

Time consistent pricing of options with embedded decisions

J. Gerer^a, G. Dorfleitner^a

^a *University of Regensburg, Department of Finance, 93040 Regensburg, Germany*

Abstract

Many financial contracts are equipped with exercise rights or other features enabling the parties to actively shape the contract's payoff. These decisions pose a great challenge for the pricing and hedging of such contracts. Yet, the literature lacks a consistent way of dealing with these decisions, and instead only provides methods for specific contracts and not transferable to other models.

In this paper we present a framework that allows us to separate the treatment of the decisions from the pricing problem and derive a general pricing principle for the price of an option with decisions by both parties. To accomplish this we present a general version of the duality between acceptance sets and pricing functions, and use it to translate the pricing problem into the language of acceptance. Expressing certain aspects of economic behavior in this language is sufficient to fully eliminate the decisions from the problem.

Further, we demonstrate why time consistent pricing functions are crucial when dealing with options with embedded decisions and how the ad-hoc pricing functions used in many contributions can be derived if time consistency is added to our minimal set of assumptions.

Keywords: derivatives pricing, embedded decisions, acceptance sets, time consistency, early exercise

JEL classification codes: G13, C61

1. Introduction

A theory of option pricing should be a rational theory that tries to explain more with less by formally deriving far reaching results from a small set of assumptions—a property that we call *scope of a theory*. A sharp distinction between assumptions and results as well as between motivation of the former and proof of the latter is crucial for the *quality of a theory*. The pursuit of a theory with such properties is witnessed by the literature from the past sixty years in this field, most notably by the titles of important papers such as “Rational theory of warrant pricing”, [Samuelson \(1965\)](#) or “Theory of rational option pricing”, and the use of Occam’s razor to “analyze the assumptions to determine which ones are necessary” ([Merton, 1973](#)).

Usually, derivative pricing theories focus on the contingency of the contracts, while often not applying the same rigor to the decisions embedded in the contract. Of course, the holder’s decision at expiry of European options is trivial, yet there exists many contracts that include complex decisions by either one or both of the parties (i.e warrants). Furthermore, every contract considered from the point of view of a hedger includes a vast number of non-trivial hedging decisions. Current option pricing literature does not handle decisions on the theory level but instead on an ad-hoc basis and—making things worse—in a way that is tightly entangled with the description and specifics of the model at hand.

In this paper we will present the foundation of a pricing theory that is powerful enough to abstract over the treatment of decisions in financial contracts and that can be derived from a small set of easily accessible assumptions and axioms. It is inspired by the current research on *acceptance sets* and their connection to coherent or convex risk or monetary utility functions ([Artzner et al., 1999](#); [Föllmer and Schied, 2002](#); [Artzner et al., 2007](#)). Even though our pricing functions are more general and not restricted to convex, coherent measures or monotone measures, we are able to derive this connection by explicitly employing the notion of an *acceptable opportunity* (a generalization of *arbitrage opportunities*, see [Carr et al., 2001](#)). Furthermore, we do not rely on a specific model for market price dynamics.

The literature on pricing of American options or warrants serves as an illustrative example of the problem we are trying to solve. The relevant publications can be classified into three stages. The first stage lies in the pre-risk neutral world, with its most prominent representatives being [Samuelson \(1965\)](#); [McKean \(1965\)](#); [Chen \(1970\)](#); [Merton \(1973\)](#). While still struggling from the lack of a meaningful notion of the price of a derivative security, these authors simply postulate the properties that the price of an American option should fulfill. [Samuelson \(1965\)](#) and [McKean \(1965\)](#) both postulate that the price cannot be lower than the values of either the exercised or continued versions of the option. [McKean \(1965\)](#) goes on to define the price as the smallest value that fulfills these properties. [Samuelson \(1965\)](#) fixes the price to be the maximum value of exercised and continued version of the option. Both authors then derive formulae for different stochastic market models. [Chen](#)

(1970) defines the American option as a series of European-style compound options, for which the holder at each time period either receives the exercise price or the next option depending on which one has a greater value. In his seminal work, Merton (1973) postulates the existence of a continuation region of a certain shape within which the price satisfies the Black-Scholes-Merton equation, and that its boundary be chosen in a price-maximizing way.

These theories can be considered to be of high quality, as the postulates and assumptions are clearly introduced as such at the beginning. The fact that formally complex properties of the price are simply postulated or defined as opposed to being consequences of simpler assumptions somehow limits their scope from a theoretical modeling perspective. This of course, does not apply to the derivation of the Black-Scholes-Merton equation. Instead, it applies to the situation concerning the theoretical treatment of the early exercise feature and the resulting decisions embedded in the contract.

Among the noteworthy contributions of the next stage are Brennan and Schwartz (1977); Parkinson (1977); Cox et al. (1979); Geske and Johnson (1984). The aim of these articles is shifted towards providing usable algorithms to calculate the price of an American option. To this end they merely adopt the theoretical foundations of earlier contributions, thus inheriting their limited scope. Also, the quality — from the perspective of the treatment of decisions — suffers from an obscure use of assumptions, possibly even their complete absence (e.g. Brennan and Schwartz, 1977).

The third stage concludes the development of a theory for the decisions embedded in American options, providing a basis for subsequent publications. Bensoussan (1984), like Samuelson (1965), but in a different framework, starts by postulating a type of complementarity problem for the price and then shows how the price is the solution to an optimal stopping problem. This is taken up by Karatzas (1988, 1989), who defines the price as the smallest amount of capital required to set up a super-replication strategy and is able to arrive at the same conclusion. In this setup, the explicit distinction between assumptions and theorems as well as between motivation and proofs warrants the quality. But more importantly, Karatzas (1989) finally fulfills the requirement of scope: His analysis opens the door to obtain all previously published results from one easily accessible definition of the price.¹

The academic literature on American options is not an isolated case. The following shortcomings are equally applicable to the literature on all other types of options with embedded decisions (chooser, passport, shout, swing options etc.) as well as many textbooks on derivative pricing. The robustness and validity of the methods used is hard to verify, often difficult to follow with rigor, and different approaches are impossible to compare without further (mathematically involved) investigation. Contributions such as Myneni (1992), providing a detailed survey

¹See Myneni (1992) for a detailed summary of the derivations and equivalences of the different formulations for the price of an American option.

and proof of equivalence for the different approaches to pricing an American option, show that these problems can be overcome with time for particular contracts in particular frameworks.

The much larger problem is the lack of any progress towards a coherent theory for options with embedded decisions. This results in a situation, in which the methods developed and the knowledge gained are not transferable; they do not actually deepen our understanding of the matter and provide no insight into the nature of the underlying problem. Instead, there remain many unanswered questions, most importantly: Are the various postulated methods consistent? Which assumptions about the decision making process are needed to derive the current results? Is the argumentation also valid in different models? Is an exclusion of the possibility of *clairvoyance* (as done by Karatzas, 1988) actually necessary for deriving the pricing equations? Will the calculated price be correct if this assumption fails to hold?

While it could be the case that answering these questions was not the intention of earlier contributions, many articles still contain a considerable amount of argumentation motivating their numerous and, at times, quite obscure and complex assumptions. To answer these questions, we offer a framework in which all the postulated price properties can be (formally) derived from much simpler principles. This renders the mostly unsatisfactory explanations superfluous and dramatically reduces the argumentative burden. As an example we will demonstrate below, how the arbitrage-free price of an option in an incomplete market can be formally derived within our theory from a small set of assumption.

Section 2 commences with the usual probabilistic setting of a filtered probability space. On top of that we will define *decision procedures*, which describe the choices made by the agent or counterparty in every possible evolution of the world. In the language of probability theory, a decision procedure is a stochastic process φ , where $\varphi_t(\omega)$ stands for the choice made at time t in the world state ω . With the concept of decision procedures we are able to describe options with embedded decisions. Traditionally, options are modeled by their cumulative discounted payoffs expressed as random variables. A natural generalization to options with embedded decisions is to consider an option's payoff f as a function that assigns every decision procedure φ a random variable $f(\varphi) : \omega \mapsto f(\varphi)(\omega)$ describing how much is paid out in each world state ω , if the agent and counterparty follow the decision procedure φ .

At the core of our formalism lies a duality between acceptance sets and pricing functions derived in Section 3. Acceptance sets are collections of payoffs that are accepted by the agent as *zero cost investments*, i.e. option contracts he or she would enter without any additional payments. Pricing functions convert future random payoffs into prices known today. The same concept applies to payoffs describing options with decisions. However, special care needs to be taken when the payoff depends on past decisions, in which case the price inherits this dependence. Our first major result is a characterization of the essential properties required to derive a bijection between acceptance sets and pricing functions.

In Section 4 we introduce *conservative acceptance* to eliminate decisions from the pricing problem, and give an example deriving the arbitrage-free price. Section 5 discusses why time consistency is essential in pricing options with decisions, provides a characterization of time consistent acceptance sets and derives the price of a general option with decisions. Section 6 concludes with a general discussion of the results.

2. Formal setup

The theory is formulated from the perspective of a single market participant, that we will refer to as *agent*, engaging in financial activities and entering contracts with other agents, called her counterparty.

Assumption 2.1 (Probabilistic world). *All possible evolutions of the world, their physical probabilities and the time-dependence of information about the evolution are described by a filtered probability space $(\Omega, \{\mathcal{F}_t\}_{t \in \mathcal{T}}, \mathbb{P})$, where all points of time are given by the totally ordered set \mathcal{T} .*

Definition 2.1 (Random variables). *Let L_t^G represent all \mathcal{F}_t -measurable random variables into the set $G \subseteq \overline{\mathbb{R}}$. We will use the abbreviations $L_t \equiv L_t^{(-\infty, \infty)}$, $L_t^- \equiv L_t^{[-\infty, \infty)}$, $L_t^+ \equiv L_t^{(-\infty, \infty]}$ and $L_t^\pm \equiv L_t^{[-\infty, \infty]}$. We will employ the convention $\infty - \infty \equiv \infty$.*

Define also the set of positive t -premiums $V_t \equiv \{x \in L_t^+ \mid 0 \stackrel{a.s.}{<} x\}$.

The values ∞ and $-\infty$ represent values higher or lower than any possible real value. Remark 2.3 will provide the rationale for the inclusion of these values.

2.1. Decisions

We assume decisions happen at predetermined times. These times are then used to identify a decision (e.g. when describing a payoff's dependence on a decision). This does not prevent us from handling more complex decision for which the point of time can also be chosen by the agent, for example in contracts with the so called *American* exercise feature.

The set of times at which decision are made is called $\mathbf{T}_d \subseteq \mathcal{T}$. At each point of time $t \in \mathbf{T}_d$ there can be exactly one decision by either the agent or the counterparty.

Remark 2.1. This poses no limitation because multiple decisions by one agent can be combined into one choice tuple and decisions by different agents cannot be effectively simultaneous in practice: The agent can either react to the counterparty's decision or not, implying that a chronological order always exists.

Decisions to be made by the agent happen at times \mathbf{T}_a and decisions by the counterparty at times $\mathbf{T}_c \equiv \mathbf{T}_d \setminus \mathbf{T}_a$.

The indexed sets D_t contain all possible choices at time t .

The decision behavior of the agents will be modeled by *decision procedures*, describing how the choices for a subset of decisions depend on the world state. The set of *decision procedures* for decisions at times $\mathbf{T} \subseteq \mathbf{T}_d$ is abbreviated by $\Phi_{\mathbf{T}}$ and defined as the set of stochastic processes taking at time t values in D_t :

$$\Phi_{\mathbf{T}} \equiv \left\{ \varphi : \mathbf{T} \times \Omega \rightarrow \bigcup_{t \in \mathbf{T}} D_t \mid \varphi_t : \Omega \rightarrow D_t, \text{ for all } t \in \mathbf{T} \right\}$$

We will use the abbreviation $\Phi \equiv \Phi_{\mathbf{T}_d}$.

2.2. Options

Before we can actually describe options (by their payoffs), we need one more assumption:

Assumption 2.2 (Cash-flows). *The timing of discounted payments is irrelevant, i.e. cash-flows are fully determined by their cumulative discounted values.*

A sufficient condition in most theories for this assumption to hold is the existence of a risk-free investment instrument. The complexity added by the use of random processes, i.e. cash flows with timing information, instead of random variables could not be justified within the goals of this work.

By Assumption 2.2 the payoff of an option with embedded decisions can be described by a function specifying the cumulative discounted cash-flow to be *received by the agent* for any possible combination of choices and world states. Additionally, we need to be able to restrict our considerations to payoffs that depend only on a subset of decisions. This is provided by the following definition.

Definition 2.2 (Payoffs). *Define $\mathcal{X}_{\mathbf{T}}^t$ as the set of \mathcal{F}_t -measurable payoffs that depend only on decisions made at times $\mathbf{T} \subseteq \mathcal{T}$:*

$$\mathcal{X}_{\mathbf{T}}^t \equiv \left\{ f : \Phi \rightarrow L_t^{\pm} \mid f(\psi) \stackrel{B}{=} f(\varphi), \text{ if } B \in \mathcal{F}_t \text{ and } \psi_t \stackrel{B}{=} \varphi_t \text{ for all } t \in \mathbf{T} \cap \mathbf{T}_d \right\}$$

Putting a set $B \in \mathcal{F}_{\infty}$ above a comparison operator means conditionally almost surely equal: $x \stackrel{B}{=} y \Leftrightarrow \mathbb{P}(\{x = y\} | B) = 1$, with $\{x = y\} \equiv \{\omega \in \Omega \mid x(\omega) = y(\omega)\}$.

We will use the abbreviations $\mathcal{X}_{\mathbf{T}} \equiv \mathcal{X}_{\mathbf{T}}^{\infty}$ and $\mathcal{X} \equiv \mathcal{X}_{\mathcal{T}}$.

In other words, the values of a payoff in the set $\mathcal{X}_{\mathbf{T}}$ are almost surely determined by decisions at times \mathbf{T} . Making different choices at times outside \mathbf{T} does not lead to different cash flows. As an important example, consider the set $\mathcal{X}_{[t, \infty)}$ containing all payoffs with no decisions before t .

Some further remarks on payoffs and their definition are as follows:

Remark 2.2. This definition gives an indirect description of payoffs. It allows payoffs (as functions from Φ to L_{∞}^{\pm}) to show arbitrarily complex and non-local dependence on the decision procedure, only to restrict this freedom at the same time. A more

straightforward approach would be to define the cash flow described by payoff f for a given decision procedure φ in the world state ω by something like $f(\varphi(\omega))(\omega)$. However, this ω -by- ω definition would be too limiting. As an example, take *stochastic integrals*, which are used extensively in the financial modeling of portfolios, trading gains and hedging. They cannot be defined in a pathwise manner (as the limits involved diverge for almost any ω) and thus cannot be handled by this naive approach. Definition 2.2 alleviates this problem and Corollary B.4 shows that the stochastic integral is in fact included in the definition.

Remark 2.3. The definition uses the set L_t^\pm and thus allows $\pm\infty$ in payoffs. This is required to assure that the price of an option can again be treated as a payoff and to enable iterative application of pricing functions. Both aspects will be of importance when dealing with time consistency (Section 5). The same could be achieved with a restriction to bounded payoffs. However, this limitation would exclude too many important applications of practical relevance.

Remark 2.4. If a random variable $x \in L_\infty^\pm$ is used in the context of payoffs, it is understood as the corresponding constant payoff given by $\psi \mapsto x$, which is an element of \mathcal{X}_\emptyset , and vice versa.

Remark 2.5. If not stated differently, all operators, relations and also suprema and infima used on payoffs are the pointwise versions of their L_∞^\pm , \mathbb{P} -almost sure variants: $fRg \Leftrightarrow \forall \varphi \in \Phi : f(\varphi) \stackrel{\text{a.s.}}{R} g(\varphi)$

Finally we need a way to consider the *effective* payoff that results if an agent or counterparty follows a decision procedure for a certain subset of decisions. These decisions can be considered *fixed* and the effective payoff does not depend on them anymore. We introduce the following notation:

Definition 2.3 (Effective payoff). *For any payoff $f \in \mathcal{X}$ and decision procedure $\varphi \in \Phi_{\mathbf{T}}$ define the effective payoff, $f[\varphi] \in \mathcal{X}_{\mathbf{T} \setminus \mathbf{T}}$ by*

$$f[\varphi](\psi) \equiv f(\varphi \mathbf{1}_{\mathbf{T}} + \psi \mathbf{1}_{\mathbf{T}_d \setminus \mathbf{T}}), \text{ for all } \psi \in \Phi.$$

3. From acceptance to pricing

3.1. Acceptable opportunities

In their ground-breaking work Artzner et al. (1999) “provide [...] a definition of risks [...] and present and justify a unified framework for the analysis, construction and implementation of measures of risk”. They make the “acceptable future random net worths” the center of attention and postulate four economically motivated axioms for acceptability, leading to coherent risk measures, which possess a general representation using “generalized scenarios”. Their framework enjoys great popularity and was generalized to convex (Föllmer and Schied, 2002) and dynamic (Artzner et al., 2007; Cheridito et al., 2006; Cheridito and Kupper, 2011) risk-measures.

Carr et al. (2001) use this notion of acceptability to address the shortcomings of many pricing theories, which either require the existence of complete markets (a questionable requirement) or are unable to predict the smallness of observed spreads, “by expanding the role played by arbitrage opportunities to acceptable opportunities”.

The problem with these contributions is that it is hard to see which axioms are actually needed to ascertain the connection between acceptance sets and pricing functions. Acceptance sets are introduced as black boxes lacking any inner structure. The desired structure is then superimposed using axioms which can be too restrictive (as in the case with coherent risk-measures, that had to be generalized to convex risk-measures by Carr et al., 2001; Föllmer and Schied, 2002).

Our approach is slightly different. We begin one step earlier by formalizing the *acceptable opportunity* directly and motivate its properties. The properties of the acceptance set can then be formally derived. This is not only more natural, but also reduces the number of axioms needed to derive the desired properties to one. Properties and connections that finally lead to a better understanding of what prices and risk premiums are and how they can be modeled and calculated.

First, two remarks about our terminology must be made:

Remark 3.1. We will use *premium* to describe the discounted net amount that is paid by the agent upon entering the option’s contract. We will often use the set of positive t -premiums V_t from Definition 2.1.

Remark 3.2. Payoffs and the options they describe should be understood in the sense of *zero cost investments or opportunities*, i.e. for the question of acceptance an option’s premium is understood to be already included in the payoff (which is possible due to Assumption 2.2).

The introduction of acceptable opportunities is based on the following assumption:

Assumption 3.1 (Acceptable opportunity). *For every payoff with no past decisions it can be answered in every world and at every point of time using only information available at that time, whether the agent accepts it (making it an acceptable opportunity) or not.*

Thus, for each time t and option f there exists an event $\alpha_t(f) \in \mathcal{F}_t$, which encodes the acceptability of f at time t . We impose only property upon α_t :

Axiom 3.1. *For all $B \in \mathcal{F}_t$ and $f \in \mathcal{X}_{[t,\infty)}$ the following holds:*

$$\mathbb{P}(\alpha_t(f)|B) = 1 \iff \left(\mathbb{P}(\alpha_t(g)|B) = 1, \text{ if } g \stackrel{B}{=} f + x \text{ for some } x \in V_t \right)$$

The “ \implies ” direction of this axiom is economically uncontroversial: If an option f is accepted in an event B , then any option is accepted whose payoff in the event B is higher than f ’s by a positive premium. The other direction is simply there

to exclude the pathological case in which an unacceptable opportunity can become acceptable by adding an arbitrarily small premium. As we shall see in Example 4.2, this property is especially important in the context of options with decisions.

Now we can define acceptance sets and derive their properties. We will call any set of payoffs $\mathcal{A} \subseteq \mathcal{X}_{[t,\infty)}$ a *t-acceptance set*.

An acceptance set should contain all acceptable payoffs. Thus, from a given function α_t we can derive the corresponding *t-acceptance set*

$$(3.1) \quad \mathcal{A} \equiv \left\{ f \in \mathcal{X}_{[t,\infty)} \mid \mathbb{P}(\alpha_t(f)) = 1 \right\}.$$

In the following sections we will work directly with acceptance sets, or more specifically *proper* acceptance sets, which have the same properties that Axiom 3.1 induces in \mathcal{A} (cf. Corollary 3.1):

Definition 3.1 (Proper acceptance sets). *A t-acceptance set \mathcal{A} is called proper if it is t-compatible (see below) and*

$$(3.2) \quad \mathcal{A} = \left\{ f \in \mathcal{X}_{[t,\infty)} \mid \{f + x \mid x \in V_t\} \subseteq \mathcal{A} \right\}.$$

Properness uses the following definition of *t-compatibility*:

Definition 3.2 (*t-compatibility*). *A non empty set X of functions from some set G into L_t^\pm is t-compatible, if for all $\{x_n\} \subseteq X$ and mutually disjoint $\{B_n\} \subseteq \mathcal{F}_t$ with $\mathbb{P}(\bigcup_n B_n) = 1$ it holds $\sum_n x_n \mathbb{1}_{B_n} \in X$, where $\mathbb{1}_{B_n}$ is the indicator function of the set $G \times B_n$.*

Corollary 3.1. *If α_t satisfies Axiom 3.1 and \mathcal{A} from Eq. (3.1) is not empty, then \mathcal{A} is a proper acceptance set.*

Proof. See Appendix A.1. □

3.2. The price of an option

Formally, we will call any function $\pi : \mathcal{X}_{[t,\infty)} \rightarrow L_t^\pm$, from the set of options with no decisions before t to the set of *t*-measurable random variables, a *t-pricing function*. It is our aim to construct a *t-pricing function* for a given *t-acceptance set*. The intuitive characterization of the price of an option could be summarized as the highest premium the agent would accept to pay for entering the contract. More precisely this describes the agent's *bid* price. If an option's *ask* price is wanted, it is given by the negative of the bid price of the reversed option.

One possible formalization of this description of the bid price is given by:

Definition 3.3 (Associated pricing function). *For any t-acceptance set \mathcal{A} define its t-pricing function $P[\mathcal{A}]$ by*

$$P[\mathcal{A}](f) \equiv \sup \left\{ x \in L_t^- \mid f - x \in \mathcal{A} \right\} \text{ for all } f \in \mathcal{X}_{[t,\infty)},$$

where \sup stands for the essential supremum.

It is easy to see that this definition actually specifies a t -pricing function, as the supremum always exists in L_t^\pm (Theorem B.1). But instead of the *highest acceptable premium*, which does not exist in general, this definition uses the supremum. Random variables and thus premiums are not totally ordered and the supremum of such sets can be far away from its elements. Figure 1 illustrates how the supremum of the set of acceptable premiums is in general *not acceptable*. However, Corollary 3.2 will ensure that for proper \mathcal{A} the use of the supremum is justified.

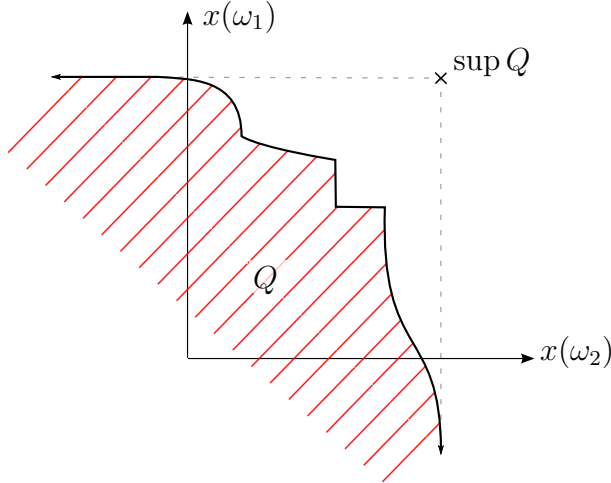


Figure 1: Example of a set of accepted premiums $Q = \{x \in L_t^- \mid f - x \in \mathcal{A}\}$ taking into account Eq. (3.2) from Definition 3.1 in a probability space with $\mathcal{F}_t = \mathcal{P}(\{\omega_1, \omega_2\})$ and no decisions. A premium x can be visualized by a 2D-point $(x(\omega_1), x(\omega_2))$.

In Section 5 we will use the term normalized pricing function:

Definition 3.4 (Normalized pricing function). *A pricing function π is called normalized, if $\pi(0) = 0$. Every pricing function π with $|\pi(0)| < \infty$ has a normalized version $x \mapsto \pi(x) - \pi(0)$.*

We do not require pricing functions to be normalized, as this would unnecessarily restrict the following results and make them more complicated without providing additional value.

3.3. The duality

In this section we present an important connection between pricing functions and acceptance sets. Similar relationships have been derived in previous literature (Artzner et al., 1999; Föllmer and Schied, 2002; Detlefsen and Scandolo, 2005; Cheridito et al., 2006). However, these relationships are formulated only for specialized versions of risk measures and acceptance sets.

Our work uses acceptance sets as the starting point of a pricing theory for options with decision and therefore deeply relies on this relationship, whose generality will carry over to the generality of our results. Thus, our focus lies in finding the most general relationship that still permits a sensible definition of the price of an option, i.e. solving the problem of the existence of the maximum and the associated question of validity of Definition 3.3 from the last section.

The relationship we will use, is a duality between proper acceptance sets and cash invariant pricing functions:

Definition 3.5 (Cash invariance). *A t -pricing function π is called cash invariant if for each $f \in \mathcal{X}_{[t,\infty)}$ and $x \in L_t^+$ it holds: $\pi(f + x) = \pi(f) + x$*

Cash invariance ensures that adding a premium to any option simply increases its price by that amount. This clearly is a desirable property for any theory and, as Theorem 3.1 will show, holds for $P[\mathcal{A}]$ if \mathcal{A} is proper.

The inverse duality operation, which derives an acceptance set from a pricing function, is the literal translation of “the agent accepts any option he would pay a non-negative premium for”:

Definition 3.6 (Dual acceptance set). *For any t -pricing function π define its dual t -acceptance set*

$$A[\pi] \equiv \left\{ f \in \mathcal{X}_{[t,\infty)} \mid 0 \leq \pi(f) \right\}.$$

Now, the complete duality can be formally stated:

Theorem 3.1 (Duality). *For any t , a bijection between the set of cash invariant t -pricing functions and the set of proper t -acceptance sets exists. The bijection and its inverse are given by Definitions 3.3 and 3.6. I.e. for any cash invariant t -pricing function π : (1) $A[\pi]$ is a proper t -acceptance set, and (2) $P[A[\pi]] = \pi$ and for any proper t -acceptance set \mathcal{A} : (3) $P[\mathcal{A}]$ is a cash invariant t -pricing function, and (4) $A[P[\mathcal{A}]] = \mathcal{A}$.*

Proof. See Appendix A.2. The key step, which exploits the inner structure of acceptance sets and enables us to arrive at these new results, is found in Lemma A.3. \square

With this theorem we can answer the question of the previous section:

Corollary 3.2. *For any proper acceptance set \mathcal{A} and option $f \in \mathcal{X}_{[t,\infty)}$ it holds:*

$$P[\mathcal{A}](f) = \max \left\{ x \in L_t^- \mid f - x \in \mathcal{A} \right\}, \text{ if } P[\mathcal{A}](f) < \infty$$

Proof. See Appendix A.3 \square

But beyond this, the duality proves that our formalization is consistent with our intuition. And more importantly, it provides the justification for using the two notions, acceptance sets and pricing functions, interchangeably in developing an

axiomatic option pricing theory. Properties that are best expressed for one of the two can be easily translated for the other.

Furthermore, it can be used to prove the important property of *locality* for any cash invariant pricing function.

Corollary 3.3. *Any cash invariant t -pricing function, π , is also local, i.e. $\pi(f) \stackrel{B}{=} \pi(g)$, if $f \stackrel{B}{=} g$ and $B \in \mathcal{F}_t$.*

Proof. See Appendix A.4. □

4. Decisions

The aim of this work is the development of a pricing theory for options with embedded decisions. The duality result from the previous section enables us to develop our theory in terms of the more directly accessible language of acceptance sets. The derivation of the pricing function then merely becomes a mechanical exercise.

4.1. The counterparty's decisions

What is the agent's price of a contract whose payoff depends on a decision by the counterparty? First we give an intuitive answer to the question, when does the agent accept such a contract, and then derive the pricing function using the duality from the last section.

The agent has neither previous knowledge about nor influence on the counterparty's decision. The counterparty has to be considered completely free in its choice. This suggests the following acceptance set, which is employed implicitly in most of the option pricing literature: The agent accepts an option if and only if the option is acceptable for any possible behavior of the counterparty. It is important to understand, that this presumes nothing about the counterparty's actual behavior.

To formalize this we need a given set of *admissible decision procedures* S for decisions by the counterparty, which, at this point, does not need to be specified further:

Definition 4.1. *For a given t -acceptance set \mathcal{A} and a set of admissible decision procedures S define the conservative acceptance set for counterparty decisions:*

$$\mathcal{A}^{\forall S} \equiv \left\{ f \in \mathcal{X}_{[t, \infty)} \mid \forall \varphi \in S : f[\varphi] \in \mathcal{A} \right\}$$

The agent's price corresponding to this acceptance set is given by the lowest price attainable by any decision procedure in S :

Theorem 4.1. *If \mathcal{A} is a proper t -acceptance set with pricing function π , then $\mathcal{A}^{\forall S}$ also is a proper t -acceptance set and its pricing function is given by*

$$P[\mathcal{A}^{\forall S}](f) = \inf_{\varphi \in S} \pi(f[\varphi]) \text{ for all } f \in \mathcal{X}_{[t, \infty)}$$

Furthermore, the agent's price for any actual decisions procedure followed by the counterparty is in general equal to or higher than this price. The difference adds to the agent's profit. However, the counterparty can make this profit arbitrarily small (if the infimum is finite).

Proof. See Appendix Section A.6. □

This type of acceptance is called *conservative* because it involves no estimate of the counterparty's behavior.

For more complex and realistic problems this kind of acceptance can be too limiting. In these cases non-conservative acceptance sets, that by definition cannot insure against every possible behavior of the counterparty, are needed. They are of interest if the counterparty is somehow limited in his or her actions, acts upon a different maxim (like maximization of another objective function, like a utility function or the value of some larger portfolio) or in cases of market access or information asymmetries, e.g. retail banking customers.

The concepts introduced in this section could be extended by a probabilistic description of the counterparty's behavior used to formulate non-conservative acceptance, where the effective payoffs are accepted in some statistical sense.

A less intrusive way to introduce non-conservative acceptance is to declare certain procedures by the counterparty that are theoretically admissible as practically impossible. Formally, this can be achieved by simply restricting the set of admissible decision procedures to a subset of $S' \subset S$. In this case the acceptance set would be $\mathcal{A}^{\forall S'}$ and Theorem 4.1 applies analogously. An important application of such non-conservative acceptance arises for problems of aligned interest between agent and counterparty as in the case of the following example.

Example 4.1. A minority share of common stock issued by a company can be understood as a call option on the company's assets. The "payoff" crucially depends on decisions by the counterparty, i.e. the majority owner and management of the issuing company. In this example, the boundaries of the applicability of conservative acceptance become obvious: even if prohibited by law, management of the company could deliberately steer into bankruptcy. By the admissibility of this procedure, the conservative stock price as given by Theorem 4.1 would be zero. The situation is more accurately described by a smaller set of admissible procedures incorporating the fact that such behavior would hurt the counterparty's own interests.

The remainder of this work is limited to conservative acceptance, which is the predominant — albeit implicitly used — method to handle decision in the current pricing literature. While beyond the scope of this article, one of the motivations for this work is to enable the development of non-conservative models.

4.2. The agent's decisions

The agent's decisions can be handled analogously by first formulating the conditions for acceptance and then deriving the price. As it is now her decision, the

only rational conduct is to accept the option if and only if there exists at least one decision procedure that makes the option acceptable.

Unlike in the counterparty case, a naive translation of this behavior would not in general produce a proper acceptance set — as demonstrated by Example 4.2. This can be cured by loosening the requirement of acceptance of $f[\varphi]$ to the acceptance of $f[\varphi] + x$ for any positive premium x :

Definition 4.2. For a given t -acceptance set \mathcal{A} and set of admissible decision procedures S define the conservative acceptance set of agent decisions:

$$\mathcal{A}^{\exists S} \equiv \left\{ f \in \mathcal{X}_{[t, \infty)} \mid \forall x \in V_t, \exists \varphi \in S : f[\varphi] + x \in \mathcal{A} \right\}$$

In contrast to the counterparty case, we also need to place a restriction on the set S in order to be able to derive an analogous dual pricing function. S has to be t -compatible (Definition 3.2). If this is the case, then the agent's price is given by the highest effective price attainable by any decision procedure:

Theorem 4.2. If S is t -compatible and \mathcal{A} is a proper t -acceptance set with pricing function π , then $\mathcal{A}^{\exists S}$ is also a proper t -acceptance set and its pricing function is given by:

$$P[\mathcal{A}^{\exists S}](f) = \sup_{\varphi \in S} \pi(f[\varphi]) \text{ for all } f \in \mathcal{X}_{[t, \infty)}$$

Furthermore, while this price will in general be higher than the price for any actual decision procedure followed by the agent, she can make this loss arbitrarily small (if the supremum is finite).

Proof. See Appendix Section A.7. □

Example 4.2. This demonstrates why a simpler definition of $\mathcal{A}^{\exists S}$, analogous to $\mathcal{A}^{\forall S}$'s definition, does not ensure properness.

Define a simple ($t = 0$)-acceptance set $\mathcal{A} \equiv \{f \mid f \geq 0\}$, an option $f : \varphi \mapsto \varphi$ with $\varphi \in S \equiv \langle -\infty, 0 \rangle$ paying an arbitrary negative number of the agent's choosing at time $t = 0$.

It is clear to see, that no decision procedure exists, such that the payoff becomes non-negative and thus accepted. Consequently, f would not be in the following alternative version of $\mathcal{A}^{\exists S}$: $f \notin \mathcal{B} \equiv \left\{ f \in \mathcal{X}_{[t, \infty)} \mid \exists \varphi \in S : f[\varphi] \in \mathcal{A} \right\}$. However, if we add any positive value x ($\in V_0$) to f , then there exists a decision procedure, which makes $f + x$ acceptable (of course $\varphi = -x$) and thus $f + x \in \mathcal{B}$. This is a violation of Eq. (3.2) from Definition 3.1 and thus \mathcal{B} is not proper.

An undesirable consequence of this fact is that f is not acceptable, yet it has price zero and thus violates Theorem 3.1.4:

$$P[\mathcal{B}](f) = \sup \left\{ x \in L_0^- \mid \exists \varphi \in S : \varphi \geq x \right\} = \sup \left\{ x \in L_0^- \mid 0 > x \right\} = 0$$

4.3. Arbitrage-free pricing

In this section, we present a derivation of the arbitrage-free price in incomplete markets (which includes the complete market as a special case) for options without decisions using our framework and the above introduced concepts.

The discounted price processes of the market's assets are modeled by an N -dimensional semi-martingale $X = (X_t)_{t \geq 0}$. Furthermore, this example only makes sense in an arbitrage-free market or, more precisely, in a market with *no free lunch with vanishing risk*. By the *fundamental theorem of asset pricing*, derived for general processes (with unbounded jumps) by [Delbaen and Schachermayer \(1998\)](#), absence of arbitrage is equivalent to the existence of a sigma-martingale measure:

$$\mathcal{M} \equiv \{Q \sim \mathbb{P} \mid X \text{ is a } Q \text{ sigma-martingale}\} \neq \emptyset.$$

The agent is allowed to hedge against her risk exposure by continuously trading in the markets. To translate this into our framework, we need to model the proceeds of the hedging activity. Let the decision procedure φ describe the number of shares of each asset held in the agent's hedging portfolio at different points of time (formally, $\mathbf{T}_a \equiv [0, \infty)$ and $D_t \equiv \mathbb{R}^N$ for any t , see [Section 2.1](#)). The payoff describing the proceeds of a strategy is defined as the stochastic integral with respect to X (the \cdot denotes scalar product between two vectors):

$$H(\varphi) = \int_0^\infty \varphi_t \cdot dX_t$$

A strategy is *admissible*, i.e. in $S \subseteq \Phi_{[0, \infty)}$, if this stochastic integral is well defined and bounded from below. The last requirement excludes so-called doubling strategies and the possibility of infinite wealth generation by trading (see [Delbaen and Schachermayer, 1994](#), and references therein).

The agent is infinitely risk averse and will only accept investments that almost surely do not lose any money. Using [Definition 4.2](#), her acceptance set is thus given by $\{f \mid f \geq 0\}^{\exists S}$. Because she will not only receive the option's actual payoff, but also the proceeds of her hedging strategy, this set equals the sum of the acceptance set of actual option payoffs, \mathcal{A} , and H , $\{f + H \mid f \in \mathcal{A}\}$, and leads to the following price:

Theorem 4.3. *In an arbitrage-free market ($\mathcal{M} \neq \emptyset$) the bid price of an upper-bounded option $f \in L_\infty$ for an agent with $\{f + H \mid f \in \mathcal{A}\} = \{f \mid f \geq 0\}^{\exists S}$ is given by*

$$P[\mathcal{A}](f) = \inf_{Q \in \mathcal{M}} \mathbb{E}_Q[f],$$

also known as the super-replication price.

Proof. See [Appendix A.8](#). □

4.4. Pricing options with decisions

The last section was an example that used the new formalism to derive a well known result. In this section we will demonstrate how to use Theorems 4.1 and 4.2 to solve the problem of pricing options with decisions by reducing it to the classical pricing of options without decision.

Our starting point is classical option pricing. We presume the existence of a t -pricing function π^0 — with dual \mathcal{A}^0 — defined for options without decisions, i.e. for $\mathbf{T}_d = \emptyset$ in our formalism. The aim is to develop a theory with an even number of decisions, $n \in \mathbb{N}$, taking place at times given by an increasing sequence $(\tau_i)_{i \leq n}$ alternating between counterparty, $\mathbf{T}_a \equiv \{\tau_1, \tau_3, \dots, \tau_{n-1}\}$, and agent, $\mathbf{T}_c \equiv \{\tau_2, \tau_4, \dots, \tau_n\}$. As in the previous sections, we first construct the acceptance set and then derive the price.

To build the full acceptance set from \mathcal{A}^0 , we will handle one the decision at a time, each with its own set of admissible decision procedures denoted by $S_i \subseteq \Phi_{\{\tau_i\}}$ for each $i \leq n$. It is clear that the temporal order of a contract's decisions plays an important role in accepting or pricing it. While later decisions can react to decisions made earlier, earlier decisions are fixed and have to take different possible later decisions into account. As a consequence, the reasoning behind the acceptance sets introduced in Sections 4.1 and 4.2 can only be meaningfully applied to the earliest decision in the contract.

This suggests the following recursive extension of the acceptance set \mathcal{A}^0 . Let \mathcal{D}_i be the acceptance set for options with decisions at or after time τ_i . For an agent's decision at τ_i an option is in \mathcal{D}_i if and only if it is in $\mathcal{D}_{i+1}^{\exists S_i}$. For a counterparty decision we have $\mathcal{D}_i = \mathcal{D}_{i+1}^{\forall S_i}$. This scheme stops at $\mathcal{D}_{n+1} \equiv \mathcal{A}^0$ and the full acceptance set is given by \mathcal{D}_1 , which expands to

$$(4.1) \quad \mathcal{D}_1 = \mathcal{A}^{0, \exists S_n \forall S_{n-1} \dots \exists S_2 \forall S_1}.$$

Equipped with this full acceptance set we can now calculate prices for options with decisions:

Theorem 4.4. *If S_i is t -compatible (Definition 3.2) for every i with $\tau_i \in \mathbf{T}_a$ and $\mathcal{A}^0 \subseteq \mathcal{X}_\emptyset$ is a proper t -acceptance set, then $\mathcal{A}^{0, \exists S_n \forall S_{n-1} \dots \exists S_2 \forall S_1}$ also is a proper t -acceptance set. Its dual pricing function is given by*

$$\pi(f) = \inf_{a_1 \in S_1} \sup_{a_2 \in S_2} \dots \inf_{a_{n-1} \in S_{n-1}} \sup_{a_n \in S_n} \pi^0(f[a_1][a_2] \dots [a_n])$$

for $f \in \mathcal{X}_{[t, \infty)}$.

Proof. See Appendix A.9. □

As motivated above, the order of decisions is important. Using an order different from that in Eq. (4.1) can lead to a different price. For example, it is worth a non-negative premium to be able to react to the counterparty's decision, due to $\sup_x \inf_y f(x, y) \leq \inf_y \sup_x f(x, y)$ for any $f : X \times Y \mapsto Z$.

5. Time consistency

So far, the acceptance and thus price of an option is based on *today's* assessment of effective payoffs corresponding to different *future* decision procedures. In general, there is no guarantee that this assessment remains valid over time.

5.1. Problems with time inconsistency

A family of pricing functions at different times exhibiting this inconsistency is called *time inconsistent*. Using a time inconsistent family of pricing functions introduces an ambiguity into the handling of options with decisions; the result will depend on the point of time at which the effects of different choices are assessed. There is, however, no valid argument to choose one point of time over the other. We will describe the shortcomings that come with each possibility, and thereby make a strong point against the use of time inconsistent methods.

Consider the case, in which the agent buys at time t_1 an option $f \in \mathcal{X}_{\{t_2\}}$ with one decision by the agent at time t_2 for a price $p_a = \sup_{\varphi \in S} \pi(f[\varphi])$ calculated using Theorem 4.4. p_a is as high as possible, with the restriction that at least one decision procedure φ^* for her exists that makes the option acceptable *today* (at time t_1). However, when the time comes for her to actually face the decision, the only rational choice is to choose a φ that maximizes the option's present (i.e. time t_2) value. For time inconsistent pricing functions, this procedure will generally be different from φ^* . The problem is that the effective option resulting from this behavior has a t_1 -price p'_a lower than p_a , due to $\pi_t(f[\varphi]) \leq \sup_{\varphi \in S} \pi(f[\varphi])$ and in retrospect the agent paid too much.

Now consider the case with a counterparty decision at t_2 . The agent accepts the price $p_c = \inf_{\varphi \in S} \pi(f[\varphi])$ for any possible future behavior of the counterparty. However, at time t_2 , the worst case behavior of the counterparty would be to follow the decision procedure ψ that minimizes the option's present (i.e. time t_2) value. For time inconsistent pricing functions, the resulting effective option $f[\psi]$ has a time t_1 -price higher than the infimum p_c . The agent did not pay her highest price and could have offered a more competitive bid price.

On the other hand, if the agent based her prices on φ and ψ from above, the resulting prices p'_a and p'_c would have the following deficiencies: She could force herself at time t_1 to stick to the decision procedure φ^* and consequently offer the more competitive bid price $p_a \geq p'_a$. In the counterparty case, she paid too much, because there are adverse decision procedures by the counterparty that only justify the lower price $p_c \leq p'_c$ and thus leave her with a strictly unacceptable position at time t_1 .

These shortcomings strongly suggest the use of time consistent families of pricing functions.

5.2. Time consistent acceptance and pricing

In this section we extend our framework to express time consistency of pricing functions and acceptance sets and derive a duality between the two, required to obtain the time consistent price of a general option with decisions in the next section. To do this conveniently we formalize the concept of a *family*:

Definition 5.1 (Families). *A time indexed set $\{x_t\}_t$, also written as x ., is called pricing family or acceptance family, if every x_t is a cash invariant t -pricing function or proper t -acceptance set, respectively.*

The duality from Theorem 3.1 extends to families. \mathcal{A} .'s dual is written as $P[\mathcal{A}]$ and π .'s dual is written as $A[\pi]$.

In the simplest case, the acceptance family is time independent. As an example take Section 4.3, where a payoff without decisions is accepted, if and only if it is non-negative. This behavior is clearly independent of time and consequently all acceptance sets of the corresponding family would be given by $\{f | f \geq 0\}$ and contain the same payoffs. In general, acceptance families exhibit more complex time dependency. For example, it could be argued that an option's acceptance in all possible states at a future time implies its acceptance today, i.e. $\mathcal{A}_s \subseteq \mathcal{A}_t$, for all $s \geq t$, which is called *weak acceptance consistency* (Artzner et al., 2007).

So far, pricing functions were only defined for options with no past decisions (cf. Section 3.2). However, working with pricing families to express time consistency requires an extension to general payoffs. The price of an option with past decisions does also depend on these past decisions, i.e. is again a payoff in the sense of Definition 2.2. The following lemma shows how this extension can be defined formally and that it works as expected when based on cash invariant pricing functions:

Lemma 5.1 (Extended cash invariant pricing function). *Let $\varphi|_{\langle -\infty, t \rangle}$ denote the restriction of φ to past decisions. Then for any $f \in \mathcal{X}$ and cash invariant t -pricing function π , the mapping $\varphi \mapsto \pi\left(f[\varphi|_{\langle -\infty, t \rangle}]\right)$ is a payoff with no present or future decisions, i.e. element of $\mathcal{X}_{\langle -\infty, t \rangle}^t$.*

From here on we identify $\pi(f)$ with this payoff, giving rise to an extended version of cash invariance:

$$(5.1) \quad \pi(f + g) = \pi(f) + g, \text{ for any } g \in \mathcal{X}_{\langle -\infty, t \rangle}^t \text{ with } g > -\infty$$

Proof. See Appendix A.11. □

A notion of time consistency we can use within our minimalist setting is the so called *recursiveness*. The price of an option equals the price of any other option paying the first option's normalized price (c.f. Definition 3.4) at some future point of time:

Definition 5.2 (Time consistent pricing family). *A pricing family π_\bullet is called time consistent if for all $s \geq t$ and $f \in \mathcal{X}$:*

$$\pi_t(\pi_s(f) - \pi_s(0)) = \pi_t(f)$$

There is another, more intuitive characterization of time consistency: If at some future point of time one option costs more than another option in every world state, the same should be true today (see also [Detlefsen and Scandolo, 2005](#), and the references therein). These two definitions are equivalent if we impose more restrictions on the pricing family:

Corollary 5.1 (Alternative time consistency formulation). *A pricing family π_\bullet with $|\pi_\bullet(0)| < \infty$ consisting of monotone pricing functions, i.e. for all t and $f, g \in \mathcal{X}$*

$$f \geq g \implies \pi_t(f) \geq \pi_t(g),$$

is time consistent if and only if for all $s \geq t$ and $f, g \in \mathcal{X}$

$$(5.2) \quad \pi_s(f) \geq \pi_s(g) \implies \pi_t(f) \geq \pi_t(g).$$

Proof. See [Appendix A.10](#). □

A time consistent acceptance family can be characterized as follows: The agent accepts an option if and only if she accepts any option, which pays the first option's normalized price calculated at some future time. This definition is equivalent to [Definition 5.2](#):

Theorem 5.1 (Time consistent acceptance family). *π_\bullet is time consistent, if and only if its dual acceptance family, \mathcal{A}_\bullet , is time consistent, which is defined by:*

$$(5.3) \quad f \in \mathcal{A}_t \iff \pi_s(f) - \pi_s(0) \in \mathcal{A}_t, \text{ for all } s \geq t.$$

Proof. See [Appendix A.12](#). □

5.3. Application to options with decisions

A time consistent theory does not suffer from the problems outlined in [Section 5.1](#). As we will demonstrate, there is no ambiguity in handling decisions and as a consequence, pricing options with decisions requires far less argumentation compared to the general (time inconsistent) case from [Section 4.4](#).

We will work in the same setting as in [Section 4.4](#), consisting of n decisions at times $\mathbf{T}_d = (\tau_i)_{i \leq n}$ with the corresponding admissible decision procedures $S_i \subseteq \Phi_{\{\tau_i\}}$ for each i , a proper acceptance family $\mathcal{A}_\bullet^0 \subseteq \mathcal{X}_\emptyset$ for options without decisions and its dual pricing family π_\bullet^0 . In contrast to [Section 4.4](#), the complex construction of the full acceptance family, which includes imposing the specific order of decision elimination, will not be necessary. It suffices to specify how options with *present*

time decisions are handled by the full acceptance family \mathcal{A} . Following the reasoning from Sections 4.1 and 4.2, we assume the following conditions for all $i \leq n$:

$$(5.4) \quad \begin{aligned} \mathcal{A}_{\tau_i} &= \mathcal{A}_{\tau_i}^{\exists S_i} \text{ and } S_i \text{ is } \tau_i\text{-compatible} && \text{if } \tau_i \in \mathbf{T}_a, \\ \mathcal{A}_{\tau_i} &= \mathcal{A}_{\tau_i}^{\forall S_i} && \text{if } \tau_i \in \mathbf{T}_c \end{aligned}$$

This fully determines \mathcal{A} , and thus the price of an option with decisions:

Theorem 5.2 (Time consistent conservative price). *If \mathcal{A}^0 is time consistent, it has only one extension \mathcal{A} , i.e. $\mathcal{A} \cap \mathcal{X}_\emptyset = \mathcal{A}^0$ that is also time consistent and satisfies Eq. (5.4). If π_\bullet^0 is normalized, then \mathcal{A} 's dual pricing family is given by*

$$\pi_t(f) = \pi_t^0 \left(\inf_{a_i \in S_i} \pi_{\tau_i}^0 \left(\sup_{a_{i+1} \in S_{i+1}} \pi_{\tau_{i+1}}^0 \left(\dots \inf_{a_{n-1} \in S_{n-1}} \pi_{\tau_{n-1}}^0 \left(\sup_{a_n \in S_n} \pi_{\tau_n}^0 (f[a_i] \dots [a_n]) \right) \right) \right) \right),$$

for all $f \in \mathcal{X}$ and t with $\tau_{i-1} < t \leq \tau_i$ and $\tau_i \in \mathbf{T}_c$. If $\tau_i < t \leq \tau_{i+1}$, then $\pi_t(f)$ is given by the same expression but without the first infimum.

Proof. See Appendix A.13. □

Compared to Theorem 4.4, this pricing function is more suitable for actual calculation. The recursive structure can in many interesting cases be used to separate the n optimization problems, therefore removing the n -exponent in the time complexity dependence on the size of S_i and enabling numeric calculations. Classic examples would be the valuation of American options (as limiting case of the Bermudan option) or options on trading gains.

6. Discussion

To illustrate how our results can be applied and to round up the discussion initiated in Section 1, we will interpret two common approaches to the pricing of American options within our framework.

One approach lies in the derivation of an optimal stopping problem from economically justifiable principles. For example, Karatzas (1988) defines the price of an American option as the minimal initial capital required to set up a portfolio never worth less than the exercise value, g , and then shows that it is given by

$$(6.1) \quad \sup_{\tau \in \mathcal{S}} \mathbb{E}[g_\tau], \text{ where } \mathcal{S} \text{ denotes the set of stopping times.}$$

Despite its advantages, this approach still suffers from the same problem of a very indirect treatment of the decision. The American exercise feature is deeply intermingled with the pricing model and thus this method of handling decisions in an option contract is neither extensible nor transferable to other problems.

These shortcomings are overcome in our framework. In the most straightforward formulation the decision of the holder lies in choosing the (stochastic) exercise time

and the option's payoff is given by $f : \varphi \mapsto g_{\varphi_0}$. For an infinitely risk averse agent that uses conservative acceptance and is continuously hedging in an arbitrage free market, Theorems 4.1 and 4.3 yield Eq. (6.1).² This modular approach provides the flexibility to experiment with different acceptance mechanism and other than the arbitrage-free pricing functions, like e.g. good-deal bounds.

Another popular approach is based on dynamic programming. Early, noteworthy contributions were Chen (1970) or the binomial method from Cox et al. (1979). The American exercise feature is approximated by a finite number of exercise dates. Following an ad-hoc argument, the price at each decision is defined as the maximum of the exercise and continuation value, which in turn depends on the price at the next decision. This approach bears no resemblance to the above approach and it does not answer the question as to why it gives the right answer either.

This can be directly translated into our framework by setting $D_t = \{1, 0\}$, where 1 stands for “exercise” and 0 for “do not exercise”. The payoff then takes the form $f : \varphi \mapsto g_{\min\{t|\varphi_t=1\}}$ and the dynamic programming problem follows directly from Theorem 5.2. In addition to conservative acceptance, it requires the pricing function π^0 to be time consistent, a crucial but not usually explicitly mentioned property.

The same result can be derived from our formulation of the first approach. The bijection between stopping times φ_0 that take only finitely many values and the set of discrete-time adapted processes taking values in $\{0, 1\}$, leads to the above formulation and makes the relationship and compatibility of the two approaches obvious.

7. Conclusion

In this article we close the argumentative gap between pricing theories for classical payoffs and a theory for options with decisions.

This is accomplished in two steps. The first step lies in the derivation of a duality between *cash invariant* pricing functions and so called *proper* acceptance. Proper acceptance sets unite some meaningful properties, but are more general than coherent or convex acceptance sets. This duality effectively shows that acceptance sets and pricing functions contain the same information.

In the second step, it allows us to solve the problem of pricing options with decisions. As acceptance sets are best suited to model agent behavior, we provide a characterization of acceptance of options with decisions corresponding to what is implicitly performed in the option pricing literature, and which we call *conservative acceptance*. Theorem 4.4 and Theorem 5.2 then translate these acceptance sets via the duality into pricing functions for options with decisions.

²While the theorem gives an infimum, following our sign conventions in Sections 2.2 and 3.2, the price for the writer is given by the negative price of $-f$, and thus by the infimum-supremum duality equals the supremum.

Conservative acceptance, however, fails to capture more complex real world situations. For example, alignment of interests between agent and counterparty (as in Example 4.1) demands a different kind of acceptance. Our contribution consists of a concise framework and a precise characterization of the status quo, i.e. conservative acceptance, implicitly assumed in today's pricing theories. Our aim is to enable further research and the development of non-conservative pricing theories.

As demonstrated in the previous section, our results enable a consistent and modular treatment of the decisions within option contracts. Instead of constantly reinventing the wheel, it suffices to motivate how options with decisions are accepted and then use the duality and our formalism to determine the price. This minimizes the argumentative burden and provides a basis to investigate new ways of treating decisions. Furthermore, the modular nature ensures that different approaches can be transferred to a wide variety of pricing methods and models.

One underdeveloped area that will significantly benefit from these methods is the pricing and hedging of options with complex decisions for *both* holder and writer. In this area, in which existing approaches fail, our framework proves most fruitful and can be used to generate new insights. Gerer and Dorfleitner (2016) apply our results to the problem of pricing and realistically hedging American options featuring a complex interplay between exercise and hedging decisions. Based on the results provided in this article, they provide the optimal solution to the full problem, which is — to our knowledge — the first of its kind.

Appendix A Proofs

Definition A.1 (Time t -acceptable premiums). $\mathcal{K}[\mathcal{A}, f] = \{x \in L_t^- \mid f - x \in \mathcal{A}\}$

A.1 Proof of Corollary 3.1

Proof. Equation (3.2) from Definition 3.1 follows directly from Axiom 3.1 and the definition of \mathcal{A} .

To prove t -compatibility assume $\mathbb{P}(\alpha_t(f_n)) = 1$ and therefore $\mathbb{P}(\alpha_t(f_n)|B_n) = 1$ for all $n \in \mathbb{N}$. Take any $x \in V_t$ and define $h \equiv x + \sum_n f_n \mathbb{1}_{B_n}$. Due to B_n 's disjointness we have $h \stackrel{B_n}{=} f_n + x$ and thus by the " \implies " direction of Axiom 3.1, $\mathbb{P}(\alpha_t(h)|B_n) = 1$ and furthermore by Corollary B.2.3 $\mathbb{P}(\alpha_t(h)) = 1$. Using the " \impliedby " direction of Axiom 3.1 yields $\mathbb{P}(\alpha_t(\sum_n f_n \mathbb{1}_{B_n})) = 1$. \square

A.2 Proof of Theorem 3.1

This proof requires the following lemmata:

Lemma A.1. *Given $h \in \mathcal{X}^t$, $C \in \mathcal{F}_\infty$ and a non-empty set of payoffs $Y \subset \mathcal{X}_T^t$, such that $h \stackrel{C}{<} \sup Y$, a sequence $(g_n) \subseteq Y$ exists as well as mutually disjoint events $\{B_n\} \in \mathcal{F}_t$, such that $B_n \subseteq \{g_n > h\} \cup \overline{C}$ and $\mathbb{P}(\bigcup_n B_n) = 1$.*

Proof. By Theorem B.1 and $Y \neq \emptyset$ a sequence $(g_n) \subseteq Y$ exists with pointwise supremum $g(\omega) \equiv \sup_n g_n(\omega)$ such that $g = \sup Y$. Define $D_n \equiv \{h < g_n\} \cup \overline{C}$ and $B_n \equiv D_n \setminus \bigcup_{m=1}^{n-1} B_m$. Disjointedness and $B_n \subseteq \{h < g_n\} \cup \overline{C}$ follow trivially. Next, show:

$$(A.1) \quad \bigcup_n B_n = \overline{C} \cup \bigcup_n \{h < g_n\} = \overline{C} \cup \{\exists n : h < g_n\} = \overline{C} \cup \{h < g\} = \overline{C} \cup (\{h < g\} \cap C)$$

The fourth from the least upper bound property of the supremum. By the hypothesis of the lemma we have $\mathbb{P}(\{h < g\} \cap C) = \mathbb{P}(C)$. Together with \mathbb{P} 's additivity for disjoint events, Eq. (A.1) proves $\mathbb{P}(\bigcup_n B_n) = 1$ \square

Lemma A.2. $\infty \in \mathcal{A}$ and $-\infty \in \mathcal{K}[\mathcal{A}, f]$ for any proper t -acceptance set \mathcal{A} and $f \in \mathcal{X}_{[t, \infty)}$.

Proof. Due to \mathcal{A} 's t -compatibility some $g \in \mathcal{A}$ exists. By Eq. (3.2) from Definition 3.1 and $\infty \in V_t$ we can follow $g + \infty = \infty = f - (-\infty) \in \mathcal{A}$. i.e. $-\infty \in \mathcal{K}[\mathcal{A}, f]$. \square

Lemma A.3. For any proper t -acceptance set \mathcal{A} and $f \in \mathcal{X}_{[t, \infty)}$:

$$0 < P[\mathcal{A}](f) \implies f \in \mathcal{A}$$

Proof. By Lemma A.2 we can apply Lemma A.1 with $h = 0$, $Y = \mathcal{K}[\mathcal{A}, f]$ and $C = \Omega$ and get a sequence $(g_n) \in \mathcal{K}[\mathcal{A}, f]$ and a sequence (B_n) with $B_n \subseteq \{0 < g_n < \infty\}$ for any $n \in \mathbb{N}$. Define $x = \sum_n g_n \mathbb{1}_{B_n}$. By $f - g_n \in \mathcal{A}$, the other properties of B_n and t -compatibility we have $\sum_n (f - g_n) \mathbb{1}_{B_n} = f - x \in \mathcal{A}$. Furthermore, due to $x > 0$ and $x \in L_t$, we have $f = (f - x) + x$ and thus by Eq. (3.2) from Definition 3.1: $f \in \mathcal{A}$. \square

Lemma A.4. For every cash invariant t -pricing function π and $B \in \mathcal{F}_t$, it holds:

$$\pi(f + \infty \cdot \mathbb{1}_{\overline{B}}) = \pi(g + \infty \cdot \mathbb{1}_{\overline{B}}) \implies \pi(f) \stackrel{B}{=} \pi(g)$$

Proof.

$$(A.2) \quad \begin{aligned} \pi(f) \stackrel{B}{=} \pi(f) + \infty \cdot \mathbb{1}_{\overline{B}} &= \pi(f + \infty \cdot \mathbb{1}_{\overline{B}}) \\ &= \pi(g + \infty \cdot \mathbb{1}_{\overline{B}}) = \pi(g) + \infty \cdot \mathbb{1}_{\overline{B}} \stackrel{B}{=} \pi(g) \end{aligned}$$

The second and fourth equation follow from cash invariance (Definition 3.5) and $\infty \cdot \mathbb{1}_{\overline{B}} \in L_t^+$. \square

Proof of Theorem 3.1.1. First we show t -compatibility. $A[\pi]$ is not empty due to $\pi(\infty) \geq 0$, which follows from $0 \in \mathcal{X}_{[t, \infty)}$, $\infty \in L_t^+$, cash invariance of π and the convention from Definition 2.1: $\pi(\infty) = \pi(0) + \infty = \infty \geq 0$.

For t -compatibility, take a sequence f_n with $\pi(f_n) \geq 0$, mutually disjoint $\{B_n\} \subseteq \mathcal{F}_t$ with $\mathbb{P}(\bigcup_n B_n) = 1$ and define $g \equiv \sum_n^\infty f_n \mathbb{1}_{B_n}$. We need to prove $\pi(g) \geq 0$. Using Corollary B.1.2 it remains to show $\pi(g) \stackrel{B_n}{\geq} 0$ for any n : The two functions $f_n + \infty \cdot \mathbb{1}_{\overline{B_n}}$ and $g + \infty \cdot \mathbb{1}_{\overline{B_n}}$ are identical and thus Lemma A.4 can be applied to prove $\pi(g) \stackrel{B_n}{=} \pi(f_n) \geq 0$.

Now prove the “ \subseteq ”-direction in Eq. (3.2) from Definition 3.1. For any f with $0 \leq \pi(f)$, $g \in \mathcal{X}_{[t, \infty)}$ and $x \in V_t$ with $g = f + x$ (pointwise) it holds:

$$0 \leq \pi(f) \leq \pi(f) + x = \pi(f + x) = \pi(g)$$

The first equality uses π ’s cash invariance (Definition 3.5).

For the “ \supseteq ”-direction: Assume for each $n \in \mathbb{N}$ it holds $f + \frac{1}{n} \in A[\pi]$, thus $0 \leq \pi(f + \frac{1}{n})$ and by π ’s cash invariance $0 \leq \pi(f) + \frac{1}{n}$. By taking the limit $n \rightarrow \infty$ yields the result: $f \in A[\pi]$. \square

Proof of Theorem 3.1.2. Take any $f \in \mathcal{X}_{[t, \infty)}$ and show:

$$\begin{aligned} P[A[\pi]](f) &= \sup\{x \in L_t^- \mid f - x \in A[\pi]\} = \sup\{x \in L_t^- \mid \pi_t(f - x) \geq 0\} \\ &= \sup\{x \in L_t^- \mid \pi_t(f) \geq x\} = \pi_t(f) \end{aligned}$$

The first two equations are A ’s and P ’s definitions and the second uses cash invariance (Definition 3.5) of π and the convention from Definition 2.1. The last equation: The upper bound property of $\pi_t(f)$ is trivial and to prove the *least* upper bound property, assume $h \in \mathcal{X}_{[t, \infty)}$ is another upper bound and define the sequence $g_n \equiv \min(n, \pi(f))$. By definition of a t -pricing function, we have $\pi(f) \in L_t^\pm$ and thus $g_n \in L_t^-$ and furthermore $\pi(f) \geq g_n$ for any $n \in \mathbb{N}$. As an upper bound h fulfills $h \geq g_n$, which in the limit $n \rightarrow \infty$ proves that $h \geq \pi(f)$ \square

Proof of 3.1.3. It is a t -pricing function by Theorem B.1. It remains to show cash invariance.

Due to Lemma A.2 the following proof can ignore the case of empty sets. Take any $x \in L_t^+$ and define $g \equiv P[\mathcal{A}](f + x) = \sup \mathcal{K}[\mathcal{A}, f + x]$, $h \equiv P[\mathcal{A}](f) = \sup \mathcal{K}[\mathcal{A}, f]$, and the set $B \equiv \{x = \infty\}$.

1) $h + x \stackrel{B}{\leq} g$: By taking into account the convention from Definition 2.1, it is easy to see that for any $n \in \mathbb{N}$ it holds $f + x - (n - \infty \cdot \mathbb{1}_{\overline{B}}) = \infty$, which is element of \mathcal{A} by Lemma A.2. Consequently, $n - \infty \cdot \mathbb{1}_{\overline{B}} \in \mathcal{K}[\mathcal{A}, f + x]$ and the upper bound property of the supremum g directly implies $n \stackrel{B}{\leq} g$ and in the limit $g \stackrel{B}{=} \infty \stackrel{B}{\geq} h + x$.

2) $h + x \stackrel{\overline{B}}{\leq} g$: Take any $y \in \mathcal{K}[\mathcal{A}, f]$ and define $z \equiv y + x \cdot \mathbb{1}_{\overline{B}} \in L_t^-$, such that $f + x - z = (f - y) \cdot \mathbb{1}_{\overline{B}} + \infty \cdot \mathbb{1}_B$. From $f - y \in \mathcal{A}$, $\infty \in \mathcal{A}$ (by Lemma A.2) and t -compatibility we derive $f + x - z \in \mathcal{A}$, i.e. $z \in \mathcal{K}[\mathcal{A}, f + x]$. The upper bound property of the supremum g and z ’s definition yield $g - x \stackrel{\overline{B}}{\geq} y$. Together with the

previous result $g \stackrel{B}{=} \infty$ we get $g - x \geq y$ and thus $g - x$ is an upper bound of $\mathcal{K}[\mathcal{A}, f]$. By its definition, h has to be the *least* upper bound of this set, thus $h \leq g - x$, which proves the remaining $h + x \stackrel{\overline{B}}{\leq} g$.

3) $g \leq h + x$: For any $y \in \mathcal{K}[\mathcal{A}, f + x]$ it holds $f - (y - x) = f + x - y$ (taking into account the convention from Definition 2.1). Thus due to $f + x - y \in \mathcal{A}$ we have $y - x \in \mathcal{K}[\mathcal{A}, f]$. The upper bound property of the supremum h entails $y - x \leq h$. Or in other words, $h + x$ is an upper bound of the set $\mathcal{K}[\mathcal{A}, f + x]$. By its definition, g has to be the *least* upper bound of the assertion follows. \square

Proof of “ \supseteq ” in Theorem 3.1.4. Take any $f \in \mathcal{A}$. It directly follows that $0 \in \mathcal{K}[\mathcal{A}, f]$, and by the upper bound property of the supremum we have $P[\mathcal{A}](f) \geq 0$. The assertion follows from Definition 3.6 of $A[P[\mathcal{A}]]$. \square

Proof of “ \subseteq ” in Theorem 3.1.4. Take any $f \in A[P[\mathcal{A}]]$. By Definition 3.6 this implies $0 \leq P[\mathcal{A}](f)$. For any $x \in V_t$, $P[\mathcal{A}]$ ’s cash invariance ensures $0 < x \leq x + P[\mathcal{A}](f) = P[\mathcal{A}](f + x)$. Now, Lemma A.3 implies $f + x \in \mathcal{A}$. Using Eq. (3.2) from Definition 3.1 we arrive at $f \in \mathcal{A}$. \square

A.3 Proof of Corollary 3.2

Proof. Define $\pi \equiv P[\mathcal{A}]$. We have to show that $\sup \mathcal{K}[\mathcal{A}, f] \in \mathcal{K}[\mathcal{A}, f]$, i.e. $\pi(f) \in L_t^-$ and $f - \pi(f) \in \mathcal{A}$. The first statements follows from Theorem B.1 and the hypotheses $\pi(f) < \infty$. Now, prove the second statement using

$$\pi(f - \pi(f)) = \pi(f) - \pi(f) \geq 0 \implies f - \pi(f) \in A[P[\mathcal{A}]] = \mathcal{A},$$

which follows from cash invariance, Definition 3.6 and Theorem 3.1.4. \square

A.4 Proof of Corollary 3.3

Proof. Assume $B \in \mathcal{F}_t$, $f, g \in \mathcal{X}$ and $f \stackrel{B}{=} g$ and define $f' = f + \infty \cdot \mathbf{1}_{\overline{B}}$ and $g' = g + \infty \cdot \mathbf{1}_{\overline{B}}$. By Lemma A.4 it remains to show $\pi(f') = \pi(g')$, which by cash invariance and Theorem 3.1.2 follows from

$$f' - x \in A[\pi] \Leftrightarrow g' - x \in A[\pi].$$

We prove the “ \Rightarrow ”-direction of this equation, the other follows analogously. Assume $f' - x \in A[\pi]$. It holds for every $y \in V_t$ and $h = (g' - x) + y$, that $h = (f' - x) + y$ and thus by Theorem 3.1.1 and Eq. (3.2) from Definition 3.1, $h \in A[\pi]$, which by the other direction of Eq. (3.2) from Definition 3.1 proves $g' - x \in A[\pi]$. \square

A.5 Properties of \inf and \sup

Lemma A.5. *The following two mappings inherit π 's cash invariance:*

$$f \mapsto \sup_{\varphi \in S} \pi(f[\varphi]) \text{ (only if } S \neq \emptyset \text{) and } f \mapsto \inf_{\varphi \in S} \pi(f[\varphi])$$

Proof of the sup-version. By Definition 3.5, we need to show for any $x \in L_t^+$ and $f \in \mathcal{X}_{[t, \infty)}$:

$$\sup_{\varphi \in S} \pi((f + x)[\varphi]) = \sup_{\varphi \in S} (\pi(f[\varphi]) + x) = x + \sup_{\varphi \in S} \pi(f[\varphi])$$

The first equation follows from the facts that fixation commutes with addition (trivially by Definition 2.3), $x[\varphi] = x$ (by Remark 2.4), and finally from cash invariance of π .

For the second equation we need $S \neq \emptyset$ and Theorem B.1, together ensuring the existence of a sequence $(\varphi_n)_{n \in \mathbb{N}} \subseteq S$, such that $g(\omega) \equiv \sup_{\varphi \in S} (\pi(f[\varphi])(\omega) + x(\omega))$ and $g = \sup_{n \in \mathbb{N}} (\pi(f[\varphi_n]) + x)$. It remains to show for all $\omega \in \Omega$:

$$(A.3) \quad g(\omega) = x(\omega) + \sup_{n \in \mathbb{N}} (\pi(f[\varphi_n])(\omega))$$

For all ω with $x(\omega) = \infty$, this follows directly from the convention introduced in Definition 2.1. For all other ω 's, we have $x(\omega) - x(\omega) = 0$ and thus Eq. (A.3) follows from translational invariance of the supremum. \square

Proof of the inf-version. This prove is completely analogous with the only difference being that the S can be empty. However, in this case the assertion follows directly from the convention $\inf \emptyset = \infty$ and the one introduced in Definition 2.1. \square

Lemma A.6. *For any $f \in \mathcal{X}_{[t, \infty)}$, $h, g \in L_t^\pm$, local (Corollary 3.3) t -pricing function π , t -compatible S and $x \in V_t$ it holds:*

If $h < \infty$ and $h \leq \sup_{\varphi \in S} \pi(f[\varphi])$, then a $\psi \in S$ exists with $h \leq \pi(f[\psi]) + x$.

If $-\infty < g$ and $\inf_{\varphi \in S} \pi(f[\varphi]) \leq g$, then a $\psi \in S$ exists with $\pi(f[\psi]) \leq g + x$.

Proof of the inf-version. As locality is invariant under a $\pi \mapsto -\pi$ substitution an inf-version of Lemma A.6 can be derived from it using the inf / sup-duality, $-\sup A = \inf(-A)$. This derivation also relies on $x \in L_t^+$. \square

Proof of the sup-version. Using cash invariance (Lemma A.5) and $h < \infty$ we have $h \stackrel{C}{<} \sup_{\varphi \in S} \pi((f + x)[\varphi])$ with $C \equiv \{-\infty < h\}$. Let $\{B_n\}$ be as in Lemma A.1 with $g_n = \pi(f[a_n] + x)$ for some sequence $\{a_n\}_{n \in \mathbb{N}} \subseteq S$.

Define $b = \sum_n^\infty a_n \mathbf{1}_{B_n}$. By the hypothesis of our lemma that S is t -compatible (Definition 3.2) and B_n 's properties we have $b \in S$. By $\sum_n^\infty a_n \mathbf{1}_{B_n} \stackrel{B_n}{=} a_n$ and

Definitions 2.2 and 2.3 we have $f[b] \stackrel{B_n}{=} f[a_n]$ and locality yields $g_n \stackrel{B_n}{=} \pi(f[b] + x)$. From $B_n \subseteq \{g_n > h\} \cup \{h = -\infty\}$ and cash invariance we obtain $h \stackrel{B_n}{\leq} \pi(f[b]) + x$ and the lemma now follows from $\mathbb{P}(\bigcup_n B_n) = 1$ and Corollary B.1.2. \square

A.6 Proof of Theorem 4.1

Proof. Define $\pi \equiv P[\mathcal{A}]$ and prove for any $f \in \mathcal{X}_{[t, \infty)}$:

$$(A.4) \quad 0 \leq \inf_{\varphi \in S} \pi(f[\varphi]) \iff (\forall a \in S : 0 \leq \pi(f[a])) \iff (\forall a \in S : f[a] \in \mathcal{A})$$

The first equality follows from the infimum's defining greatest lower bound property and the second from Theorem 3.1.4 and \mathcal{A} 's properness.

Furthermore, by Theorem 3.1.3 we follow that π is cash invariant and thus by Lemma A.5 (Section A.5) the inf-pricing function is also cash invariant. Equation (A.4) shows that its dual acceptance set, which is automatically proper by Theorem 3.1.1, equals $\mathcal{A}^{\forall S}$.

The first sentence of the second paragraph follows directly from the lower bound property of the supremum and the last sentence from Lemma A.6 (with $g \rightarrow \inf_{\varphi \in S} \pi(f[\varphi])$). \square

A.7 Proof of Theorem 4.2

Proof of Theorem 4.2. The proof follows in analogy to the proof of Theorem 4.1 (Section A.6). We are required to show for any $f \in \mathcal{X}_{[t, \infty)}$:

$$0 \leq \sup_{\varphi \in S} \pi(f[\varphi]) \iff (\forall x \in V_t : \exists a \in S : f[a] + x \in \mathcal{A})$$

To prove " \Leftarrow ", note that for any $n \in \mathbb{N}$ an $a_n \in S$ with $f[a_n] + \frac{1}{n} \in \mathcal{A}$ exists. Theorem 3.1.4, π 's cash invariance (by Theorem 3.1.3) and the upper bound property of the supremum imply $0 \leq \pi(f[a_n] + \frac{1}{n}) \leq \frac{1}{n} + \sup_{\varphi \in S} \pi(f[\varphi])$. The assertion then follows after taking the limit $n \rightarrow \infty$.

For the other direction, due to cash invariance and locality (Corollary 3.3) we can apply Lemma A.6 with $h = 0$. It shows that for any $x \in V_t$ an $a \in S$ such that $0 \leq \pi(f[a]) + x$ exists and thus by cash invariance and Theorem 3.1.4 we have $f[a] + x \in \mathcal{A}$.

If the supremum is finite, the upper bound property of the supremum and Lemma A.6 (with $h \rightarrow \sup_{\varphi \in S} \pi(f[\varphi])$) ensure, that for any finite $x \in V_t$, there exists some $\psi \in S$ with:

$$\sup_{\varphi \in S} \pi(f[\varphi]) - x \leq \pi(f[\psi]) \leq \sup_{\varphi \in S} \pi(f[\varphi])$$

This proves the last statement of the theorem. \square

A.8 Proof of Theorem 4.3

This proof uses the abbreviation $\varphi \cdot X = \int_0^\infty \varphi_t \cdot dX_t$.

Proof. Take any $f \in L_\infty$ bounded from above by some $z \in \mathbb{R}$, define $\mathcal{C} \equiv \mathcal{K}[\mathcal{A}, f]$ (Definition A.1) and show:

$$\begin{aligned} \mathcal{C} &= \left\{ x \in L_0^- \mid f - x + H \in \{g \mid g \geq 0\}^{\exists S} \right\} \\ (A.5) \quad &= \left\{ x \in L_0^- \mid \forall y \in V_0, \exists \varphi \in S : f - x + \varphi \cdot X + y \geq 0 \right\} \end{aligned}$$

The first equation follows from the definition of \mathcal{A} . The second uses Definition 4.2 and H 's definition.

Now, define $\mathcal{D} = \left\{ x \in L_0^- \mid \exists \varphi \in S : f + \varphi \cdot X \geq x \right\}$ and show

$$(A.6) \quad \sup \mathcal{D} = \sup \mathcal{C}$$

The “ \leq ” inequality follows directly from the trivial $\mathcal{D} \subseteq \mathcal{C}$. To prove the reverse equality by $\sup \mathcal{C}$'s *least* upper bound property, it remains shown that $\sup \mathcal{D}$ is an upper bound of \mathcal{C} : For any $x \in \mathcal{C}$ and $n \in \mathbb{N}$ it follows from Eq. (A.5) that there exists a φ in S , such that $f + \varphi \cdot X \geq x - \frac{1}{n}$. Consequently, $x - \frac{1}{n} \in \mathcal{D}$ and thus $x - \frac{1}{n} \leq \sup \mathcal{D}$, which also holds in the limit $n \rightarrow \infty$.

We can complete the proof:

$$\pi(f) = \sup \mathcal{C} = \sup \mathcal{D} = \inf_{Q \in \mathcal{M}} \mathbb{E}_Q[f]$$

The first equations uses Definition 3.3 of the pricing function, the second uses Eq. (A.6), and the third Theorem 5.12 of Delbaen and Schachermayer (1998, p. 246) with the substitution $g \rightarrow -f$ and the infimum-supremum duality. \square

A.9 Proof of Theorem 4.4

This proof requires the following lemma

Lemma A.7. *Given a proper \mathcal{A}^0 with its dual pricing function π^0 , it holds for all $f \in \mathcal{X}_\emptyset$ and t -acceptance sets \mathcal{A} with $\mathcal{A} \cap \mathcal{X}_\emptyset = \mathcal{A}^0$*

$$P[\mathcal{A}](f) = \pi^0(f).$$

Proof. By Definition 3.3 and $\mathcal{A} \cap \mathcal{X}_\emptyset = \mathcal{A}^0$ it remains to show $\mathcal{K}[\mathcal{A}, f] = \mathcal{K}[\mathcal{A} \cap \mathcal{X}_\emptyset, f]$ for any $f \in \mathcal{X}_\emptyset$, which follows from

$$f - x \in \mathcal{X}_\emptyset \text{ for all } x \in L_t^-, f \in \mathcal{X}_\emptyset.$$

\square

Proof of Theorem 4.4. By Theorems 4.1 and 4.2 it holds for any i with $\tau_i \in \mathbf{T}_c$ and $f \in \mathcal{X}_{[t, \infty)}$:

$$P[\mathcal{D}_i](f) = \inf_{a_i \in S_i} \sup_{a_{i+1} \in S_{i+1}} P[\mathcal{D}_{i+2}](f[a_i][a_{i+1}])$$

We can use this relation recursively until we reach $P[\mathcal{D}_{n+1}](g) = P[\mathcal{A}^0](g)$, with $g \equiv f[a_1] \dots [a_n]$. By Definition 2.3 it is easy to see that $g \in \mathcal{X}_\emptyset$ and thus using $\mathcal{A}^0 \subseteq \mathcal{X}_\emptyset$ and Lemma A.7 we have $P[\mathcal{A}^0](g) = \pi^0(g)$. \square

A.10 Proof of Corollary 5.1

Proof of the “if”. Define $g = \pi_s(f) - \pi_s(0)$ and $B \equiv \{\pi_s(f) > -\infty\}$ ($B = \{g > -\infty\}$ due to $|\pi_s(0)| < \infty$). Assuming Eq. (5.2) we can prove Definition 5.2 by showing

$$\pi_s(g) = \pi_s(f).$$

1) We have

$$\pi_s(g) \stackrel{B}{=} \pi_s(0 + g \cdot \mathbf{1}_B) \stackrel{B}{=} \pi_s(0) + \pi_s(f) - \pi_s(0) = \pi_s(f)$$

The first equation uses locality (by Corollary 3.3), the second cash invariance and $g \cdot \mathbf{1}_B \in L_t^+$ and the third $|\pi_s(0)| < \infty$.

2) $\pi_s(g) \stackrel{\overline{B}}{\geq} \pi_s(f)$ is trivial.

3) Note that $g \leq \infty \cdot \mathbf{1}_B - n$ for all $n \in \mathbb{N}$. Using monotonicity and cash invariance we have $\pi_s(g) \leq \infty \cdot \mathbf{1}_B - n + \pi_s(0)$ for any n and therefore due to $|\pi_s(0)| < \infty$:

$$\pi_s(g) \stackrel{\overline{B}}{\leq} -\infty \stackrel{\overline{B}}{=} \pi_s(f) \quad \square$$

Proof of the “only if”. Assume π_s satisfies Definition 5.2 and show

$$\pi_s(f) \geq \pi_s(g) \implies \pi_t(\pi_s(f)) \geq \pi_t(\pi_s(g)) \implies \pi_t(f) \geq \pi_t(g)$$

The first implication follows from monotonicity of π_t and the second from time consistency. \square

A.11 Proof of Lemma 5.1

Proof. The first statement follows directly from the lemma below. To prove Eq. (5.1), take any $\varphi \in \Phi$ and show:

$$\pi(f + g)(\varphi) = \pi\left((f + g)\left[\varphi|_{\langle -\infty, t \rangle}\right]\right) = (\pi(f) + g)(\varphi)$$

The second equality follows from cash invariance of π , $g \in \mathcal{X}_{\langle -\infty, t \rangle}$ and $g(\varphi) \in L_t^+$. \square

Lemma A.8. *For any $f \in \mathcal{X}_{\mathbf{T}}$ and cash invariant t -pricing function π , the mapping $\varphi \mapsto \pi\left(f\left[\varphi|_{\langle -\infty, t \rangle}\right]\right)$ is an element of $\mathcal{X}_{\mathbf{T} \cap \langle -\infty, t \rangle}^t$.*

Proof. That $\pi\left(f\left[\varphi|_{\langle -\infty, t \rangle}\right]\right)$ is well-defined and \mathcal{F}_t -measurable, follows from the definition of a t -pricing function in Section 3.2.

To prove the remaining properties in Definition 2.2 take $B \in \mathcal{F}_t$ and $\varphi, \psi \in \Phi$ with $\psi_s \stackrel{B}{=} \varphi_s$ for all $s \in \mathbf{T} \cap \langle -\infty, t \rangle \cap \mathbf{T}_d$. The required equation

$$\pi\left(f\left[\varphi|_{\langle -\infty, t \rangle}\right]\right) \stackrel{B}{=} \pi\left(f\left[\psi|_{\langle -\infty, t \rangle}\right]\right)$$

follows directly from locality (Corollary 3.3) and

$$f\left[\varphi|_{\langle -\infty, t \rangle}\right](\rho) \stackrel{B}{=} f\left[\psi|_{\langle -\infty, t \rangle}\right](\rho), \text{ for any } \rho \in \Phi,$$

which in turn follows from $f \in \mathcal{X}_{\mathbf{T}}$ and Definitions 2.2 and 2.3. \square

A.12 Proof of Theorem 5.1

Proof of time consistency of $P[\mathcal{A}_\bullet]$. Define $\pi_\bullet \equiv P[\mathcal{A}_\bullet]$. By Lemma A.8 we have $\pi_s(f) \in \mathcal{X}_{[t,s)} \subseteq \mathcal{X}_{[t,\infty)}$ and thus by Definition 3.3 it remains to show $\mathcal{K}[\mathcal{A}_t, f] = \mathcal{K}[\mathcal{A}_t, \pi_s(f) - \pi_s(0)]$. Take any $x \in L_t^-$ and prove:

$$f - x \in \mathcal{A}_t \iff \pi_s(f - x) - \pi_s(0) \in \mathcal{A}_t \iff (\pi_s(f) - \pi_s(0)) - x \in \mathcal{A}_t$$

The two equivalences follow from Eq. (5.3) and cash invariance (Theorem 3.1.3 and Definition 3.5). \square

Proof of time consistency of $A[\pi_\bullet]$. Define $\mathcal{A}_\bullet \equiv A[\pi_\bullet]$ and show:

$$f \in \mathcal{A}_t \iff \pi_t(f) \geq 0 \iff \pi_t(\pi_s(f) - \pi_s(0)) \geq 0 \iff \pi_s(f) - \pi_s(0) \in \mathcal{A}_t$$

The first and third equivalences follow from Definition 3.6. The second from π 's time consistency. \square

A.13 Proof of Theorem 5.2

This proof requires the following lemma

Lemma A.9. *Take a proper \mathcal{A}_\bullet^0 with dual π_\bullet^0 and some \mathcal{A}_\bullet with dual π_\bullet that satisfies $\mathcal{A}_\bullet \cap \mathcal{X}_\emptyset = \mathcal{A}_\bullet^0$. Taking into account Lemma 5.1, it holds $\pi_t(f) = \pi_t^0(f)$ for any payoff with no decisions at or after time t .*

Proof. This follows directly from Lemmas A.7 and 5.1 and $f[\varphi|_{(-\infty, t)}] \in \mathcal{X}_\emptyset$ for all $\varphi \in \Phi$. \square

Proof of Theorem 5.2. Define $\pi_\bullet \equiv P[\mathcal{A}_\bullet]$. Show for all $j \in \{i, \dots, n\}$ with $\tau_j \in \mathbf{T}_a$, $\tau_{j-1} < s \leq \tau_j$ and $g \in \mathcal{X}$ or $s = \tau_{j-1}$ and $g \in \mathcal{X}_{\mathcal{T} \setminus \{s\}}$:

$$\pi_s(g) = \pi_s(\pi_{\tau_j}(g)) = \pi_s^0(\pi_{\tau_j}(g)) = \pi_s^0\left(\inf_{a_j \in S_j} \pi_{\tau_j}(g[a_j])\right)$$

The first equation follows from \mathcal{A}_\bullet 's time consistency, π_\bullet^0 's normalization and Theorem 5.1. The second follows from Lemma A.9 with $\pi_{\tau_j}(g) \in \mathcal{X}_{(-\infty, s)}$ (by Lemma A.8). And the third from Remark 2.5, Eq. (5.4) and Theorem 4.2.

For j with $\tau_j \in \mathbf{T}_c$, Theorem 4.1 would be needed instead of Theorem 4.2 and inf instead of sup.

To prove the theorem we start with $s \rightarrow t$ and $g \rightarrow f$. Then we can apply this equation recursively with $j \rightarrow j + 1$, $s \rightarrow \tau_j$ due to $g[a_j] \in \mathcal{X}_{\mathcal{T} \setminus \{\tau_j\}}$. \square

Appendix B Mathematical theorems

Theorem B.1 (Existence of the essential supremum). *Suppose that the measure space $(\Omega, \mathcal{F}, \mu)$ is σ -finite. Then the essential supremum of a collection \mathcal{S} of measurable functions into the set \mathbb{R} exists. Furthermore, if \mathcal{S} is nonempty then a sequence $(f_n)_{n \in \mathbb{N}}$ in \mathcal{S} exists such that its pointwise supremum equals (almost surely) the essential supremum of \mathcal{S} .*

Proof. See Chapter V.18 in Doob (1994). \square

Corollary B.1 (Properties of conditionally almost sure). *Let \mathbf{D} be a countable collection of sets with positive probability and $A, B \in \mathbf{D}$.*

- (1) $\mathbb{P}(B \setminus C) = 0 \iff \mathbb{P}(C|B) = 1$
- (2) $(\forall A \in \mathbf{D} : \mathbb{P}(C|A) = 1) \implies \mathbb{P}(C|\bigcup \mathbf{D}) = 1$
- (3) $\mathbb{P}(B|A) = \mathbb{P}(C|B) = 1 \implies \mathbb{P}(C|A) = 1$

Proof of Corollary B.1.1. The statement follows from:

$$\mathbb{P}(B) = \mathbb{P}((B \setminus C) \cup (B \cap C)) = \mathbb{P}(B \setminus C) + \mathbb{P}(C \cap B) = \mathbb{P}(B \setminus C) + \mathbb{P}(C|B)\mathbb{P}(B)$$

Using Corollaries B.2.1 and B.2.2, additivity of \mathbb{P} for disjoint sets and the definition of conditional probability. \square

Proof of Corollary B.1.2. By Corollary B.1.1 it follows from the antecedent for any $A \in \mathbf{D}$: $\mathbb{P}(A \setminus C) = 0$. By using Corollary B.2.3 with subadditivity of \mathbb{P} we can follow $\mathbb{P}((\bigcup \mathbf{D}) \setminus C) = 0$, which again with Corollary B.1.1 yields the result. \square

Proof of Corollary B.1.3. Using Corollary B.1.1 we can follow from the assumption: $\mathbb{P}(A \setminus B) = \mathbb{P}(B \setminus C) = 0$. By using Corollary B.2.4 with monotonicity and subadditivity of \mathbb{P} , it follows $\mathbb{P}(A \setminus C) = 0$, which again with Corollary B.1.1 yields the result. \square

Corollary B.2 (Properties of sets). *For all sets A, B, C and collections of sets \mathbf{D} , it holds: (1) $(A \setminus B) \cap (A \cap B) = \emptyset$, (2) $(A \setminus B) \cup (A \cap B) = A$, (3) $\bigcup_{E \in \mathbf{D}} (E \setminus A) = (\bigcup \mathbf{D}) \setminus A$, and (4) $A \setminus C \subseteq (A \setminus B) \cup (B \setminus C)$.*

The proofs are left to the reader. \square

Corollary B.3 (Convergence in probability preserves equality almost surely). *For any set $B \in \mathcal{F}_\infty$ and two sequences of random variables $(f_n)_{n \in \mathbb{N}}$ and $(g_n)_{n \in \mathbb{N}}$ that converge in probability to f and g , respectively, it holds $f \stackrel{B}{=} g$, if $B \subseteq \{f_n = g_n\}$.*

Proof. Simple set manipulations yield:

$$B \setminus \{f = g\} = B \cap \{\exists m \in \mathbb{N} : |f - g| > \frac{1}{m}\} = \bigcup_{m \in \mathbb{N}} B \cap \{|f - g| > \frac{1}{m}\}$$

Consequently, by Corollary B.1.1 and monotonicity of \mathbb{P} , $f \stackrel{B}{=} g$ is equivalent to

$$(B.1) \quad \mathbb{P}\left(B \cap \{|f - g| > \frac{1}{m}\}\right) = 0 \text{ for all } m \in \mathbb{N}$$

We prove Eq. (B.1) in two steps. First show for any $m, n \in \mathbb{N}$:

$$(B.2) \quad B \cap \{|f - g| > \frac{1}{m}\} \subseteq \{|f - f_n| + |g - g_n| > \frac{1}{m}\} \subseteq \{|f - f_n| > \frac{1}{2m}\} \cup \{|g - g_n| > \frac{1}{2m}\}$$

The first “ \subseteq ” uses $B \subseteq \{f - g = (f - f_n) - (g - g_n)\}$, together with the triangle inequality. By definition of convergence in probability we can then prove:

$$\mathbb{P}\left(B \cap \{|f - g| > \frac{1}{m}\}\right) \leq \mathbb{P}\left(\{|f - f_n| > \frac{1}{2m}\}\right) + \mathbb{P}\left(\{|g - g_n| > \frac{1}{2m}\}\right) \rightarrow 0$$

□

Corollary B.4 (Integration preserves equality almost surely). *For any set $B \in \mathcal{F}_\infty$, integrator X and integrands a_1, b_2 , it holds $\int a_1 dX \stackrel{B}{=} \int a_2 dX$, if $B \subseteq \{a_1 = a_2\}$.*

Proof. Let $(a_i^n)_{n < \infty}$ be a sequence of simple integrands approximating a_i . We can choose them such that $a_1^n = a_2^n$ on B for all n . Define $A_i \equiv \int_t^T a_i dX \equiv \text{plim}_{n \rightarrow \infty} A_i^n$ (limit in probability) with $A_i^n \equiv \int_t^T a_i^n dX$. These integrals of simple integrands are defined path-wise and thus $B \subseteq \{A_1^n = A_2^n\}$ for each n . By Corollary B.3 it holds $P(A_1 = A_2 | B) = 1$. □

References

- Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance* 9(3), 203–228.
- Artzner, P., F. Delbaen, J.-M. Eber, D. Heath, and H. Ku (2007). Coherent multiperiod risk adjusted values and bellman’s principle. *Annals of Operations Research* 152(1), 5–22.
- Bensoussan, A. (1984). On the theory of option pricing. *Acta Applicandae Mathematica* 2(2), 139–158.
- Brennan, M. and E. Schwartz (1977). The valuation of American put options. *Journal of Finance* 32(2), 449–462.
- Carr, P., H. Geman, and D. B. Madan (2001). Pricing and hedging in incomplete markets. *Journal of Financial Economics* 62(1), 131 – 167.
- Chen, A. (1970). A model of warrant pricing in a dynamic market. *The Journal of Finance* 25(5), 1041–1059.
- Cheridito, P., F. Delbaen, and M. Kupper (2006). Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability* 11(3), 57–106.
- Cheridito, P. and M. Kupper (2011). Composition of time-consistent dynamic monetary risk measures in discrete time. *International Journal of Theoretical and Applied Finance* 14(01), 137–162.
- Cox, J. C., S. A. Ross, and M. Rubinstein (1979). Option pricing: A simplified approach. *Journal of Financial Economics* 7(3), 229 – 263.
- Delbaen, F. and W. Schachermayer (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen* 300(1), 463–520.
- Delbaen, F. and W. Schachermayer (1998). The fundamental theorem of asset pricing for unbounded stochastic processes. *Mathematische Annalen* 312(2), 215–250.
- Detlefsen, K. and G. Scandolo (2005). Conditional and dynamic convex risk measures. *Finance and Stochastics* 9(4), 539–561.
- Doob, J. (1994). *Measure Theory*, Volume 143 of *Graduate Texts in Mathematics Series*. Springer.
- Föllmer, H. and A. Schied (2002). Convex measures of risk and trading constraints. *Finance and Stochastics* 6(4), 429–447.
- Gerer, J. and G. Dorfleitner (2016). Optimal discrete hedging of American options using an integrated approach to options with complex embedded decisions. *SSRN eLibrary*.
- Geske, R. and H. E. Johnson (1984). The American put option valued analytically. *Journal of Finance* 39(5), 1511–1524.
- Karatzas, I. (1988). On the pricing of American options. *Applied Mathematics and Optimization* 17(1), 37–60.
- Karatzas, I. (1989). Optimization problems in the theory of continuous trading. *SIAM Journal on Control and Optimization* 27(6), 1221–1259.
- McKean, H. P. (1965). Appendix: A free-boundary problem for the heat-equation arising from a problem of mathematical economics. *Industrial Management Review* 6(2), 32–39.
- Merton, R. C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4(1), 141–183.
- Myneni, R. (1992). The pricing of the American option. *Annals of Applied Probability* 2(1), 1–23.
- Parkinson, M. (1977). Option pricing: the American put. *Journal of Business* 50(1), 21–36.
- Samuelson, P. A. (1965). Rational theory of warrant pricing. *Industrial Management Review* 6(2), 13–31.