tau\right]\qquad\text{among all stopping times }\tau$$

where $\mathcal{P}$ is the set of all measures we think possible. Here we asked if and under what conditions it is possible to interchange the order of minimizing over the distributions and maximizing over all stopping times. The result is a minimax theorem under rather general assumptions on the payoff process X and standard assumptions on the set of measures $\mathcal{P}$. It incorporates an explicit but still universal construction given in [Delbaen, 03] for time-consistent sets of measures.

This theorem allows us to identify the optimal stopping time for payoff processes which are multiple prior sub- or supermartingales and have a finite investment horizon. It also helps us determine the worst case distribution in the setting of κ-ambiguity for options with an increasing payoff in the underlying and with this turns our ambiguous decision problem into a classical purely risky one.

It would be desirable to find further applications of this theorem. Since the crux of applying this theorem is mainly a minimization over the drift it would be preferable to find classical solutions that depend on the drift of the payoff process. However the drift has to be stochastic and since in classical solutions the drift is commonly assumed constant for simplicity, it might be a good idea to look at processes in a setting with stochastic interest rates.

Chapter 4

Learning for Convex Risk Measures with Increasing Information

4.1 Introduction

Reaching decisions concerning risky projects in a dynamic system, an agent faces new information consecutively influencing her assessment of risk instantaneously.

In this article, we answer the question how anticipation of risk evolves over time when an agent gathers information. We show that, in the limit, all uncertainty is revealed but risk remains if the agent perceives risk in terms of time-consistent dynamic convex risk measures and, hence, generalize the famous Blackwell-Dubins Theorem to convex risk measures. We then relax the time-consistency assumption and show the result to still be valid. Hereto, a fundamental assumption is existence of a reference distribution that fixes impossible and sure events by virtue of equivalence of distributions under consideration.

Coherent risk measures were introduced by virtue of an axiomatic ansatz in [Artzner et al., 99] in a static setting and have been generalized to a dynamic framework in [Riedel, 04]. Tangible problems in this setup are inter alia discussed in [Riedel, 09]. The equivalent theory of multiple prior preferences in a static setup is introduced in [Gilboa & Schmeidler, 89]; a dynamic generalization is given in [Epstein & Schneider, 03]. Applying coherent risk measures substantially decreases model risk as they do not assume a specific probability distribution to hold but assume a whole set of equally likely probability models. Moreover, they possess a simple robust representation. However, as they assume homogeneity, coherent risk measures do not account for liquidity risk. Though in financial applications, the Basel II accord requires a "margin of conservatism", coherent risk measures are far too conservative when estimating risk of a project as they result in a worst case approach. Furthermore, popular examples of risk measures, as e.g. entropic risk, are not coherent.

Hence, it seems worthwhile to consider a more sophisticated axiomatic system: [Föllmer & Schied, 04] introduce convex risk measures as a generalization of coherent ones relaxing the homogeneity assumption. Equivalently, [Maccheroni et al., 06a] generalize multiple prior preferences to variational preferences. Convex risk measures are applied to a dynamic setup in [Föllmer & Penner, 06] for a stochastic payoff in the last period or, equivalently, in [Maccheroni et al., 06b] in terms of dynamic variational preferences. [Cheridito et al., 06] applies dynamic convex risk measures to stochastic payoff processes. Given a set of possible probabilistic models, convex risk measures are less conservative than coherent ones. Dynamic convex risk measures as well as dynamic variational preferences possess a robust representation in terms of minimal penalized expectation. The minimal penalty, serving as a measure for uncertainty aversion, uniquely characterizes the risk measure or, respectively, the preference. Conditions on the minimal dynamic penalty characterize time-consistency of the dynamic convex risk measure.

A parametric learning model in an uncertain environment for dynamic coherent risk measures or, equivalently, dynamic multiple priors as introduced in [Epstein & Schneider, 03], is elaborated in [Epstein & Schneider, 07]. The main virtue of this article is to introduce learning based on experience to convex risk measures models. First, we try to introduce learning in a constructive approach: we design a minimal penalty function and plug it into the robust representation: Since the penalty might be seen as some inverse likelihood of a specific prior distribution, we first apply a quite simple and intuitive learning mechanism to the penalty. We calculate the likelihood of a distribution given past experience and use this as updated penalty. The intuition behind this approach is quite simple: observing good events, distributions of a payoff process that are "stochastically more dominated", i.e. put more weight on bad events, become more unlikely, i.e. have a higher penalty. However, besides its intuitive appeal, it turns out that this procedure does not result in a penalty function as it is backwards oriented and a penalty function, by definition, incorporates probability distributions of the future movement of the payoff process. In a second, more sophisticated approach, we model a penalty incorporating projections of "past" likelihoods on future distributions. Here, we make use of the conditional relative entropy as penalty function: we achieve a proper penalty that penalizes distributions according to "distance" from the "most likely" distribution serving as reference distribution. However, the convex risk measure in terms of this penalty turns out not to be time-consistent in general as shown by a counterexample. In [Epstein & Schneider, 07], time-consistency is not an issue as multiplicity of priors is not introduced in terms of multiple equally likely distributions of the payoff process as e.g. in [Riedel, 09] or [Maccheroni et al., 06a], but in terms of multiple distributions on the parameter space.

Our further approach is not constructive but takes the robust representation of a risk measure in terms of minimal penalty for granted. As the main result of this article we achieve a generalization of the famous Blackwell-Dubins Theorem in [Blackwell & Dubins, 62] from conditional probabilities to time-consistent dynamic convex risk measures. We pose a condition on the minimal penalty in the robust representation, always satisfied by coherent risk measures, forcing the convex risk measure to converge to the conditional expected value under the true underlying distribution. Intuitively, this result states that, eventually, the uncertain distribution is revealed or, in other words, uncertainty diminishes as information is gathered but risk remains. The agent, as she has learned about the underlying distribution, is again in the framework of being an expected utility maximizer with respect to the true underlying distribution. We have hence achieved *learning as an intrinsic property* of dynamic convex risk measures.

Our generalization of the Blackwell-Dubins Theorem serves as an alternative approach to limit behavior of time-consistent dynamic convex risk measures as the one in [Föllmer & Penner, 06]. The result particularly states the existence of a limiting risk measure. As an example we consider dynamic entropic risk measures or, equivalently, dynamic multiplier preferences. We, however, show a Blackwell-Dubins type result to hold, even if we relax the time-consistency assumption. Again, we obtain existence of a limiting risk measure but in a more general manner than [Föllmer & Penner, 06] for not necessarily time-consistent convex and coherent risk measures.

[Schnyder, 02] discusses H.P. Minsky's theory of financial instability, a huge portion of which is caused by herding on financial markets. Besides, herding is usually one of the major objections towards Basel II. Our result however shows that, in the long run, there is hardly any chance to circumvent herding behavior.

The article is considered in a parametric setting. However, the second part can be restated in a non parametric setting. It is structured as follows: The next section formally introduces the underlying probabilistic model. Section 4.3 elaborately discusses robust representation of dynamic (time-consistent) convex risk measures. Constructive approaches to learning in terms of dynamic minimal penalty as well as their shortcomings are stated in Section

4.4. Section 4.5 generalizes the Blackwell-Dubins Theorem to conditional expectations. The following two sections then apply this result to coherent and convex risk measures first in the time-consistent case and then in the case without time-consistency. Section 4.8 states examples. Then we conclude.

4.2 Model

For our model we start with a discrete time set $t \in \{0, ..., T\}$ where T is an infinite time horizon. We will now construct an underlying filtered reference space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_t, \mathbb{P}^{\theta_0})$ and define risky projects X:

We fix $(S, \mathcal{A})$ as a measurable space where S describes the possible states of the world at a fixed point in time t and define Ω to be all possible states of the world, formally the set of sequences of elements of S. For this let $S_t = S$ for all $t \in \{0, ..., T\}$ and then define $\Omega := \bigotimes_{t=0}^{T} S_t$. On this space let $\mathcal{F}$ be the product σ-field generated by all projections $\pi_t : \Omega \to S_t$ and let the elements of the filtration $\mathcal{F}_t$ be generated by the sequence $\pi_1, ..., \pi_t$. Additionally define all sequences up to time t by $\Omega^t := \bigotimes_{s=0}^{t} S_s$. Denote generic elements on these spaces by $s_t \in S_t$, $s \in \Omega$, and $s^t \in \Omega^t$.

Let Θ be a set of parameters where every $\theta \in \Theta$ uniquely defines a distribution $\mathbb{P}^\theta$ on $(\Omega, \mathcal{F})$ with filtration $(\mathcal{F}_t)_t$ and fix $\mathbb{P}^{\theta_0}$ as a reference distribution which can be seen as the true distribution of the states. For all $\theta \in \Theta$, $\mathbb{P}^\theta$ is assumed to be equivalent to $\mathbb{P}^{\theta_0}$. Let $\mathcal{M}^e(\mathbb{P}^{\theta_0})$ denote the set of all distributions on $(\Omega, \mathcal{F})$ equivalent to $\mathbb{P}^{\theta_0}$. Assume that all these can be achieved by parameters $\theta \in \Theta$, i.e. $\mathcal{M}^e(\mathbb{P}^{\theta_0}) = \{\mathbb{P}^\theta | \theta \in \Theta\}$. For $\mathbb{P}^\theta \in \mathcal{M}^e(\mathbb{P}^{\theta_0})$ let $\mathbb{P}^\theta(\cdot|\mathcal{F}_t)$ denote the distribution conditional on $\mathcal{F}_t$. Due to our assumption to only consider distributions equivalent to $\mathbb{P}^{\theta_0}$, the reference distribution merely fixes the null-sets of the model, i.e. distinct agents at least agree on impossible and sure events. This assumption has no influence on the stochastic structure of the distributions it just tells the decision makers what sure or impossible events are. An economic interpretation of this assumption was given by Epstein and Marinacci in [Epstein & Marinacci, 06]. They related it to an axiom on preferences first postulated by Kreps in [Kreps, 79]. He claimed that if an agent is ambivalent between an act x and $x \cup x'$ then he should also be ambivalent between $x \cup x''$ and $x \cup x' \cup x''$. Meaning if the possibility of choosing x' in addition to x brings no extra utility compared to just being able to choose x, then also no additional utility should arise from being able to choose x' supplementary to $x \cup x''$.

Furthermore we define $X : \Omega \to \mathbb{R}$ to be an $\mathcal{F}$-measurable random variable which can be interpreted as a payoff at final time T. Assume X being essentially bounded with $\operatorname{ess\,sup} |X| = \kappa > 0$. Having constructed the filtered reference space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P}^{\theta_0})$ as above, the sets of almost surely bounded $\mathcal{F}$-measurable and $\mathcal{F}_t$-measurable random variables are denoted by $L^\infty := L^\infty(\Omega, \mathcal{F}, \mathbb{P}^{\theta_0})$ and $L^\infty_t := L^\infty(\Omega, \mathcal{F}_t, \mathbb{P}^{\theta_0})$, respectively. All equations have to be understood $\mathbb{P}^{\theta_0}$-almost surely.

Remark 4.2.1. *As we will see in course of the article, the parametric setting is only needed in the first part on the constructive approach to learning. All statements in the second part, the generalization of the Blackwell-Dubins theorem, can be posed in terms of an arbitrary underlying filtered space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P}_0)$ with distributions in $\mathcal{M}^e(\mathbb{P}_0)$, where $\mathbb{P}_0$ denotes the reference distribution, i.e. in a non-parametric setting. Moreover, for these results, we do not need the particular structure of Ω in terms of a product of marginal spaces S_t. We however follow the parametric approach throughout to obtain a unified appearance.*

4.3 Dynamic Convex Risk Measures

In this article, we apply the theory of convex risk measures as set out in [Föllmer & Penner, 06] for end-period payoffs. For payoff processes, convex risk measures are described in [Cheridito et al., 06]. We do not consider the axiomatic approach to convex risk but take the robust representation of dynamic convex risk measures or, equivalently, of dynamic variational preferences as given.

Definition 4.3.1 (Dynamic Convex Risk & Penalty Functions)**.** (a) *A family* $(\rho_t)_t$ *of mappings* $\rho_t : L^\infty \to L^\infty_t$ *is called a* dynamic convex risk measure *if each component* ρ_t *is a conditional convex risk measure, i.e. for all* $X \in L^\infty$, ρ_t *can be represented in terms of*

$$\rho_t(X) = \operatorname*{ess\,sup}_{\mathbb{Q} \in \mathcal{M}^e(\mathbb{P}^{\theta_0})} \left(\mathbb{E}^{\mathbb{Q}} \left[-X | \, \mathcal{F}_t \right] - \alpha_t(\mathbb{Q}) \right),$$

where $(\alpha_t)_t$ *denotes the dynamic penalty function, i.e. a family of mappings* $\alpha_t : \mathcal{M}^e(\mathbb{P}^{\theta_0}) \to L^\infty_t$, $\alpha_t(\mathbb{Q}) \in \mathbb{R}_+ \cup \infty$, *closed and grounded. For technical details on the penalty see [Föllmer & Schied, 04].*
(b) *Equivalently, we define the* dynamic concave monetary utility function $(u_t)_t$ *by virtue of* $u_t := -\rho_t$, *i.e.*

$$u_t(X) := \operatorname*{ess\,inf}_{\mathbb{Q} \in \mathcal{M}^e(\mathbb{P}^{\theta_0})} \left(\mathbb{E}^{\mathbb{Q}} \left[X | \, \mathcal{F}_t \right] + \alpha_t(\mathbb{Q}) \right).$$

Remark 4.3.2. (a) *By Theorem* 4.5 *in [Föllmer & Penner, 06], the above robust representation in terms of* $\mathcal{M}^e(\mathbb{P}^{\theta_0})$ *is sufficient to capture all time-consistent dynamic convex risk measures.*
(b) *Assuming risk neutrality but uncertainty aversion, no discounting, and no intermediate payoff,* $(u_t)_t$ *is the robust representation of dynamic variational preferences as introduced in [Maccheroni et al., 06b]. In this sense, all our results also hold equivalently for dynamic variational preferences. However, we have chosen to concentrate on dynamic convex risk measures here.*

Assumption 4.3.3. *In the robust representation, we assume the penalty* α_t *to be given by the* minimal penalty $\alpha_t^{\min}$. *The minimal penalty is introduced in terms of acceptance sets in [Föllmer & Penner, 06], p.64: For every* $\mathbb{Q} \in \mathcal{M}^e(\mathbb{P}^{\theta_0})$

$$\alpha_t^{\min}(\mathbb{Q}) := \operatorname*{ess\,sup}_{X \in L^\infty : \rho_t(X) \leq 0} \mathbb{E}^{\mathbb{Q}} \left[-X | \, \mathcal{F}_t \right].$$

As stated in the respective references, every dynamic convex risk measure $(\rho_t)_t$ can be expressed in terms of the above robust representation, uniquely by virtue of the minimal penalty and vice versa. The notion of minimal penalty is justified by the fact that every other penalty representing the same convex risk measure a.s. dominates the minimal one, cp. [Föllmer & Penner, 06]'s Remark 2.7. Throughout, we assume a representation in terms of the minimal penalty $(\alpha_t^{\min})_t$.

Remark 4.3.4 (Equivalent Notation). *In our parametric set-up, a distribution $\mathbb{P}^\theta$ of the process is uniquely defined by a parameter $\theta \in \Theta$. Hence, we write*

$$\rho_t(X) = \operatorname*{ess\,sup}_{\theta\in\Theta} \left(\mathbb{E}^{\mathbb{P}^\theta}\left[-X|\,\mathcal{F}_t\right] - \alpha_t^{\min}(\theta)\right).$$

Further assumptions on the risk measure under consideration will be posed when necessary.

Remark 4.3.5 (On Coherent Risk). *As set out in the references, the robust representation of coherent risk is a special case of the robust representation of convex risk when the penalty is trivial, i.e. for all t it holds*

$$\alpha_t(\theta) = \begin{cases} 0 & \text{if } \mathbb{P}^\theta(\cdot|\mathcal{F}_t) \in \tilde{\mathcal{Q}}(\cdot|\mathcal{F}_t), \\ \infty & \text{else} \end{cases}$$

for $\tilde{\mathcal{Q}}$ the set of prior distributions induced by all θ in some set $\tilde{\Theta} \subset \Theta$. Throughout, $\tilde{\mathcal{Q}}$ is assumed to be convex and weakly compact or, equivalently, $\tilde{\Theta}$ is assumed to be such.

The following definition is a major assumption needed in order to solve tangible economic problems under convex risk.

Definition 4.3.6 (Time-Consistency). *A dynamic convex risk measure $(\rho_t)_t$ is called* time-consistent *if, for all $t, s \in \mathbb{N}$, it holds*

$$\rho_t = \rho_t(-\rho_{t+s})$$

or, equivalently, $u_t = u_t(u_{t+s})$.

Remark 4.3.7. *For the special approach here, [Cheridito et al., 06] show that it suffices to consider $s = 1$ in the above definition.*

Remark 4.3.8. *As inter alia shown in [Föllmer & Penner, 06], Theorem* 4.5, *time-consistency of $(\rho_t)_t$ is equivalent to a condition on the minimal penalty $(\alpha_t^{\min})_t$ called* no-gain condition *in [Maccheroni et al., 06b].*

We now introduce a special class of dynamic convex risk measures that will be used in several examples later on: Dynamic entropic risk measures. Therefore, we first have to introduce:

Definition 4.3.9 (Relative Conditional Entropy)**.** *For $\mathbb{P} \ll \mathbb{Q}$, we define the* relative entropy *of $\mathbb{P}$ with respect to $\mathbb{Q}$ at time $t \geq 0$ as*

$$H_t(\mathbb{P}|\mathbb{Q}) := \mathbb{E}^{\mathbb{P}}\left[\log Z_t\right],$$

where $(Z_t)_t$ by virtue of $Z_t := \frac{d\mathbb{P}}{d\mathbb{Q}}|_{\mathcal{F}_t}$ denotes the density process of $\mathbb{P}$ with respect to $\mathbb{Q}$. Furthermore, we define the conditional relative entropy *of $\mathbb{P}$ with respect to $\mathbb{Q}$ at time $t \geq 0$ as*

$$\hat{H}_t(\mathbb{P}|\mathbb{Q}) := \mathbb{E}^{\mathbb{P}}\left[\log \frac{Z_T}{Z_t}\middle| \mathcal{F}_t\right] = \mathbb{E}^{\mathbb{Q}}\left[\frac{Z_T}{Z_t}\log \frac{Z_T}{Z_t}\middle| \mathcal{F}_t\right]\mathbb{I}_{\{Z_t>0\}}.$$

Definition 4.3.10 (Entropic Risk Measures)**.** *Given reference model $\mathbb{Q} \in \mathcal{M}^e(\mathbb{P}_0)$. Let $\delta > 0$. We say that dynamic convex risk $\rho_t^e(X)$ of a random variable $X \in L^\infty$, is obtained by a* dynamic entropic risk measure *given reference model $\mathbb{Q} \in \mathcal{M}^e(\mathbb{P}^{\theta_0})$ if it is of the form*

$$\rho_t^e(X) = \operatorname*{ess\,sup}_{\mathbb{P}\in\mathcal{M}^e(\mathbb{P}^{\theta_0})} \left(\mathbb{E}^{\mathbb{P}}[-X|\mathcal{F}_t] - \delta\hat{H}_t(\mathbb{P}|\mathbb{Q})\right). \tag{4.1}$$

Equivalently, dynamic multiplier preferences $(u_t^e)_t$ *are defined by virtue of*

$$u_t^e(X) = \operatorname*{ess\,inf}_{\mathbb{P}\in\mathcal{M}^e(\mathbb{P}^{\theta_0})} \left(\mathbb{E}^{\mathbb{P}}[X|\mathcal{F}_t] + \delta\hat{H}_t(\mathbb{P}|\mathbb{Q})\right). \tag{4.2}$$

Remark 4.3.11. *The variational formula for relative entropy implies*

$$\rho_t^e(X) = \delta\log(\mathbb{E}^{\mathbb{Q}}[e^{-\frac{1}{\delta}X}|\mathcal{F}_t]).$$

Intuitively, an entropic risk measure means that the agent in an uncertain setting believes the reference model $\mathbb{Q}$ as most likely and distributions "further away" as more unlikely. Again, we can write $(\rho_t^e)_t$ by virtue of

$$\rho_t^e(X) = \operatorname*{ess\,sup}_{\theta \in \Theta} \left(\mathbb{E}^{\mathbb{P}^\theta}[-X|\mathcal{F}_t] - \delta \hat{H}_t(\theta|\eta) \right),$$

where $\mathbb{P}^\eta$ defines the reference model.

4.4 A Constructive Approach to Learning

In this section, we try to explicitly develop a learning mechanism by virtue of penalty functions that are then used for the robust representation of dynamic convex risk measures. We will encounter, that this is not an eligible approach to model learning as it is still not clear how to explicitly form a penalty. In a later section, we will just take the robust representation as given and pose the question what can be said about learning when distinct properties of the penalty are assumed.

4.4.1 The Intuition of Learning via Penalties

In a first, intuitive approach, we explicitly introduce a learning mechanism to the penalty $(\alpha_t)_t$ in terms of a likelihood function. The fundamental idea is that the penalty might be viewed as a measure for the likelihood of a distribution. In the extreme case of coherent risk, this means

- $\alpha_t(\theta) = \infty$: $\mathbb{P}^\theta$ is not possible,
- $\alpha_t(\theta) = 0$: $\mathbb{P}^\theta$ is among the most likely.

In general, the larger α_t, the less likely the respective distribution. Stated in other terms, $(\alpha_t)_t$ is a measure for uncertainty aversion: given two penalties $(\alpha_t^1)_t$ and $(\alpha_t^2)_t$, the a.s. larger one corresponds to the less uncertainty averse agent. In the entropic case, $\alpha_t(\theta) = H_t(\mathbb{P}^\theta|\mathbb{P}^{\bar{\theta}})$, the conditional relative entropy of $\mathbb{P}^\theta$ with respect to $\mathbb{P}^{\bar{\theta}}$ at time t, the agent considers $\mathbb{P}^{\bar{\theta}}$ most likely as $H_t(\mathbb{P}^{\bar{\theta}}|\mathbb{P}^{\bar{\theta}}) = 0$ and distributions "further away" as more and more unlikely.

In the coherent case characterized by a trivial penalty, learning means to alternate the sets $\tilde{\mathcal{Q}}_t := \{\mathbb{P} \in \tilde{\mathcal{Q}} \mid \mathbb{P}(\cdot|\mathcal{F}_t) \in \tilde{\mathcal{Q}}(\cdot|\mathcal{F}_t)\}$, $t = 0, ..., T$ of conditional priors on which the penalty has value zero: when more information is available and hence, more might be known about the distribution that rules the world, $\tilde{\mathcal{Q}}_t \supset \tilde{\mathcal{Q}_{t+1}}$, i.e. penalty is increasing in t. For some cut off value β, an intuitive approach would be in terms of some likelihood function l:

$$\alpha_t(\theta) = \begin{cases} 0 & \text{if } l(\mathbb{P}^\theta|\theta, \mathcal{F}_t) \geq \beta, \\ \infty & \text{else.} \end{cases}$$

As a direct generalization to convex risk measures, one might consider the log-likelihood $-\log(l(\mathbb{Q}|\theta, \mathcal{F}_t))$ as penalty. It will turn out that this approach is not eligible since a penalty defined in terms of likelihood functions is not feasible. Hence, we come up with a distinct ansatz in which penalty is given by relative conditional entropy. We then achieve a dynamic convex risk measure but run into trouble regarding time-consistency. A model defined as above serves as a measure theoretic fundament of H.P. Minsky's theory of financial instability: A sequence of "good" events causes the penalty to be smaller for distributions that stochastically dominate for the payoff under consideration. Upon observing favorable events, the agent thinks that nature has become kinder. This might help to understand underestimation of risk leading to bubbles and financial instability in times of growth and financial success.

4.4.2 Special Case: Explicit Learning for Coherent Risk

[Epstein & Schneider, 07] introduce learning for coherent risk in terms of likelihood ratio tests. As we will see later, they do not consider the sets of priors $(\mathcal{Q}_t)_t$ as for example in [Riedel, 04] but the process $\mathcal{P}_t(\mathcal{F}_t)$ of one-step ahead conditional beliefs, formally introduced below, as these immediately represent the learning process. Moreover, [Epstein & Schneider, 07] distinguish between information that can be learned and information that cannot: Information that can be learned is incorporated in a the set of priors not being singleton, information that cannot be learned is incorporated in the set of likelihood functions not being singleton.

Formally, let the state space be given by $S^T := \otimes_{t=1}^T S_t$, $S_t = S$, Θ as in the general model. The space of parameters will be slightly modified, i.e. every $\theta \in \Theta$ uniquely characterizes a distribution on S and not on Ω; however, this modification is restricted to the current subsection. Let $\mathcal{Q}_0 \subset \mathcal{M}(\Theta)$ be the set of priors on Θ and $\mathcal{L}$ the set of likelihoods, i.e. every $l \in \mathcal{L}$ satisfies $l(\cdot|\theta) \in \mathcal{M}(S)$ and $l(s_t|\cdot)$ is $\mathcal{F}_t$-measurable for $s_t \in S_t$. Set $s^t = (s_1, \ldots, s_t)$, $s_i \in S_i$. Every $\mu_0 \in \mathcal{Q}_0$ together with a family of likelihoods $(l_1, l_2, \ldots) \in \mathcal{L}^\infty$ induces a prior $\mathbb{P} \in \mathcal{M}^e(\mathbb{P}_0)$ of the payoff process or, equivalently, the process $(p_t)_t$ of one-step-ahead conditionals

$$p_t(\cdot|s^t) = \int_\Theta l(\cdot|\theta) d\mu_t(\theta|s^t) \quad \in \mathcal{M}(S_{t+1}),$$

where μ_t is derived from μ_0 as described below and $\mu_t(\cdot|s^t) \in \mathcal{Q}_t(s^t)$, the set of posterior beliefs on Θ given history s^t. Hence, multiplicity of beliefs is described by

$$\begin{aligned} \mathcal{P}_t(s^t) &= \left\{ p_t(\cdot|s^t) = \int_\Theta l(\cdot|\theta) d\mu_t(\theta) \;\middle|\; \mu_t \in \mathcal{Q}_t^\alpha(s^t), l \in \mathcal{L} \right\} \\ &:= \int_\Theta \mathcal{L}(\cdot|\theta) d\mathcal{Q}_t^\alpha(\theta). \end{aligned}$$

To complete the model, it leaves to show how $(\mu_0; l_1, \ldots)$ induce μ_t or, equivalently, how $\mathcal{Q}_t(s^t)$ is obtained. For $(\mu_0; l_1, \ldots)$, the posteriors are obtained by Bayesian updating:

$$\begin{aligned} & d\mu_t(\cdot, s^t, \mu_0, l^t) \\ &\quad = \frac{l_t(s_t|\cdot)}{\int_\Theta l_t(s_t|\tilde{\theta}) d\mu_{t-1}(\tilde{\theta}, s^{t-1}, \mu_0, l^{t-1})} d\mu_{t-1}(\cdot, s^{t-1}, \mu_0, l^{t-1}). \end{aligned}$$

Then, the posteriors are achieved by virtue of a likelihood ratio test in terms of the unconditional data density:

$$\mathcal{Q}_t^\alpha(s^t) := \left\{ \mu_t(s^t, \mu_0, l^t) \,\middle|\, \mu_0 \in \mathcal{Q}_0, l^t \in \mathcal{L}^t, \int \prod_{j=1}^t l_j(s_j|\theta) d\mu_0(\theta) \right.$$
$$\left. \geq \beta \max_{\bar{\mu}_0 \in \mathcal{Q}_0, \bar{l}^t \in \mathcal{L}^t} \int \prod_{j=1}^t \bar{l}_j(s_j|\theta) d\bar{\mu}_0(\theta) \right\}$$

for some bound $\beta \in \mathbb{R}^+$.

Remark 4.4.1. *Conceptually, there is a huge difference between the approach in [Epstein & Schneider, 07] and [Gilboa & Schmeidler, 89]: In the latter, the term "multiple priors" means multiple distributions of the payoff stream, all being equally likely, in the former, it means multiple distributions of the parameter, i.e. multiple distributions on the distributions of the payoff stream. Hence, [Epstein & Schneider, 07] is a generalization of [Gilboa & Schmeidler, 89] as the latter framework is achieved with $\mathcal{Q}_0 = \{\mu_0\}$ with μ_0 the uniform distribution on some subset of Θ. In that case we have a trivial α and hence a coherent risk measure. Intuitively, a uniform distribution on a subset of Θ corresponds to the agent believing all distributions in that subset being equally likely and the others impossible.*

Nevertheless, fruitful insights from [Epstein & Schneider, 07] can be gained for our approach in particular the incorporation of a likelihood ratio test. We go a step closer to [Gilboa & Schmeidler, 89] and introduce a single distribution on Θ inducing a unique penalty for a dynamic convex risk measure.

4.4.3 A First, Particularly Intuitive Approach: Simplistic Learning

As stated above, multiple prior preferences mean the agent has a uniform distribution on a subset of Θ: She is sure about which parameters are possible and which not, but has no tendency towards their likeliness. In a way, this corresponds to a non-informative weighting or a trivial penalty function α_0. We act on this non-informative approach and assume the following penalty at time zero: Let $\tilde{\Theta} \subset \Theta$. The penalty corresponding to this distribution is given by:

$$\alpha_t(\theta) = \begin{cases} 0 & \text{if } \theta \in \tilde{\Theta}, \\ \infty & \text{else.} \end{cases}$$

Hence, initially the convex risk measure is actually coherent:

$$\rho_0(X) := \operatorname*{ess\,sup}_{\theta\in\Theta} \left\{ \mathbb{E}^{\mathbb{P}^\theta}[-X] - \alpha_0(\theta) \right\} = \operatorname*{ess\,sup}_{\theta\in\tilde{\Theta}} \mathbb{E}^{\mathbb{P}^\theta}[-X].$$

We now come up with a simple learning mechanism directly defining the dynamic penalty function $(\alpha_t)_t$ in terms of likelihoods. At $t = 0$, we have already characterized the penalty. Furthermore, we set

$$\alpha_1(\theta) := -\ln\left(\frac{l(s_1|\theta)}{\sup_{\bar{\theta}} l(s_1|\bar{\theta})}\right) = -\ln\left(\frac{\mathbb{Q}^\theta(s_1)}{\sup_{\bar{\theta}} \mathbb{Q}^{\bar{\theta}}(s_1)}\right),$$

where $s_1 = s^1$ and

$$\alpha_2(\theta) = -\ln\left(\frac{l(s^2|\theta)}{\sup_{\bar{\theta}} l(s^2|\bar{\theta})}\right) = -\ln\left(\frac{\mathbb{Q}^\theta(s_1)\mathbb{Q}^\theta(s_2|\theta, s^1)}{\gamma_2}\right),$$

where $\gamma_2 := \sup_{\theta\in\Theta} \mathbb{Q}^\theta(s_1)\mathbb{Q}^\theta(s_2|\theta, s^1)$.

Definition 4.4.2. *We say that the penalty $(\alpha_t)_t$ in the robust representation of the convex dynamic risk measure $(\rho_t)_t$ is achieved by* simplistic learning, *if it is of the form:*

$$\alpha_t(\theta) := -\ln\left(\frac{\prod_{i=1}^t \mathbb{Q}^\theta(s_i|\theta, s^{i-1})}{\gamma_t}\right),$$

where $\gamma_t := \sup_{\theta\in\Theta} \prod_{i=1}^t \mathbb{Q}^\theta(s_i|\theta, s^{i-1})$.

Remark 4.4.3 (On improperness of simplistic learning)**.** *$(\alpha_t)_t$ achieved by simplistic leaning is* not *a feasible penalty function.*

Proof. A penalty at t should include the conditional distributions from t onwards as seen in the definition. In our likelihood approach α_t only depends on distributions up to time t, i.e. already realized entities of the density process. □

4.4.4 A Second, More Sophisticated Approach: Entropic Learning

We now incorporate the likelihood function in the relative entropy in order to achieve a risk measure based on the well known and elegant entropic risk measures.

Here, we assume $\theta = (\theta_t)_t \in \Theta$; every entity θ_t characterizes a distribution in $\mathcal{M}(S_t)$ possibly dependent on $(\theta_i)_{i<t}$. The family $\theta = (\theta_t)_t$ then defines a prior $\mathbb{P}^\theta \in \mathcal{M}^e(\mathbb{P}^{\theta_0})$. Set $\theta^t := (\theta_1, \dots, \theta_t)$ analogous to s^t.

In the foregoing section, we have seen the major problem to be that our "penalty" was only contingent on the past evolution of the density process. There is however a whole bunch of possibilities to estimate the future by use of past information. A prominent route is by virtue of maximum likelihood estimator.

Definition 4.4.4 (Experience Based Learning)**.** (a) *Given likelihood l. Being at time t, learning is said to be* naive *if the estimator $\hat{\theta}_t$ for θ_t is achieved solely by taking into account maximum likelihood for the observation s_t at time t.*
(b) *Learning is called* intermediate *or* experience based at level m, *if $\hat{\theta}_t$ is the maximum likelihood estimator of the last m observations* $(s_{t-m}, \dots, s_t)$

$$\mathrm{MLE}_{-m} \in \arg\max_{\theta_t \in \Theta} l(s_{t-m}, \dots, s_t | \theta_t, \hat{\theta}^{t-1}, s^{t-m-1}).$$

(c) *Learning is said to be of* maximum likelihood *type, if, at any t, $\hat{\theta}_t$ is the maximum likelihood estimator of the whole history.*

Note that the naive estimator is just the intermediate one at level zero. Furthermore, note that our definition of experience based maximum likelihood. In the next definition, we characterize how learning results in a distribution for the payoff.

Definition 4.4.5 (Learning Distributions)**.** *Being at time t, having obtained $\hat{\theta}_t$ and the foregoing estimators $(\hat{\theta}_i)_{i<t}$, the reference family $\hat{\theta}$ of parameters is achieved by*

$$\hat{\theta}_i = \begin{cases} \hat{\theta}_i & i \leq t, \\ \hat{\theta}_t & i > t. \end{cases}$$

Having seen how agents learn about the best fitting distribution, we now formally introduce entropic learning for which dynamic entropic risk measures in Definition 4.3.10 serve as a vehicle: We choose the best fitting distribution as reference distribution in the conditional relative entropy.

The agent's variational utility incorporating learning is in our setup given by a convex risk measure with an entropic penalty function:

Definition 4.4.6 (Experience Based Entropic Risk)**.** *A penalty $(\hat{\alpha}_t)_t$ is said to be achieved by* experience based entropic learning *if given as*

$$\hat{\alpha}_t(\eta) := \delta \hat{H}_t(\mathbb{P}^\eta | \mathbb{P}^{\hat{\theta}})$$

for $\delta > 0$ and $\hat{\theta} = (\hat{\theta}_t)_t$ achieved as in Definition 4.4.5, $\eta = (\eta_t)_t \in \Theta$. The resulting convex risk measure $(\hat{\rho}_t)_t$ incorporating this very penalty function is then called experience based entropic risk*.*

Remark 4.4.7. *$(\hat{\alpha}_t^\theta)_t$ is well defined as penalty; this is inter alia shown in [Föllmer & Schied, 04]. Due to our construction, the penalty now incorporates conditional distributions of future movements.*

Remark 4.4.8. *When the parameter is also the realization of an entity in the density process, e.g. in a tree (cp. the example below), relative entropy can directly be written as*

$$\hat{\alpha}_t(\theta) = \mathbb{E}^{\mathbb{P}^\theta} \left[\ln \left(\frac{d\mathbb{P}^\theta}{d\mathbb{P}^{\theta_0}} \Big/ \frac{d\mathbb{P}^{\hat{\theta}}}{d\mathbb{P}^{\theta_0}} \right) \middle| \mathcal{F}_t \right].$$

Remark 4.4.9. *Naive entropic learning reflects the tendency of the agent to forget (or ignore) about the distant past and just assume the present to be the* *best estimator of the underlying model. This learning mechanism is then of course particularly adjuvant in explaining a bubble as it is harder to see that the financial system moves away from the fundamentals.*

Despite [Epstein & Schneider, 07] we do not consider multiplicity of likelihoods here. Hence, we do not incorporate information that cannot be learned upon in our model. Though real world applications with several true parameters, e.g. in incomplete financial markets with a multiplicity of equivalent martingale measures, would be modeled in terms of multiple likelihoods. However, our main result in this section on "time-*in*consistency" of experience based entropic risk would not change when extending the model to multiple likelihoods.

Proposition 4.4.10. *The model is well defined, i.e. for every t, $\hat{\rho}_t$ is a conditional convex risk measure.*

Proof. As can easily be seen, the model satisfies the axioms of convex risk measures: $\hat{\rho}_t : L^\infty \to L_t^\infty$ and

- $\hat{\rho}_t$ is monotone, i.e. $\hat{\rho}_t(X) \leq \hat{\rho}_t(Y)$ for $X \geq Y$ a.s.
- $\hat{\rho}_t$ is cash-invariant, i.e. $\hat{\rho}_t(X + m) = \hat{\rho}_t(X) - m$ $\forall m \in L_t$, $X \in L_T$
- $\hat{\rho}_t$ is convex as a function on L_T

□

As inter alia shown [Föllmer & Penner, 06], Proposition 4.4, dynamic entropic risk measures are time-consistent when the reference distribution is not learned but fixed at the beginning. However, now that the reference distribution is also stochastic, we achieve:

Proposition 4.4.11. *Experience based entropic risk is in general not time-consistent.*

Proof. As proof we construct the following counterexample showing an experience based entropic risk measure which is not time-consistent. □

Example 4.4.12 (Entropic Risk in a Tree)**.** *Since our example is mainly for demonstration purposes we restrict ourselves to a simple Cox-Ross-Rubinstein model with 3 time periods. Each time period is independent of those before. One could imagine that in every time period a different coin is thrown and the result of the coin toss determines the realization in the tree, e.g. from heads results up and from tails down. The payoffs of our random variable* X *are limited to the last time-period and are as shown in the figure below. For tractability reasons we also confine ourselves to a single likelihood function* $l(\cdot \mid \theta)$*. For the same reason we will also use the extreme case of naive updating which means our reference distribution will merely depend on the last observed event in our tree. The probability for going up in this tree will always be assumed to lie in the interval* $[a, b]$ *where* $0 < a \leq b < 1$*.*

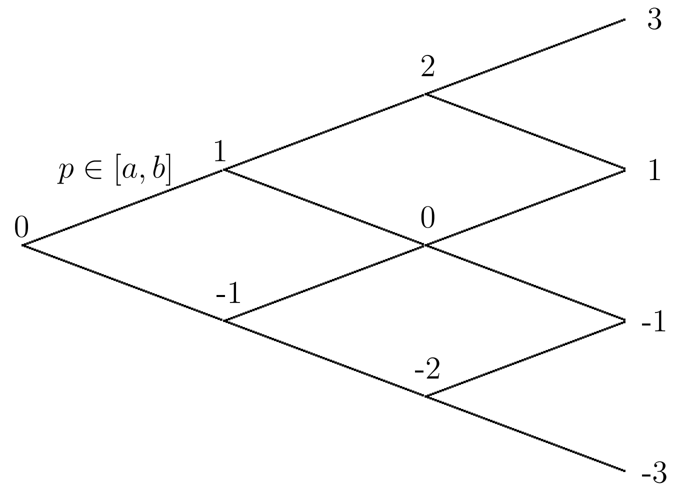

Figure 4.1: Cox-Ross-Rubinstein Model

Time-period 2: *Since we want to show a contradiction to time-consistency we will show that the recursive formula*

$$\hat{\rho}_t(X) = \hat{\rho}_t(-\hat{\rho}_{t+s}(X)) \qquad \textit{for all } t \in [0, T] \textit{ and } s \in \mathbb{N}$$

is violated. So we start with the calculation of $\rho_2(X)$ for the different sets in $\mathcal{F}_2$

$$\begin{aligned}
&\hat{\rho}_2(X)(\text{up},\text{up})\\
&\quad = \operatorname*{ess\,sup}_{p\in[a,b]} \mathbb{E}\left[-X \mid \mathcal{F}_2\right](\text{up},\text{up}) - \mathbb{E}\left[\ln\left(\frac{\theta_2}{\theta_2^*}\right) \mid \mathcal{F}_2\right](\text{up},\text{up})\\
&\quad = \sup_{p\in[a,b]}\left(-3p - 1 + p - p\ln\frac{p}{b} - (1-p)\ln\left(\frac{1-p}{1-b}\right)\right)\\
&\quad = \ln\left(be^{-3} + (1-b)e^{-1}\right),
\end{aligned}$$

where the reference distribution $\mathbb{P}^{\theta^*}$ induced by θ^* is determined by the following maximization:

$$\theta^* = (\theta_0^*, \theta_1^*, \theta_2^*), \quad \theta_2^* \in \arg\max_{\theta_2\in[a,b]} l(\text{up} \mid \theta_2)$$

giving us the maximum-likelihood estimator for what happened in the last time-period which we also think is the right distribution for the next time-period.

The result of this computation can also be obtained by using a variational form which can for example be found in [Föllmer & Penner, 06] and takes the following form

$$\hat{\rho}_t(X) = \ln \mathbb{E}^{\mathbb{P}^{\theta^*}}\left[\exp(-X) \mid \mathcal{F}_t\right],$$

where $\mathbb{P}^{\theta^*}$ is again the reference distribution the decision maker establishes by looking at the past, which, as we look at naive learning, will again only be what happened in the last period. Since this gives way for an easier and quicker computation we will use this form for the following calculations:

$$\begin{aligned}
\hat{\rho}_2(X)(\text{down},\text{up}) &= \ln \mathbb{E}^{\mathbb{P}^{\theta^*}}\left[\exp(-X) \mid \mathcal{F}_2\right](\text{down},\text{up})\\
&= \ln\left(be^{-1} + (1-b)e^{1}\right),
\end{aligned}$$

$$\begin{aligned}
\hat{\rho}_2(X)(\text{up},\text{down}) &= \ln \mathbb{E}^{\mathbb{P}^{\theta^*}}\left[\exp(-X) \mid \mathcal{F}_2\right](\text{up},\text{down})\\
&= \ln\left(ae^{-1} + (1-a)e^{1}\right).
\end{aligned}$$

Here one can nicely observe the extremeness of the naive learning approach. Even though the decision maker in these two calculations is located at the same vertex in the tree he has very different beliefs about the probability of going up or down which causes strong shifts in his risk conception.

In the case of going first down then up he clearly believes up will be more probable in the next step. This is visible in his choice of reference measure $\mathbb{P}^{\theta^*}$ *in the penalty function which he sets* b *for going up and* $1-b$ *for going down.*

In contrast to this the decision maker who has observed up and then down will put more weight on the probability of going down in the next step and therefore sets his reference measure a *for up and* $1-a$ *for down.*

For the last possible event in time 2 our risk-measure takes the following value:

$$\begin{aligned}\hat{\rho}_2(X)(\text{down},\text{down}) &= \ln \mathbb{E}^{\mathbb{P}^{\theta^*}}\left[\exp(-X) \mid \mathcal{F}_2\right](\text{down},\text{down}) \\ &= \ln\left(ae^1 + (1-a)e^3\right).\end{aligned}$$

Time-period 1: *If for the next time-period we maintain the assumption of time-consistency and make use of the recursive formula, using the variational form as we did above will yield*

$$\begin{aligned}\hat{\rho}_1(X)(\text{up}) &= \hat{\rho}_1(-\hat{\rho}_2(X))(\text{up}) = \ln \mathbb{E}^{\mathbb{P}^{\theta^*}}[\exp(\hat{\rho}_2(X)) \mid \mathcal{F}_1](\text{up}) \\ &= \ln\left(b\left(be^{-3} + (1-b)e^{-1}\right) + (1-b)\left(ae^{-1} + (1-a)e^1\right)\right) \\ &= \ln\left(b2e^{-3} + (a+b)(1-b)e^{-1} + (1-a)e^1\right).\end{aligned}$$

Now if we calculate $\hat{\rho}_1(X)(up)$ *without the time-consistency assumption meaning we cannot use the recursive formula we obtain the following equation:*

$$\begin{aligned}\hat{\rho}_1(X)(up) &= \operatorname*{ess\,sup}_{p,q\in[a,b]} \mathbb{E}^{p,q}\left[-X \mid \mathcal{F}_1\right](up) - \mathbb{E}^{p,q}\left[\ln\left(\frac{\theta_1\theta_2}{\theta_1^*\theta_2^*}\right) \mid \mathcal{F}_1\right](up) \\ &= \ln\left(b2e^{-3} + 2b(1-b)e^{-1} + (1-b)2e^1\right).\end{aligned}$$

This clearly is not the same as we obtained under the assumption of time-consistency. However if our dynamic experience based entropic risk measure *were time-consistent these calculations should give us the same results.* Hence this example clearly shows us that the assumption of our risk measure being time-consistent only leads up to contradictions and can therefore not be true.

To emphasize the reason for these inconsistencies set $Z_t := \frac{d\mathbb{P}^{\theta_1}}{d\mathbb{P}^{\theta_2}}\Big|_{\mathcal{F}_t}$, *where* $\mathbb{P}^{\theta_i}$ *is the reference distribution the agent obtains at time i when looking at past realizations and then maximizing the respective likelihood function. Then for instance for* $t = 1$ *and* $\omega = up$ *we obtain:*

$$\begin{aligned}
&\hat{\rho}_1(-\hat{\rho}_2(X - \ln\frac{Z_T}{Z_2}))(up)\\
&= \ln\left[\mathbb{E}^{\mathbb{P}^{\theta_1}}\left[\exp\left(\rho_2\left(X_3 - \ln\frac{Z_3}{Z_2}\right)\right)\right] \mid \mathcal{F}_1\right](up)\\
&= \ln\left[b\mathbb{E}^{\mathbb{P}^{\theta_2}}\left[e^{-X}\frac{Z}{Z_2} \mid \mathcal{F}_2\right](up, up)\right.\\
&\qquad \left.+(1-b)\mathbb{E}^{\mathbb{P}^{\theta_2}}\left[e^{-X}\frac{Z}{Z_2} \mid \mathcal{F}_2\right](up, down)\right]\\
&= \ln\left[b\left(be^{-3}\frac{bbb}{bbb}\frac{bb}{bb} + (1-b)e^{-1}\frac{bb(1-b)}{bb(1-b)}\frac{bb}{bb}\right)\right.\\
&\qquad + (1-b)\left(ae^{-1}\frac{b(1-b)b}{b(1-b)a}\frac{b(1-b)}{b(1-b)}\right.\\
&\qquad\qquad \left.\left.+(1-a)e^{1}\frac{b(1-b)(1-b)}{b(1-b)(1-a)}\frac{b(1-b)}{b(1-b)}\right)\right]\\
&= \ln\left[b2e^{-3} + 2b(1-b)e^{-1} + (1-b)2e^{1}\right] = \rho_1(X)(up),
\end{aligned}$$

which, if $\frac{Z_T}{Z_i} \neq 1$ *(generally true), clearly contradicts time-consistency.*

In this special case for example the measure $\mathbb{P}^{\theta_1}$ *corresponds to the measure assigning the probability b to up in every time period, whereas* $\mathbb{P}^{\theta_2}$ *is the measure assigning b to up in the first 2 time periods and a in the last. That is why e.g.* $Z_3(up, down, up) = \frac{b(1-b)b}{b(1-b)a}$ *and* $\frac{Z_3}{Z_2}(up, down, up) = \frac{b}{a}$.

4.4.5 Lack of Time Consistency

As we have seen in the foregoing paragraph our definition of experience based entropic risk does not result in a time-consistent dynamic convex risk measure. This insight is somewhat disappointing as time consistency is a prosperous vehicle to solve tangible problems. On the other hand, [Schied, 07] shows that a meaningful theory of convex risk can even be achieved in a not generally time-consistent setting.

We have to pose the following question: Does there exist any learning model for the reference distribution such that dynamic entropic risk becomes time-consistent?

Remark 4.4.13. *The major issue that might come into mind is the independence of the reference distribution of future histories. As we will see, this is basically the reason for the general impossibility result below. Furthermore, the worst-case distribution chosen by nature is heavily dependent on the reference distribution. As the latter one may change in a broad variety of manners, there is no good reason to expect nature to choose in a time-consistent way.*

Next, we pose the most general definition of learning in entropic set-ups.

Definition 4.4.14. *A reference distribution $\mathbb{P}^{\tilde{\theta}}$ for experience based entropic risk is said to be obtained by* general learning *if the family $(\tilde{\theta}_t)_t$ is a family of random variables, i.e. not deterministically fixed from scratch. We call the resulting dynamic convex risk measure $(\tilde{\rho}^g_t)_t$ defined by virtue of $\tilde{\alpha}^g_t := \hat{H}_t(\cdot|(\tilde{\theta}_t)_t)$ in the robust representation* general experience based entropic risk.

We see that our definition of experience based entropic risk satisfies the above definition as in that context learning takes place in terms of maximum likelihood.

Using this general definition of learning, we can show an impossibility result for time-consistency of general experience based entropic risk.

Proposition 4.4.15. *General experience based entropic risk $(\tilde{\rho}^g_t)_t$ is in general not time-consistent.*

Proof. Let $\tilde{\theta} = (\tilde{\theta}_1, \ldots)$ be obtained by general learning and ${}^t\tilde{\theta}$ such that $\mathbb{P}^{{}^t\tilde{\theta}} = \mathbb{P}^{\tilde{\theta}}(\cdot|_{\mathcal{F}_t})$. Let $Z_{t+1} := \frac{d\mathbb{Q}^{{}^t\tilde{\theta}}}{d\mathbb{Q}^{{}^{t+1}\tilde{\theta}}}\Big|_{\mathcal{F}_{t+1}}$. Then, we have

$$
\begin{aligned}
\tilde{\rho}^g_t(X) &= \ln \mathbb{E}^{\mathbb{Q}^{{}^t\tilde{\theta}}}\left[e^{-X}\middle|\,\mathcal{F}_t\right] \\
&= \ln \mathbb{E}^{\mathbb{Q}^{{}^t\tilde{\theta}}}\left[e^{\ln \mathbb{E}^{\mathbb{Q}^{{}^t\tilde{\theta}}}\left[e^{-X}\middle|\mathcal{F}_{t+1}\right]}\middle|\,\mathcal{F}_t\right] \\
&= \ln \mathbb{E}^{\mathbb{Q}^{{}^t\tilde{\theta}}}\left[e^{\ln \mathbb{E}^{\mathbb{Q}^{{}^{t+1}\tilde{\theta}}}\left[\frac{Z_T}{Z_{t+1}}e^{-X}\middle|\mathcal{F}_{t+1}\right]}\middle|\,\mathcal{F}_t\right] \\
&= \ln \mathbb{E}^{\mathbb{Q}^{{}^t\tilde{\theta}}}\left[e^{-(-\rho_{t+1}(X-\ln(\frac{Z_T}{Z_{t+1}})))}\middle|\,\mathcal{F}_t\right] \\
&= \tilde{\rho}^g_t(-\tilde{\rho}^g_{t+1}(X-\ln(\frac{Z_T}{Z_{t+1}}))) \\
&\neq \tilde{\rho}^g_t(-\tilde{\rho}^g_{t+1}(X)),
\end{aligned}
$$

if $\frac{Z_T}{Z_{t+1}} \neq 1$ a.s., i.e. if, intuitively speaking, learning actually takes place and, hence, the reference distributions at distinct time periods differ. □

The foregoing result immediately implies our main intuition for experience based entropic risk not being time-consistent though quite puzzling as entropic risk measures are broadly used as standard example for time-consistent convex risk.

Remark 4.4.16 (Main Intuition)**.** *The minimal penalty function uniquely defines a risk measure. Changing the reference distribution due to learning results in a different minimal penalty and hence, a distinct risk measure. Hence, an experience based entropic risk measure is actually a family of dynamic entropic risk measures and our definition of time-consistency is not even applicable.*

4.4.6 A Retrospective – In Between

In this section, we have stated a constructive approach to learning for convex risk measures. We have encountered several problems in doing that:

- In our first intuitive approach, we ran into problems in defining a penalty function not entirely contingent on the past evolution of the density process.
- In our second one, we ran into time-consistency problems.

In a way, in the next section, we put the cart before the horse: We just take the robust representation in terms of minimal penalty of time-consistent dynamic convex risk measures as given and ask ourselves what can be said about "learning" in that respect. We will show an equivalent to the fundamental Blackwell-Dubins Theorem for convex risk measures. As will be seen, this result will be equivalently satisfied whenever the true parameter is eventually learned upon as defined in the subsequent subsection. Our result states some kind of herding behavior as every market participant will eventually perceive risk in the same manner.

4.4.7 Learning for a given Time-Consistent Convex Risk Measure

We now want to encounter, whether we actually have to construct a learning mechanism or if learning is not already incorporated in some sense in the concept of a time-consistent convex risk measure.

Remark 4.4.17. *We have stated that the time-consistency problem encountered so far in learning models is due to the fact that penalties are not just random variables but random itself, i.e. also the functional form depends on the observations. This assumption in general contradicts time-consistency as we actually may achieve distinct risk measures at a particular point in time. However, the basis for learning is already incorporated in convex risk as the domain of penalty consists of Bayesian updated distributions of the process.*

Let us hence assume a true underlying parameter $\theta_0 \in \Theta$ and the agent evaluates risk in terms of robust representation of time-consistent dynamic convex risk $(\rho_t)_t$ with *minimal* penalty $(\alpha_t^{\min})_t$. We then state the following definition:

Definition 4.4.18. *We say that θ_0 is* eventually learned upon *if*

$$\left|\rho_t(X) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X|\mathcal{F}_t]\right| \to 0 \quad \mathbb{P}^{\theta_0} - a.e.$$

for $t \to \infty$.

Proposition 4.4.19. *The above definition is satisfied if and only if*

$$\lim_{t\to\infty}\left|\rho_t(X) - \int_{S_{t+1}} -\rho_{t+1}(X)\mathbb{P}^{\theta_0}(ds_{t+1}|\mathcal{F}_t)\right| = 0 \quad \mathbb{P}^{\theta_0} - a.e.$$

Proof. cp. [Klibanoff et al., 08], Proposition 5. □

In the time-consistent case, the following assertion is equivalent to Definition 4.4.18:

Proposition 4.4.20. *Given a time-consistent dynamic convex risk measure $(\rho_t)_t$, then θ_0 is eventually learned upon if and only if*

$$\alpha_t^{\min}(\theta) \overset{t\to\infty}{\longrightarrow} 0 \quad \mathbb{P}^{\theta_0} - a.e$$

for all θ such that $\alpha_0^{\min}(\theta) < \infty$.

Proof. As $(\rho_t)_t$ is assumed to be time-consistent, it holds for all t

$$\rho_t = \rho_t(-\rho_{t+1})$$

or, more elaborately, for all X

$$\begin{aligned} & \rho_t(X) \\ = \; & \sup_{\theta\in\Theta}\left\{\mathbb{E}^{\mathbb{P}^\theta}\left[-X|\,\mathcal{F}_t\right] - \alpha_t^{\min}(\theta)\right\} \\ = \; & \sup_{\theta\in\Theta}\left\{\mathbb{E}^{\mathbb{P}^\theta}\left[-\rho_{t+1}(X)|\,\mathcal{F}_t\right] - \alpha_t^{\min}(\theta)\right\}. \end{aligned}$$

As further for all X

$$\begin{aligned}
& \int_{S_{t+1}} -\rho_{t+1}(X)\mathbb{P}^{\theta_0}(ds_{t+1}|\mathcal{F}_t) \\
= \; & \mathbb{E}^{\mathbb{P}^{\theta_0}}\left[-\rho_{t+1}(X)|\,\mathcal{F}_t\right] \\
= \; & \sup_{\theta\in\Theta}\left\{\mathbb{E}^{\mathbb{P}^{\theta}}\left[-\rho_{t+1}(X)|\,\mathcal{F}_t\right] - \bar{\alpha}_t^{\min}(\theta)\right\},
\end{aligned}$$

where $(\bar{\alpha}_t^{\min})_t$ is defined as

$$\bar{\alpha}_t^{\min}(\theta) := \begin{cases} 0 & \text{if } \theta = \theta_0 \\ \infty & \text{else,} \end{cases}$$

the proof follows readily: $\alpha_t^{\min}(\theta) \stackrel{t\to\infty}{\longrightarrow} \bar{\alpha}_t^{\min}(\theta)$ by Proposition 4.4.19. Theorem 5.4.(4) in [Föllmer & Penner, 06] then shows equivalence to a vanishing limit given time-consistency. □

In the subsequent sections, we show the notion of being eventually learned upon to be satisfied by convex risk measures in case of time-consistency and under less stringent assumptions in terms of Blackwell & Dubins.

4.5 Adaption of Blackwell-Dubins Theorem

As a cornerstone for our main result on convergence of dynamic convex risk measures, we first generalize the famous Blackwell-Dubins theorem, cp. [Blackwell & Dubins, 62], from conditional probabilities to conditional expectations of risky projects. As set out in the model section, we assume existence of a reference distribution $\mathbb{P}^{\theta_0}$, $\theta_0 \in \Theta$, as in [Blackwell & Dubins, 62]. This reference has to be interpersonally being agreed upon.

Proposition 4.5.1. *Let $\mathbb{P}^{\theta}$ be absolutely continuous with respect to $\mathbb{P}^{\theta_0}$ for* *some $\theta \in \Theta$,*[1] *X as in the definition of the model, then*

$$\left|\mathbb{E}^{\mathbb{P}^{\theta}}[X \mid \mathcal{F}_t] - \mathbb{E}^{\mathbb{P}^{\theta_0}}[X \mid \mathcal{F}_t]\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-almost surely for } t \to \infty.$$

Proof. For improving readability denote $\mathbb{P}^{\theta_0}$ by $\mathbb{P}$ and $\mathbb{P}^{\theta}$ by $\mathbb{Q}$.

Given $\mathbb{P}$ and $\mathbb{Q}$, $\mathbb{Q}$ being assumed absolutely continuous with respect to $\mathbb{P}$, i.e. $\frac{d\mathbb{Q}}{d\mathbb{P}} = q$, and for every n, $\frac{d\mathbb{Q}(\cdot|\mathcal{F}_t)}{d\mathbb{P}(\cdot|\mathcal{F}_t)} = q(\cdot|\mathcal{F}_t)$. Then, the following line of equations holds:

$$\begin{aligned} \mathbb{E}^{\mathbb{Q}}[X|\mathcal{F}_t] &= \mathbb{E}^{\mathbb{Q}(\cdot|\mathcal{F}_t)}[X] \\ &= \mathbb{E}^{\mathbb{P}(\cdot|\mathcal{F}_t)}[q(\cdot|\mathcal{F}_t)X] \end{aligned}$$

and hence

$$\begin{aligned} \left|\mathbb{E}^{\mathbb{Q}}[X|\mathcal{F}_t] - \mathbb{E}^{\mathbb{P}}[X|\mathcal{F}_t]\right| &= \left|\mathbb{E}^{\mathbb{P}(\cdot|\mathcal{F}_t)}\left[(q(\cdot|\mathcal{F}_t) - 1)X\right]\right| \\ &\leq \kappa \left|\mathbb{E}^{\mathbb{P}(\cdot|\mathcal{F}_t)}\left[(q(\cdot|\mathcal{F}_t) - 1)\right]\right| \\ &= \kappa \left|\int (q(\cdot|\mathcal{F}_t) - 1)\, \mathbb{P}(d \cdot |\mathcal{F}_t)\right|, \end{aligned}$$

which converges to zero $\mathbb{P}$-a.s. by Blackwell-Dubins theorem as $(\mathcal{F}_t)_t$ is assumed to be a filtration and, hence, an increasing family of σ-fields. □

Remark 4.5.2. *As we see in the proof, the parametric setting is not needed. The assertion can be shown in the same fashion in a non-parametric setting. The same holds true for subsequent results.*

4.6 Time-Consistent Risk Measures

We will now show a Blackwell-Dubins type result for coherent as well as convex risk measures in case time-consistency is assumed. We see that the risk measure eventually equals the expected value under the true parameter; in this sense, uncertainty vanishes but risk remains.

[1] Note that we have assumed all distributions induced by parameters $\theta \in \Theta$ to be equivalent. In particular, all those are absolutely continuous with respect to each other and this assumption is no restriction within our setup. Also note that the respective θ does not have to be θ_0.

4.6.1 Time-Consistent Coherent Risk

Let $(\rho_t)_t$ be a time-consistent coherent risk measure possessing robust representation

$$\rho_t(X) = \sup_{\theta \in \tilde{\Theta}} \mathbb{E}^{\mathbb{P}^{\theta}}[-X \; |\mathcal{F}_t],$$

with $\tilde{\Theta} \subset \Theta$ assumed to be a convex and compact set of parameters inducing a weakly compact and convex set of priors $\tilde{\mathcal{Q}} \subset \mathcal{M}^e(\mathbb{P}^{\theta_0})$.

Proposition 4.6.1. *For every essentially bounded $\mathcal{F}$-measurable random variable X and time-consistent coherent risk measure $(\rho_t)_t$ we have*

$$\left|\rho_t(X) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X \; |\mathcal{F}_t]\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-almost surely for } t \to \infty.$$

Proof. Thanks to the assumption of time-consistency and compactness there exists a parameter $\theta^* \in \tilde{\Theta}$ such that $\rho_t(X) = \mathbb{E}^{\mathbb{P}^{\theta^*}}[-X \; |\mathcal{F}_t]$ for all $t \in \{0, ..., T\}$ resulting in the following equation

$$\left|\rho_t(X) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X \; |\mathcal{F}_t]\right| = \left|\mathbb{E}^{\mathbb{P}^{\theta^*}}[-X \; |\mathcal{F}_t] - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X \; |\mathcal{F}_t]\right|$$

and this converges to zero as t increases and $\mathbb{P}^{\theta^*} \sim \mathbb{P}^{\theta_0}$ by Proposition 4.5.1. □

Remark 4.6.2. *Note that we have not assumed $\theta_0 \in \tilde{\Theta}$.*

Remark 4.6.3. *The assumption that $\tilde{\Theta}$ is weakly compact is a very crucial assumption, as it assures that the supremum is actually attained. Additionally it is a necessary property for our result to hold, which is shown in the Proposition 4.6.4.*

Proposition 4.6.4. *Weak compactness of the set $\{\mathbb{P}^{\theta} | \theta \in \tilde{\Theta}\}$ of priors is a necessary condition for our result in Proposition 4.6.1 to hold.*

Proof. For the proof, see the counterexample in Section 4.8.2. □

4.6.2 Time-Consistent Convex Risk

Let $(\rho_t)_t$ be a time-consistent dynamic convex risk measure, hence, possessing the following robust representation:

$$\rho_t(X) = \operatorname*{ess\,sup}_{\theta \in \Theta} \left\{ \mathbb{E}^{\mathbb{P}^\theta}[-X|\mathcal{F}_t] - \alpha_t^{\min}(\theta) \right\}$$

with dynamic minimal penalty $(\alpha_t^{\min})_t$.

Assumption 4.6.5. *We assume $(\rho_t)_t$ to be* continuous from below *for all t, i.e. for every sequence of random variables $(X_j)_j$, $X_j \in L^\infty$ for all j, with $X_j \nearrow X \in L^\infty$ we have $\lim_{j\to\infty} \rho_t(X_j) = \rho_t(X)$.*

Remark 4.6.6. *In the coherent case, continuity from below is equivalent to weak compactness of the set $\{\mathbb{P}^\theta | (\alpha_t(\theta))_t = 0\} = \{\mathbb{P}^\theta | \theta \in \tilde{\Theta}\}$ of priors as inter alia shown in [Riedel, 09].*

This assumption has technical advantages as it ensures the supremum to be achieved in the robust representation of ρ_t. A proof is given in Theorem 1.2 of [Föllmer et al., 07]. It is also shown that continuity from below implies continuity from above. To sum up: continuity from above is equivalent to the existence of a robust representation. Continuity from below (which generalizes the compactness assumption in the coherent case) is equivalent to the existence of a robust representation in terms of a distinct prior distribution, the so called worst case distribution.

From an economic point of view, continuity from below results from a feature of preferences already claimed in [Arrow, 71] and related to this assumption by [Chateauneuf et al., 05]. The condition on preferences we need to ask for in order to obtain this feature is called Monotone Continuity: If an act f is preferred over an act g then a consequence x is never that bad that there is no small p such that x with probability p and f with probability $(1-p)$ is still preferred over g. The same is true for good consequences mixed with g.

Formally this means, for acts $f \succ g$, a consequence x and a sequence of events $\{E_n\}_{n\in\mathbb{N}}$ with $E_1 \supseteq E_2 \supseteq ...$ and $\cap_{n\in\mathbb{N}} E_n = \emptyset$ there exists an $\bar{n} \in \mathbb{N}$ such that

$$\begin{bmatrix} x \text{ if } s \in E_{\bar{n}} \\ f(s) \text{ if } s \notin E_{\bar{n}} \end{bmatrix} \succ g \quad \text{and} \quad f \succ \begin{bmatrix} x \text{ if } s \in E_{\bar{n}} \\ g(s) \text{ if } s \notin E_{\bar{n}} \end{bmatrix}.$$

Now with the help of this assumption we can show the Blackwell-Dubins result for time-consistent convex risk measures:

Proposition 4.6.7. *For every essentially bounded $\mathcal{F}$-measurable random variable X and time-consistent convex risk measure $(\rho_t)_t$, continuous from below, it holds*

$$\left|\rho_t(X) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X \mid \mathcal{F}_t]\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-almost surely for } t \to \infty$$

if there exists $\theta \in \Theta$ such that $\alpha_t^{\min}(\theta) \to 0$ $\mathbb{P}^{\theta_0}$-almost surely and $\alpha_0^{\min}(\theta) < \infty$.

Remark 4.6.8 (On the Assumption)**.** *By the main assumption in Proposition 4.6.7 there ought to be some θ such that the penalty vanishes in the long run. This intuitively means that, eventually, nature at least has to pretend some distribution to be the correct one. We see that this is satisfied e.g. in the coherent or in the entropic case.*

The assertion then states that it does not matter which risk measure was chosen as long as the penalty is finite in the beginning. In the time-consistent case, the penalty then vanishes for all those parameters and the convex risk eventually will be coherent.

As we will see later, in the non-time-consistent case, nature has to pay a price for not choosing a distribution time-consistently as in that case penalty has to vanish for the true underlying parameter. To conclude: when nature chooses the worst case distribution time-consistently, she merely has to pretend some distribution to be the underlying one. If she does not choose the worst case measures at any stage time-consistently, she has to reveal the true underlying distribution in the long run.

Remark 4.6.9. *By Theorem* 5.4 *in [Föllmer & Penner, 06] due to time-consistency the assumption* $\alpha_t^{\min}(\theta) \to 0$ $\mathbb{P}^{\theta_0}$*-almost surely for some* $\theta \in \Theta$ *is equivalent to* $\alpha_t^{\min}(\theta) \to 0$ $\mathbb{P}^{\theta_0}$*-almost surely for all* $\theta \in \Theta$ *with* $\alpha_0(\theta) < \infty$*.*

Proof of the proposition. By our assumptions on $(\rho_t)_t$ there exists $\theta^* \in \Theta$ such that the assertion becomes

$$\left|\mathbb{E}^{\mathbb{P}^{\theta^*}}[-X|\mathcal{F}_t] - \alpha_t^{\min}(\theta^*) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X|\mathcal{F}_t]\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-a.s.}$$

By the foregoing proposition on coherent risk, we know that this assertion holds if and only if

$$\left|\alpha_t^{\min}(\theta^*)\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-a.s.}$$

Theorem 5.4 in [Föllmer & Penner, 06] implies this convergence being equivalent to

$$\left|\alpha_t^{\min}(\theta)\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-a.s.}$$

for some $\theta \in \Theta$ such that $\alpha_0(\theta) < \infty$ as assumed to hold in the assertion. □

Corollary 4.6.10. *By Proposition 4.4.20 under the conditions of Proposition 4.6.7,* θ_0 *is eventually learned upon.*

Again, note that we have not assumed θ_0 such that $\alpha_0(\theta_0) < \infty$.

Corollary 4.6.11. *Every dynamic time-consistent convex risk measure* $(\rho_t)_t$ *satisfying the assumptions of the Proposition 4.6.7 is asymptotically precise as in the sense of [Föllmer & Penner, 06], i.e.* $\rho_t(X) \to \rho_\infty(X) = -X$*, and vice versa. In particular, this holds for the coherent case as* $t \to \infty$*.*

Proof. By the assumption of continuity from below, we know that a worst case measure in the robust representation of $(\rho_t)_t$ is actually achieved. By Theorem 5.4 (5) in [Föllmer & Penner, 06] we have that $\rho_t(X) \to \rho_\infty(X) \geq -X$ as we have assumed $\alpha_t^{\min}(\theta_0) \to 0$. Then the assertion is shown by Proposition 5.11 in [Föllmer & Penner, 06]. □

Remark 4.6.12. *In [Föllmer & Penner, 06] time-consistency is directly used to show the existence of the limit $\rho_\infty := lim_{t\to\infty}\rho_t$. As, by assumptions on X in the model, $\lim_{t\to\infty}(E^{\mathbb{P}^{\theta_0}}[-X\ |\mathcal{F}_t])$ exists we achieve existence of ρ_∞ from our result not directly from time-consistency. In our proposition the convergence of the α corresponds to asymptotic precision, however starting at a different point of view. The question now is if time-consistency is a necessary condition for our result to hold. If so, we have gained nothing, if not, we have a more general existence result for ρ_∞ than [Föllmer & Penner, 06]. We will tackle the problem of necessity of time-consistency for our result within the next section.*

Proposition 4.6.13. *$(\rho_t)_t$ being continuous from below is a necessary condition for the result in Theorem 4.6.7 to hold.*

Proof. In Proposition 4.6.4 we show necessity of weak compactness of the set of priors for coherent risk measures. However, weak compactness is equivalent to continuity from below and coherent risk measures are particular examples for convex ones. This proofs the assertion. □

Remark 4.6.14. *In Proposition 4.6.7, if there does not exist θ such that $\alpha_t^{\min}(\theta) \to 0$ but $\alpha_t^{\min}(\theta^*) \le c \in \mathbb{R}_+$ for all $t \ge n_0$ for some $n_0 \in \mathbb{N}$ then there is at least an upper bound on the remaining uncertainty:*

$$|\rho_t(X) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X|\mathcal{F}_t]| \le c$$

as $t \to \infty$.

4.7 Not Necessarily Time-Consistent Risk Measures

We will now achieve a Blackwell-Dubins type result for dynamic coherent and convex risk measures for which we do not pose the time-consistency assumption. However, we still assume the dynamic risk measure to be continuous from below, i.e. in the coherent case the set of priors to be weakly compact. We can still show that anticipation of risk converges to the expected value of a risky project X as defined in the model with respect to the underlying parameter θ_0.

4.7.1 Non Time-Consistent Coherent Risk

We will now restate the result in a manner that time-consistency is not needed. We however need to assume that learning takes place; which is a more liberal assumption than time-consistency as seen in Section 4.8.3.

Definition 4.7.1. (a) *Given a dynamic convex risk measure $(\rho_t)_t$, continuous from below but not necessarily time-consistent, we call a distribution $\mathbb{P}^{\theta_t^*}$* instantaneous worst case distribution at t *if it satisfies*[2]

$$\rho_t(X) = \mathbb{E}^{\mathbb{P}^{\theta_t^*}}[-X|\,\mathcal{F}_t] - \alpha_t^{\min}(\theta_t^*).$$

(b) *We say* learning takes place *if there exists a $\theta \in \Theta$, $\mathbb{P}^{\theta} \sim \mathbb{P}^{\theta_0}$, such that the instantaneous worst case measures $\mathbb{P}^{\theta_t^*} \to \mathbb{P}^{\theta}$ weakly for $t \to \infty$. In the coherent case we need $\theta \in \tilde{\Theta}$ as the penalty is infinite otherwise.*

In this very definition, we see however, that the agent does not have to learn the true underlying parameter θ_0. In this sense, nature might mislead her to a wrong parameter.

We can now relax the time-consistency assumption in the main result of this article. Note that time-consistency is a special case of Definition 4.7.1 given continuity from below as in that case the sequence of instantaneous worst case parameters is constant. Hence, we achieve the more general result:

Proposition 4.7.2. *Let $(\rho_t)_t$ be a not necessarily time-consistent dynamic coherent risk measure for which learning takes place. Then*

$$\left|\rho_t(X) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X\ |\mathcal{F}_t]\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-almost surely for } t \to \infty.$$

[2]Note, that existence is locally guaranteed by continuity from below. As we however have not assumed time-consistency, the instantaneous worst case distributions at each time period may differ, hence global existence is not necessarily fulfilled.

Proof. To make things clearer we will write the proof in terms of penalty functions and not in terms of priors. We know that a coherent risk measure has a robust representation of a convex risk measure with a penalty

$$\alpha_t^{\min}(\theta) = \begin{cases} 0 & \text{if } \mathbb{P}^\theta(\cdot|\mathcal{F}_t) \in \tilde{\mathcal{Q}}(\cdot|\mathcal{F}_t), \\ \infty & \text{else} \end{cases}$$

where $\tilde{\mathcal{Q}}$ is the set of priors, i.e. $\tilde{\mathcal{Q}} = \{\mathbb{P}^\theta|(\alpha_t^{\min}(\theta))_t = 0\}$ uniquely defining the coherent risk measure. As we are in the case of a coherent risk measure, we particularly have $\alpha_t^{\min}(\theta_t^*) = 0$.

First, note that in case $\alpha_t^{\min}(\theta) \to \infty$ for all $\theta \in \tilde{\Theta}$[3], our convergence result cannot hold, as $\lim_{t\to\infty} \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X|\mathcal{F}_t]$ exists and is finite by assumption.

Secondly, in the time-consistent (coherent as well as convex) case, it suffices to assume $\alpha_t^{\min}(\bar{\theta}) \to 0$ for some $\bar{\theta} \in \Theta$. This assumption in the time-consistent case is equivalent to $\alpha_t^{\min}(\theta) \to 0$ for all θ for which $\alpha_0^{\min}(\theta) < \infty$ by Theorem 5.4 in [Föllmer & Penner, 06].

Let us now turn to the proof itself: As $\tilde{\mathcal{Q}}$ is assumed to be weakly compact and non-empty, i.e. there exists a distribution that has penalty zero, we achieve an instantaneous worst case distribution at each time step, i.e. at any t, there exists $\theta_t^* \in \Theta$ s.t.

$$\rho_t(X) = \mathbb{E}^{\mathbb{P}^{\theta_t^*}}[-X|\,\mathcal{F}_t] - \alpha_t^{\min}(\theta_t^*) = \mathbb{E}^{\mathbb{P}^{\theta_t^*}}[-X|\mathcal{F}_t].$$

Of course, due to "time-*in*consistency", we might have $\theta_i^* \neq \theta_j^*$ for $i \neq j$.

The proof is completed by showing the following convergence[4]

$$\mathbb{E}^{\mathbb{P}^{\theta_n^*}}[-X|\mathcal{F}_t] \to \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X|\mathcal{F}_\infty] \qquad \text{for } n, t \to \infty.$$

[3]Of course, convergence is trivial in this case due to triviality of the penalty function.

[4]By our assumptions we know:

- $\mathbb{E}^{\mathbb{P}^{\theta_n^*}}[-X|\mathcal{F}_t] \to \mathbb{E}^{\mathbb{P}^{\theta}}[-X|\mathcal{F}_t]$ for $n \to \infty$ as $\theta_n^* \to \theta$ by Portemonteau's Theorem.
- $\mathbb{E}^{\mathbb{P}^{\theta_n^*}}[-X|\mathcal{F}_t] \to \mathbb{E}^{\mathbb{P}^{\theta_n^*}}[-X|\mathcal{F}_\infty]$ for $t \to \infty$ by Proposition 4.5.1.

The question now is, whether the result also holds when letting $n, t \to \infty$ at once.

In the time-consistent case, where $\theta_i^* = \theta_j^*$ for all i, j, this is immediate by Proposition 4.5.1.

In order to do this we look at the following equation for $n \geq t$ which uses the projectivity of the density, i.e. of the Radon-Nikodym derivative:

$$\mathbb{E}^{\mathbb{P}^{\theta_n^*}}[-X|\mathcal{F}_t] = \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X \frac{d\mathbb{P}^{\theta_n^*}}{d\mathbb{P}^{\theta_0}}\Big|_{\mathcal{F}_n} |\mathcal{F}_t].$$

Define the following sequence of random variables $Y_n := -X \frac{d\mathbb{P}^{\theta_n^*}}{d\mathbb{P}^{\theta_0}}\Big|_{\mathcal{F}_n}$. These have finite expectation and thanks to our assumption that learning takes place and the original Blackwell-Dubins result we have

$$\mathbb{P}^{\theta_0}[\lim_{n\to\infty} Y_n = -X] = \mathbb{P}^{\theta_0}[-X \frac{d\mathbb{P}^{\theta_\infty^*}}{d\mathbb{P}^{\theta_0}}\Big|_{\mathcal{F}_\infty} = -X] = 1.$$

Then, by Lemma 4.7.4, the assertion follows. □

Remark 4.7.3. *Again, note that we have not assumed $\theta_0 \in \tilde{\Theta}$.*

In the foregoing proof, we need a general martingale convergence result as stated in [Blackwell & Dubins, 62], Theorem 2. We know from Doob's famous martingale convergence result that

$$\mathbb{E}^{\mathbb{P}^\theta}[X|\mathcal{F}_t] = \lim_{t\to\infty} \mathbb{E}^{\mathbb{P}^\theta}[X|\mathcal{F}_\infty] \quad a.s.$$

under suitable assumptions. The question is: If $X_n \nearrow_n X$ in some sense, is it true that

$$\mathbb{E}^{\mathbb{P}^\theta}[X|\mathcal{F}_\infty] = \lim_{n,t\to\infty} \mathbb{E}^{\mathbb{P}^\theta}[X_n|\mathcal{F}_t] \quad a.s.?$$

A positive answer is given in the following lemma.

Lemma 4.7.4. *Fix θ. Let $(Y_n)_n$ be a sequence of $\mathcal{F}$-measurable random variables such that $\mathbb{E}^{\mathbb{P}^\theta}[\sup_n |Y_n|] < \infty$. Assume $Y_n \to_{n\to\infty} Y$ almost surely for some $\mathcal{F}$-measurable random variable Y. Then, it holds*[5]

$$\lim_{n,t\to\infty} \mathbb{E}^{\mathbb{P}^\theta}\left[Y_n|\,\mathcal{F}_t\right] = \mathbb{E}^{\mathbb{P}^\theta}\left[Y|\,\mathcal{F}\right].$$

[5]The convergence in the assertion of the lemma can also be shown in L^1.

Proof. We re-sample the proof in [Blackwell & Dubins, 62]: For $k \in \mathbb{N}$, set $G_k := \sup\{Y_n | n \geq k\}$. If $n \geq k$, we hence have $Y_n \leq G_k$ and thus

$$\mathbb{E}^{\mathbb{P}^\theta}\left[Y_n | \mathcal{F}_t\right] \leq \mathbb{E}^{\mathbb{P}^\theta}\left[G_k | \mathcal{F}_t\right] \tag{4.3}$$

for all t. Together with Doob's martingale convergence result and Lebesgue's theorem, we achieve

$$\begin{aligned} z &:= \lim_{j\to\infty} \sup_{n,t\geq j} \mathbb{E}^{\mathbb{P}^\theta}\left[Y_n | \mathcal{F}_t\right] \\ &\overset{(4.3)}{\leq} \lim_{j\to\infty} \sup_{t\geq j} \mathbb{E}^{\mathbb{P}^\theta}\left[G_k | \mathcal{F}_t\right] \\ &= \lim_{t\to\infty} \mathbb{E}^{\mathbb{P}^\theta}\left[G_k | \mathcal{F}_t\right] \\ &\overset{\text{Doob}}{=} \mathbb{E}^{\mathbb{P}^\theta}\left[G_k | \mathcal{F}\right] \end{aligned}$$

and

$$z \leq \lim_{k\to\infty} \mathbb{E}^{\mathbb{P}^\theta}\left[G_k \mid \mathcal{F}\right] \overset{\text{Lebesgue}}{=} \mathbb{E}^{\mathbb{P}^\theta}\left[Y \mid \mathcal{F}\right].$$

In the same token,

$$x := \lim_{j\to\infty} \inf_{t,n\geq j} \mathbb{E}^{\mathbb{P}^\theta}\left[Y_n | \mathcal{F}_t\right] \geq \mathbb{E}^{\mathbb{P}^\theta}\left[Y | \mathcal{F}\right],$$

which completes the proof since

$$x = \lim_{j\to\infty} \inf_{t,n\geq j} \mathbb{E}^{\mathbb{P}^\theta}\left[Y_n | \mathcal{F}_t\right] \leq \lim_{j\to\infty} \sup_{n,t\geq j} \mathbb{E}^{\mathbb{P}^\theta}\left[Y_n | \mathcal{F}_t\right] = z.$$

□

Remark 4.7.5 (On Blackwell-Dubins Type Learning)**.** *Blackwell-Dubins applies for learning models but does not necessarily result in time-consistency as this notion is now motivated as a special case of our notion of θ_0 to be eventually learned upon.*

We have built a bridge between the first and the second part of this article: in the first part we have achieved dynamic convex risk measures by virtue of learning that did not turn out to be time-consistent. Hence, we have shown, that our result even holds for those models, e.g. entropic learning.

Remark 4.7.6. *Note, that the above new version of the fundamental result particularly holds for time-consistent dynamic coherent risk measures as then such a limiting θ as in the Definition 4.7.1(b) always exists, the worst case one. However, we particularly have an existence result for the limit $\rho_\infty := \lim_{t\to\infty} \rho_t$ in the non time-consistent case and thus a more general existence result than in [Föllmer & Penner, 06].*

4.7.2 Non Time-Consistent Convex Risk

As in the case of coherent risk measures, we now state our generalization of the Blackwell-Dubins theorem when the dynamic convex risk measure is *not* assumed to be time-consistent. As in the coherent case, we assume that learning takes place, i.e. there exists $\theta \in \Theta$ such that the instantaneous worst case $\theta_t^* \to \theta$ as $t \to \infty$. Furthermore, we have to assume $\alpha_t^{\min}(\theta_t^*) \to 0$ as $n \to \infty$:[6] As in the foregoing proof, we achieve convergence of the conditional expectations under the family of instantaneous worst case distributions to the conditional expectation under θ_0.

Proposition 4.7.7. *For every risky project X as set out in the model and dynamic convex risk measure $(\rho_t)_t$, continuous from below but not necessarily time-consistent, we have*

$$\left|\rho_t(X) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-X \ |\mathcal{F}_t]\right| \to 0 \quad \mathbb{P}^{\theta_0}\text{-almost surely for } t \to \infty$$

if learning takes place for an instantaneous worst case sequence $(\theta_t^*)_t$ toward some $\theta \in \Theta$ and we have

$$\alpha_t^{\min}(\theta_t^*) \to 0.$$

Proof. Applying the procedure used in the proof of Proposition 4.7.2 to the proof of Proposition 4.6.7 shows the assertion. $\square$

[6]Note, again, we do not have to assume $\alpha_t^{\min}(\theta_0) \to 0$.

4.8 Examples

In this section, we first consider dynamic entropic risk measures as a prominent economic example of time-consistent dynamic convex risk measures. In the second part we state a counterexample serving as proof for Proposition 4.6.4 and 4.6.13. As a last point, we consider a dynamic risk measure that is not time-consistent.

4.8.1 Entropic Risk

Here, we will have a look at time-consistent dynamic entropic risk measure $(\rho_t^{\mathrm{e}})_t$. Recall its Definition 4.3.10 in terms of

$$\rho_t^{\mathrm{e}}(X) := \delta \log \mathbb{E}\left[e^{-\gamma X} \middle| \mathcal{F}_t \right]$$

for some model parameter $\delta > 0$. A fundamental result shows that the robust representation of dynamic entropic risk is given in terms of conditional relative entropy as penalty function, i.e. for all n, we have

$$\alpha_t^{\min}(\theta) = \frac{1}{\gamma} \hat{H}_t(\mathbb{P}^\theta | \mathbb{P}^\eta) := \frac{1}{\gamma} \mathbb{E}^{\mathbb{P}^\theta} \left[\ln \frac{Z_T}{Z_t} \middle| \mathcal{F}_t \right],$$

where $Z_t := \frac{d\mathbb{P}^\theta}{d\mathbb{P}^\eta}\Big|_{\mathcal{F}_t}$, the Radon-Nikodym derivative of $\mathbb{P}^\theta$ with respect to $\mathbb{P}^\eta$ conditional on $\mathcal{F}_t$.

The fundamental Blackwell-Dubins Theorem immediately shows that

$$\left| \mathbb{P}^\theta(\cdot | \mathcal{F}_t) - \mathbb{P}^\eta(\cdot | \mathcal{F}_t) \right| \to 0$$

for every θ, η. Hence, we have that $\frac{Z_T}{Z_t} \to 1$ $\mathbb{P}^{\theta_0}$-a.s. for $t \to \infty$ and hence

$$\alpha_t^{\min}(\theta) \to 0$$

showing Proposition 4.6.7 to hold. This is an alternative way to show the last assertion in Theorem 6.3 in [Föllmer & Penner, 06] directly.

4.8.2 Counterexample

To show necessity of continuity from below in Proposition 4.6.7 we consider the following example introduced in [Föllmer & Penner, 06]:

The underlying probability space consists of the state space $\Omega = (0,1]$ endowed with the Lebesgue measure $\mathbb{P}^{\theta_0}$ and a filtration $(\mathcal{F}_t)_t$ generated by the dyadic partitions of Ω. This means $\mathcal{F}_t$ is generated by the sets $J_{t,k} := (k2^{-t}, (k+1)2^{-t}]$ for $k = 0, ..., 2^{t-1}$. In this setting [Föllmer & Penner, 06] construct a time-consistent coherent and therefore convex risk measures with $\alpha_t^{\min}(\theta_0) \to 0$ $\mathbb{P}^{\theta_0}$-a.s. of the following form:

$$\rho_t(X) = -\operatorname{ess\,sup}\{m \in L_t^\infty \ | m \leq X\}.$$

That this sequence from all properties assumed in Proposition 4.6.7 is only missing continuity from below (here equivalent to weak compactness of priors) can be seen in the following way: Let t be arbitrary but fixed and X defined by virtue of

$$X(\omega) = \begin{cases} 0 & \text{for } \omega \in (0, (2^t - 1)2^{-t}], \\ 1 & \text{else.} \end{cases}$$

Then we can construct a sequence $(X_n)_n$, $X_n \nearrow X$, such that $\rho_t(X_n) = 0$ for all n but $\rho_t(X) = -X \neq 0$. This shows $(\rho_t)_t$ not being continuous from below.

Now we still have to show that for this construction the statement of our proposition is not fulfilled. To verify this look at a set A assumed to be $\mathcal{F} := \sigma(\bigcup_{t \geq 0} \mathcal{F}_t)$-measurable such that $\mathbb{P}^{\theta_0}[A] > 0$ and $\mathbb{P}^{\theta_0}[A^c \cap J_{t,k}] \neq 0$ for all t and k. For this set, it holds

$$\lim_{t \to \infty} \left| \rho_t(\mathbb{1}_A) - \mathbb{E}^{\mathbb{P}^{\theta_0}}[-\mathbb{1}_A \ |\mathcal{F}_t] \right| = \lim_{t \to \infty} \left| 0 + \mathbb{P}^{\theta_0}[A \ |\mathcal{F}_t] \right| = \mathbb{P}^{\theta_0}[A] > 0$$

and hence necessity of the continuity assumption is shown.

The skeptical reader might now object that such a set A might not exist. For sake of completeness we briefly quote a set A from [Föllmer & Penner, 06] that satisfies our assumptions: Let A be defined by virtue of its complement

$$A := \left(\bigcup_{t=1}^{\infty} \bigcup_{k=1}^{2^t-1} U_{\epsilon_t}(k2^{-t}) \right)^c ,$$

where U_{ϵ_t} denotes the ϵ_t-neighborhood and $\epsilon_t \in]0, 2^{-2t}]$.

4.8.3 A Non Time-Consistent Example

Here, we consider the entropic learning model introduced in Definition 4.4.6 explicitly in terms of $\Omega = \otimes_t S_t$. Let $\mathbb{P}^\theta$ denote the distribution induced by $\theta = (\theta_t)_t$, θ_t inducing a marginal distribution in $\mathcal{M}(S_t)$. Though the model looks quite similar to dynamic entropic risk measures, we briefly recall it: Let the robust representation of a dynamic convex risk measure $(\hat{\rho}_t)_t$ be given by virtue of the penalty

$$\hat{\alpha}_t^{\min}(\theta) := \delta \hat{H}_t(\mathbb{P}^\theta | \mathbb{P}^{\hat{\theta}}),$$

$\delta > 0$ and $\hat{\theta} = (\hat{\theta}_t)_t$ be achieved as in Definition 4.4.5: for $t \in \mathbb{N}$, $\hat{\theta}_t$ is the maximum likelihood estimator of the foregoing observations and $\hat{\theta}_i := \hat{\theta}_t$ for $i > t$. Restricting ourselves to the iid case, we know that we achieve $\hat{\theta}_t \to \bar{\theta}_0$, $\mathbb{P}^{\theta_0}$-a.s., where $\theta_0 = (\bar{\theta}_0)_t$ for some $\bar{\theta}_0$ inducing a marginal distribution in $\mathcal{M}(S_t)$. By definition, $(\hat{\rho}_t)_t$ is a dynamic convex risk measure. As shown in Proposition 4.4.15, $(\hat{\rho}_t)_t$ is not time-consistent. By standard results on conditional entropic risk measures, $(\hat{\rho}_t)_t$ is continuous from below.

Furthermore, Proposition 4.7.7 is applicable and hence, our generalization of Blackwell-Dubins' theorem holds for experience based entropic risk. Indeed: By definition of the penalty and our considerations in Section 4.8.1, $\hat{\alpha}_t^{\min}(\theta) \to 0$ as $t \to \infty$ for all $\theta \in \Theta$. Secondly, as the maximum likelihood estimator is asymptotically stable, i.e. $\hat{\theta}_t \to \bar{\theta}_0$, the conditional reference distributions $\mathbb{P}^{\hat{\theta}}(\cdot|\mathcal{F}_t)$ converge. Thus, the worst case instantaneous distributions $\mathbb{P}^{\theta_t^*}$ converge as in Definition 4.7.1 due to continuity of the entropy and as the effective domain of the penalty is given by conditional distributions, a fact that is made particularly precise in [Maccheroni et al., 06b].[7]

4.9 Conclusions

The major contribution of our results is to carry over the famous Blackwell-Dubins theorem from probability distributions to convex risk measures. It is particularly striking that the results still hold when time-consistency is not posed as an assumption.

Hereto, the present article is twofold: In the first part, we show that explicitly constructing dynamic convex risk measures by virtue of a penalty emerging from a learning mechanism and inserted in the robust representation of convex risk measures leads to time-consistency problems. In the second part, we have then assumed a time-consistent dynamic convex risk measure for granted and asked the question of limit behavior; more elaborately its convergence to the expected value under the true underlying distribution.

We therefore introduced a generalization of the famous Blackwell-Dubins theorem on "Merging of Opinions" to conditional expected values. Existence of a worst case distribution due to continuity from below and time-consistency then allowed for a further generalization to coherent and convex risk measures. In particular, we have obtained the existence of the limiting risk measure ρ_∞ in that case.

By virtue of a counterexample, we have shown necessity of continuity from below for our result. However, we have shown that time-consistency is not necessary for the result to hold. In particular, we have obtained a more general existence result for the limiting risk measure ρ_∞ than in [Föllmer & Penner, 06]. Our generalization of the Blackwell-Dubins theorem

[7] The notation is quite misleading at this point: the worst case instantaneous distributions $\mathbb{P}^{\theta_t^*} \in \mathcal{M}^e(\mathbb{P}^{\theta_0})$ as in Definition 4.7.1 is a distribution on $(\Omega, \mathcal{F})$ as θ_t^* is an element of Θ and not a "marginal" parameter as the above θ_ts.

was shown to be equivalent to the notion of the parameter being eventually learned upon and the notion of asymptotic precision in [Föllmer & Penner, 06] in the time-consistent case.

Further research should be conducted in the direction of our results. First, of course, the riddle of explicitly constructing convex risk measures by virtue of the penalty function is still to solve; in particular, how a learning mechanism might be introduced without destroying the assumption of time-consistency. Weaker notions of time-consistency that are satisfied in a "learning" environment should be introduced along with a comprehensive theory allowing for solutions of tangible economic and social problems.

In the article at hand, we have considered risky projects with final payoffs, i.e. random variables of the form $X \in \mathcal{F}$. We have shown convergence of convex risk measures to the conditional expected value with respect to the true underlying distribution: a generalization of the Blackwell-Dubins theorem to (not necessarily time-consistent) convex risk measures for final payoffs. To us it seems being an interesting, yet challenging, task to generalize our result to the case of convex risk measures for stochastic payoff processes $(X_t)_t$ with respect to some filtration $(\mathcal{F}_t)_t$, where each X_t denotes the stochastic payoff in period t. [Cheridito et al., 06] introduce dynamic convex risk measures for these stochastic processes and elaborately discuss time-consistency issues but do not inspect limiting behavior. A major difficulty in the case of stochastic processes is that the assumption of equivalent distributions should be replaced by local equivalence, cp. [Riedel, 09]. Hence, the main question turns out to be if the result still holds assuming local instead of global equivalence as done here.

Chapter 5

Closing Remarks

Within the three chapters of this thesis we have studied several problems arising in the context of dynamic decision problems under Knightian Uncertainty. Each chapter discusses its respective topic in detail and ends with a conclusion summarizing its results. Nevertheless for completeness we will briefly restate our achievements at this point:

First, we presented an alternative characterization for time-consistent sets of measures on finite trees. It allows to express our set of measures through a set of predictable processes which in return again defines a time-consistent set of measures. This representation is unique up to the choice of a martingale basis. Trying to generalize our assumptions in standard ways in order to achieve a more universal representation failed, showing the scope of this characterization.

In the third chapter we studied if and under what conditions a duality theorem for optimal stopping problems holds in the multiple priors framework of [Epstein & Schneider, 03]. The result is a minimax theorem for rather general assumptions on the payoff process and standard assumptions on the set of measures. We use this theorem to identify the worst case measure in the setting of κ-ambiguity adapted to our framework and apply it to multiple prior super- and submartingales determining the optimal stopping times.

Finally, we have considered dynamic convex risk measures when information is gathered in course of time. We have generalized the fundamental Blackwell-Dubins theorem from [Blackwell & Dubins, 62] to not necessarily time-consistent dynamic convex risk measures and have thus shown their convergence to conditional expected values with respect to the true underlying distribution: Intuitively the result shows that uncertainty vanishes but risk endures.

Bibliography

[Anscombe & Aumann, 63] ANSCOMBE, F.J. & AUMANN, R.J. (1963): *A Definition of Subjective Probability* , Annals of Mathematical Statistics, Vol. 34, pp. 199-205.

[Arrow, 71] ARROW, K. (1971): *Essays in the Theory of Risk Bearing*, Markham Publishing Co., Chicago.

[Artzner et al., 99] ARTZNER, P.; DELBAEN, F.; EBER, J.-M. & HEATH, D. (1999): *Coherent Measures of Risk*, Mathematical Finance, Vol.9, No.3, pp.203-228.

[Artzner et al., 07] ARTZNER, P.; DELBAEN, F.; EBER, J.-M.; HEATH, D. & KU, H. (2007): *Coherent Multiperiod Risk Adjusted Values and Bellman's Principle*, Annals of Operations Research, Vol.152, No1, pp.5-22.

[Blackwell & Dubins, 62] BLACKWELL, D. & DUBINS, L. (1962): *Merging of Opinions with Increasing Information*, Ann. Math. Statistics, Vol. 33, No. 3, pp. 882-886.

[Chateauneuf et al., 05] CHATEAUNEUF, A.; MACCHERONI, F.; MARINACCI, M. & TALLON, J.-M. (2005): *Monotone Continuous Multiple Priors* , Economic Theory, Vol. 26, No.4, pp. 973-982.

[Chen & Epstein, 02] CHEN, Z. & EPSTEIN, L. (2002): *Ambiguity, Risk, and Asset Returns in Continuous Time* , Econometrica Vol. 70, No. 4, pp. 1403-1443

[Cheridito et al., 06] CHERIDITO, P.; DELBAEN, F. & KUPPER, M. (2006): *Dynamic Monetary Risk Measures for Bounded Discrete-Time Processes*, Electronic Journal of Probability, Vol.11, No.3, pp.57-106.

[Cheridito & Stadje, 07] CHERIDITO, P. & STADJE, M. (2009): *Time-inconsistency of VaR and time-consistent alternatives*, Finance Research Letters, Vol.6, No.1, pp.40-46.

[Chudjakow & Vorbrink, 09] CHUDJAKOW, T. & VORBRINK, J. (2009): *Exercise Strategies for American Exotic Options under Ambiguity*, IMW-Working paper

[Czichowsky & Schweizer, 09] CZICHOWSKY, C. & SCHWEIZER, M. (2009): *Closedness in the Semimartingale Topology for Spaces of Stochastic Integrals with Constrained Integrands*, FINRISK Working paper.

[Delbaen, 03] DELBAEN, F. (2003):*The Structure of m-Stable Sets and in Particular of the Set of Risk Neutral Measures*, In Yor, M., Emery, M. (eds.): In Memoriam Paul-André Meyer - Séminaire de Probabilités XXXIX. Berlin Heidelberg New York: Springer, pp. 215-258.

[Dothan, 90] DOTHAN, M. (1990): *Prices in Financial Markets*, 1. edition, Oxford University Press.

[El Karoui, 81] EL KAROUI, N. (1981): *Les Aspectes Probabilistes du Contrôle Stochastique*, Lecture Notes in Mathematics, Vol. 876, pp. 73-238. Springer, Berlin.

[Ellsberg, 61] ELLSBERG, D. (1961): *Risk, Ambiguity, and the Savage Axioms* , Quarterly Journal of Economics, Vol. 75, pp. 643-669.

[Engelage, 09] ENGELAGE, D. (2009): *Optimal Stopping under Convex Risk*, in *Essays on Coherent and Convex Risk Measures*, Ph.D. Thesis, Bonn Graduate School of Economics.

[Epstein & Marinacci, 06] EPSTEIN, L. & MARINACCI, M. (2006): *Mutual Absolute Continuity of Multiple Priors* , Journal of Economic Theory, Vol.137, No.1, pp.716-720.

[Epstein & Schneider, 03] EPSTEIN, L. & SCHNEIDER, M. (2003): *Recursive Multiple Priors*, Journal of Economic Theory, Vol.113, pp.1-13.

[Epstein & Schneider, 07] EPSTEIN, L. & SCHNEIDER, M. (2007): *Learning under Ambiguity*, Working Paper, University of Rochester.

[Föllmer & Schied, 04] FÖLLMER, H. & SCHIED, A. (2004): *Stochastic Finance, An Introduction in Discrete Time*, 2nd edition, Walter DeGruyter.

[Föllmer & Penner, 06] FÖLLMER, H. & PENNER, I. (2006): *Convex Risk Measures and the Dynamics of their Penalty Functions*, Statistics and Decisions, Vol.24, pp.61-96.

[Föllmer et al., 07] FÖLLMER, H.; SCHIED, A. & WEBER, S. (2007): *Robust Preferences and Robust Portfolio Choice*, ORIE, Cornell University.

[Gilboa & Schmeidler, 89] GILBOA, I. & SCHMEIDLER, D. (1989): *Maximin Expected Utility with non-unique Prior*, Journal of Mathematical Economics, Vol.18, pp.141-153.

[Hansen & Sargent, 01] HANSEN, L. & SARGENT, T. (2001): *Robust control and Model Uncertainty*, American Economic Review Papers and Proceedings, Vol. 91, pp. 60-66.

[Karatzas & Kou, 98] KARATZAS, I. & KOU, S. (1998): *Hedging American Contingent Claims with Constrained Portfolios*, Finance and Stochastics, Vol. 2, pp. 215-258.

[Karatzas & Shreve, 98] KARATZAS, I. & SHREVE, S. (1998): *Methods of Mathematical Finance*, New York: Springer

[Klibanoff et al., 08] KLIBANOFF, P.; MARINACCI, M. & MUKERJI,S. (2008): *Recursive Smooth Ambiguity Preferences*, Journal of Economic Theory, forthcoming.

[Knight, 21] KNIGHT, F. H. (1921): *Risk, Uncertainty and Profit*, Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company, 1921.

[Kreps, 79] KREPS, D. (1979): *A Representation Theorem for 'Preference for Flexibility'* , Econometrica, Vol. 47, No.3, pp.565-577.

[Maccheroni et al., 06a] MACCHERONI, F.; MARINACCI, M. & RUSTICHINI, A. (2006): *Ambiguity Aversion, Robustness, and the Variational Representation of Preferences*, Econometrica, Vol.74, No.6, pp.1447-1498.

[Maccheroni et al., 06b] MACCHERONI, F.; MARINACCI, M. & RUSTICHINI, A. (2006): *Dynamic Variational Preferences*, Journal of Economic Theory, Vol.128, pp. 4-44.

[McNeil et al., 05] MCNEIL, A.; FREY, R. & EMBRECHTS, P. (2005): *Quantitative risk management : concepts, techniques, and tools*, Princeton University Press.

[von Neumann & Morgenstern, 44] VON NEUMANN, J. & MORGENSTERN, O. (1944): *Theory of Games and Economic Behavior* , Second Edition, Princeton University Press, Princeton, NJ, 1947.

[Revuz & Yor, 91] REVUZ, D. & YOR, M. (1991): *Continuous Martingales and Brownian Motion*, Springer-Verlag, Berlin.

[Riedel, 04] RIEDEL, F. (2004): *Dynamic Coherent Risk Measures*, Stochastic Processes and their Applications, Vol.112, pp.185-200.

[Riedel, 09] RIEDEL, F. (2009): *Optimal Stopping for Multiple Priors*, Econometrica, Vol. 77(3), pp.857-908.

[Roorda & Schumacher, 05] ROORDA, B. & SCHUMACHER, H. (2005): *Time Consistency Conditions for Acceptability Measures, with an Application to Tail Value at Risk*, Insurance: Mathematics and Economics, Vol. 40(2), pp. 209-230.

[Savage, 54] SAVAGE, L.J. (1954): *The Foundation of Statistics* , John Wiley and Sons, New York. Revised and enlarged edition, Dover, New York, 1972

[Schied, 07] SCHIED, A. (2007): *Optimal Investment for Risk- and Ambiguity-Averse Preferences: A Duality Approach*, Finance and Stochastics, Vol.11, pp.107-129.

[Schnyder, 02] SCHNYDER, M. (2002): *Die Hypothese finanzieller Instabilität von Hyman P. Minsky – Ein Versuch der theoretischen Abgrenzung und Erweiterung.*

[Shiryayev, 78] SHIRYAYEV, A. N. (1978): *Optimal Stopping Rules*, Springer-Verlag, New York.

[Treviño, 09] TREVIÑO-AGUILAR, E. (2009) *$\mathcal{T}$-Systems and the lower Snell Envelope*, arXiv: 'Quantative Finance Papers', No 0902.4245v1.

Kurzer Lebenlauf

Monika Bier

Geboren am 19. März 1980

Abitur 1999 am Gymnasium am Stadtpark, Krefeld

Studium der Mathematik an der
Rheinischen-Friedrich-Wilhelms-Universität Bonn

Diplom-Mathematikerin 12/2005

Promotionsstudium an der Rheinischen-Friedrich-Wilhelms-Universität Bonn und der Universtät Bielefeld ab 3/2006

Promotion zum Dr. rer. pol. an der Wirtschaftswissenschaftlichen Fakultät, Universtät Bielefeld, 9/2010