



HAL
open science

Formalizing the Cox–Ross–Rubinstein Pricing of European Derivatives in Isabelle/HOL

Mnacho Echenim, Hervé Guiol, Nicolas Peltier

► **To cite this version:**

Mnacho Echenim, Hervé Guiol, Nicolas Peltier. Formalizing the Cox–Ross–Rubinstein Pricing of European Derivatives in Isabelle/HOL. *Journal of Automated Reasoning*, 2020, 64 (4), pp.737-765. 10.1007/s10817-019-09528-w . hal-02990376

HAL Id: hal-02990376

<https://hal.science/hal-02990376v1>

Submitted on 9 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Formalizing the Cox-Ross-Rubinstein pricing of European derivatives in Isabelle/HOL

Mnacho Echenim · Hervé Guiol · Nicolas Peltier

Received: date / Accepted: date

Abstract We formalize in the proof assistant Isabelle essential basic notions and results in financial mathematics. We provide generic formal definitions of concepts such as markets, portfolios, derivative products, arbitrages or fair prices, and we show that, under the usual no-arbitrage condition, the existence of a replicating portfolio for a derivative implies that the latter admits a unique fair price. Then, we provide a formalization of the Cox-Rubinstein model and we show that the market is complete in this model, i.e., that every derivative product admits a replicating portfolio. This entails that in this model, every derivative product admits a unique fair price. In addition, we provide Isabelle functions to compute the fair price of some derivative products.

Keywords Proof assistants · Financial mathematics · Discrete pricing · Isabelle/HOL

1 Introduction

The basic securities that are traded on financial markets (such as shares on the equity market or bonds on the fixed-income market) have a price that is submitted to the law of supply and demand, and depends on the needs of financial actors. Things are not that simple for all securities traded on financial markets, and in particular, determining the price of so-called *derivative*

M. Echenim
Univ. Grenoble Alpes, CNRS, LIG. 38000 Grenoble, France
E-mail: Mnacho.Echenim@univ-grenoble-alpes.fr

H. Guiol
Univ. Grenoble Alpes, CNRS, TIMC. 38000 Grenoble, France
E-mail: Herve.Guiol@univ-grenoble-alpes.fr

N. Peltier
Univ. Grenoble Alpes, CNRS, LIG. 38000 Grenoble, France
E-mail: Nicolas.Peltier@univ-grenoble-alpes.fr

products can be a far from trivial task. A derivative product is a security the value of which depends on that of one or several underlying securities; a typical example is a vanilla call on a share, which gives its holder the right, but not the obligation, to buy the share on a predetermined date at a predetermined price. Obviously, the price of a derivative product should depend on that of its underlyings, but what exactly is this dependency? In fact, is there even a unique price for any derivative? An intuitive answer to the second question is that, similarly to basic securities, the price of a derivative should be unique: if this were not the case, an investor could buy the derivative at the lower price and simultaneously sell it at the higher price, making a profit without investing any money or taking any risks. The investor would have exploited what is called an *arbitrage opportunity* (see, e.g., [14]), and although such opportunities do exist on financial markets, they are exploited by financial actors called arbitragists and tend to disappear quickly. This is the reason why many results in quantitative finance rely on a no-arbitrage hypothesis. Such a hypothesis also permits to provide a more precise definition of what a price for a derivative should be: this should be any value that is neither so high as to induce an arbitrage opportunity for the seller of the derivative, nor so low as to induce an arbitrage opportunity for the buyer. Any price satisfying these conditions is called a fair price for the derivative.

One of the most important results in financial mathematics is the proof by Black, Scholes and Merton [4, 17] that, in the so-called Black-Scholes model of an equity market, every derivative admits a unique fair price, as well as a formula permitting to compute this price. This model is based on the assumption that the price of a risky security can be modeled using so-called *Brownian motions*. Along with the no-arbitrage hypothesis, the authors also suppose that (1) the market is *frictionless*, meaning that securities can be bought or sold with no transaction costs, and (2) investors can buy and sell any amount of the securities, meaning that the quantity of a security held in a portfolio can be any real number, even a negative one if the security has been sold short (i.e., sold by an investor not owning the asset). Since then, there have been a wide variety of mathematical models devised for the pricing of derivative products, adapting the hypotheses of the Black-Scholes-Merton model or modeling other markets, such as the foreign-exchange or commodities markets.

A discrete-time model of an equity market was introduced by Cox, Ross and Rubinstein [6]. This model is based on hypotheses similar to those of the Black-Scholes-Merton model, in which time is continuous, and can actually be viewed as a discrete-time approximation of this model. Evaluating the price of a derivative in this model is more complex than in the continuous time model since no explicit price formula exists. This entails that the CRR model is not frequently used for the pricing of simpler derivatives. But several financial institutions still rely on this model for the pricing of more complex derivatives, such as *American options*, which can be exercised by their buyer at *any* time until the option expires.

In this paper, we present a formalization in Isabelle/HOL [20] of (1) basic notions in financial mathematics such as markets, portfolios etc. (2) the

definition of fair prices for derivative products on equity financial markets, (3) the proof of the uniqueness of fair prices when a replicating portfolio exists in a viable market, and (4) an algorithm to compute fair prices under a risk-neutral probability space. We also formalize the Cox-Ross-Rubinstein model and prove that in this model, every European derivative admits a replicating portfolio, i.e., a portfolio with a value identical to the payoff of the derivative. This work can be viewed as the start of a broader effort to formalize financial mathematics. This first step already contains a large amount of financial notions that are necessary for any advanced financial modeling. It also contains fundamental results on fair prices that, although they may seem intuitive, require an advanced background in Probability theory to be stated formally. The entire formalization is over 13000 lines long. The work presented here strictly subsumes the formalization carried out in [9], which was mainly devoted to the proof that in the model of a market defined by Cox, Ross and Rubinstein [6], every derivative product admits a replicating portfolio. The results presented in this paper can be found in many financial mathematics textbooks [21, 15, 3], with one main difference. Almost all derivative products expire at some point in time (this time is called the *maturity* of the product); in most textbooks, results are presented by considering an arbitrary derivative with a given maturity T , and taking the finite probability space with outcomes consisting of all sequences of T coin tosses. Here we formalize a setting in which *any* derivative can be priced, and use Isabelle's codatatypes [5] to consider non-denumerable probability spaces with outcomes consisting of all infinite streams of coin tosses. This could pave the ground for the handling of derivatives with infinite maturity, such as so-called *perpetual American options*, see, e.g., [21, Sect. 5.4].

Related work. Many results related to financial mathematics have already been formalized in Isabelle. Large parts of Probability theory have been formalized, building up on [11]; and results and concepts frequently used in financial mathematics such as Markov processes or the Central Limit Theorem are also available in Isabelle [12, 1]. Broadly speaking, this work can be viewed as part of an effort on the formalization of concepts in Economy. Several notions and theorems have been formalized in *Social Choice Theory* [18, 23, 8], as well as in *Game Theory* [22] and *Microeconomics* [16]. To the best of our knowledge, other than [9], there has been no formalization of financial mathematics itself.

Organization. This paper is organized as follows. Section 2 contains basic financial notions as well as a summary of the notions from Probability theory that will be used throughout the paper and are already formalized in Isabelle. In Section 3 we define equity markets in discrete time, introducing the notions of portfolios and their values, as well as trading strategies which represent the only reasonable portfolios that can be constructed. Arbitrage opportunities are introduced in Section 4, they permit to define the notion of a fair price for a derivative, and we show that if a derivative admits a *replicating portfolio*, i.e., a portfolio whose value at maturity is identical to the

derivative payoff, then the fair price for this derivative is unique. Section 5 is devoted to the definition of risk-neutral probability spaces, which are based on the existence of *martingales*, and permit to represent the fair price of a derivative as an expectation. In Section 6, these results are applied to the Cox-Ross-Rubinstein model, and an explicit formula for computing the fair price of any derivative is provided. Section 7 contains a detailed illustrative example, showing how in the Cox-Ross-Rubinstein model, a replicating portfolio is computed, and the fair price of a derivative is obtained. The theory files described in this paper are available on the *Archive of Formal Proofs*, at <https://www.isa-afp.org/entries/DiscretePricing.html>.

2 Preliminary notions

2.1 Some notions in finance

We begin by briefly reviewing some basic standard definitions about equity markets. This treatment is mainly based on Shreve [21], Vol. 1. An equity market consists of a set of *assets* or *securities* that can be traded at prices that evolve with time. An actor trading on different assets will own a *portfolio* containing different quantities of the traded assets. These quantities are real numbers that can be positive if the corresponding asset was bought, or negative if the asset was the object of a so-called short sale (the asset was sold by an actor not actually owning it). A portfolio can be *static* if its composition is fixed once and for all, and *dynamic* if its composition can evolve over time. Clearly, most portfolios on markets are dynamic. Among the dynamic portfolios, those of a particular interest are the *trading strategies*; these are the dynamic portfolios for which the composition at time t is a random variable that only depends on the available information up to time t . In other words, a trading strategy is one for which investments cannot depend on information that has not occurred yet. From a financial point of view, trading strategies are thus meant to represent portfolios for which no insider trading can occur. A portfolio in which cash is only invested at inception, after which all future trades are financed by buying or selling assets in the portfolio is a *self-financing portfolio*. An arbitrage represents a “free lunch”: it is defined as a self-financing trading strategy with a 0 initial investment that offers a risk-free possibility of making a profit. A market is *viable* if it offers no arbitrage opportunities.

Some of the securities that can be traded are basic securities that are generally categorized depending to their risk level. Intuitively, the more the price of an asset fluctuates, the riskier the asset, since it is more difficult to predict the return an investor would obtain by buying the asset. A class of securities with no risk is the class of *bonds*. These are assets based on a debt,

such as Treasury bills, that are assumed to guarantee a given return to an investor¹, and are thus considered as *risk-free assets*.

Other securities that are traded are *derivative* securities. These are securities that are characterized by their expiry date or maturity, and *exercise times*, which are the times when cash is exchanged between the buyer and seller of the derivatives. The amount of cash to be exchanged is called the *payoff* of the derivative product, and it depends on the evolution and values of some underlying securities. In this paper we will focus on *European* derivatives, which can only be exercised at the maturity, see, e.g., [14].

An *option* is an example of a derivative that can be viewed as an insurance: when it is exercised, it gives its owner the right—but not the obligation—to trade an asset at a given price. The best-known options are the call and the put options. A call (resp. put) option gives its owner the right, at time T , to buy (resp. sell) the underlying security at the so-called *strike price* K , thus guaranteeing that there is a cap (resp. floor) on the price that will be paid (resp. received) at a future time for the security. In practice, when at time T the price of the underlying security, denoted by S_T , is greater than the strike price K , the buyer of a call receives $S_T - K$ from the seller of the option, and buys the security on the market for S_T , in effect only spending K to obtain it. When $S_T < K$, the seller of the call does not deliver any cash, as the buyer will directly buy the security on the market for a value that is less than K . Thus, a call option is a derivative that, at maturity T , delivers a payoff of $(S_T - K)^+ \stackrel{\text{def}}{=} \max(0, S_T - K)$. In a similar way, a put option delivers at time T a payoff of $(K - S_T)^+$.

Once a derivative is sold, the seller is meant to invest the cash by creating a trading strategy, in order to be able to pay the required amount of money when the derivative is exercised. The question is then: how much should a buyer be expected to pay for a given derivative?

- No seller will be willing to sell it at a price so low that there is a risk when the option expires the seller will lose money paying what is owed: such a price would be unfair to the seller.
- If this price is so high that it is clear the seller can always pay what is owed and sometimes even make a profit, the buyer will search for another actor selling the derivative at a cheaper price. Such a price is unfair to the buyer.

What it means precisely for a price to be too high or too low will be made clear later; a price that is neither too high nor too low is a *fair price*. As we will see, under some quite natural hypotheses, there is a case where the answer to this question is straightforward. This is when the seller is capable of creating a trading strategy that generates at exercise time *exactly* the payoff of the derivative. Such a trading strategy is called a *replicating portfolio*. In this case, the fair price for the derivative is the investment needed to initiate the

¹ It can be argued that this assumption is incorrect because that there is always a nonzero probability that investors will not be paid what they are owed. But because these bills are backed by national governments, this probability is very close to 0.

trading strategy. A market is *complete* if every derivative admits a replicating portfolio; in a complete market, every derivative admits a unique fair price.

The construction of replicating portfolios is clearly not straightforward, and it may not be guaranteed that such portfolios actually exist. An answer to the existence of replicating portfolios for European options was given by Fischer Black and Myron Scholes, and by Robert Merton in [4, 17], in the so-called Black-Scholes-Merton model. In their model, the equity market consists of two assets: A risky asset, the stock, that pays no dividends and whose evolution is described by a *geometric Brownian motion* (see, e.g., [14]); and a risk-free asset with a deterministic price evolution. Their main result is that, under some simple market hypotheses such as identical bidding and asking prices and the absence of arbitrage opportunities, a European option over a single stock can be replicated with a portfolio consisting of the stock and a cash account. Their proof is based on the construction of a dynamic portfolio, the composition of which changes continuously (it is called a *delta-neutral portfolio*), which is guaranteed to replicate the option under consideration. Along with the construction of replicating portfolios, the authors provide a formula that permits to compute the fair price of any European option.

The Cox-Ross-Rubinstein model [6] that we consider in Section 6 of this paper can be considered as an approximation of the Black-Scholes-Merton model to the case where time is no longer continuous but discrete; i.e., to the case where securities are only traded at discrete times $1, 2, \dots, n, \dots$. In this setting, the evolution of the stock price is described by a geometric random walk, which can be viewed as a discrete version of the geometric Brownian motion: if the stock has a price s at time n , then at time $n + 1$, this price is either $u \cdot s$ (upward movement) or $d \cdot s$ (downward movement). The factors u and d must satisfy the relation: $d < 1 + r < u$, where r is the rate of the risk-free asset (meaning that the price of the risky asset can either move upward or downward relatively to the risk-free asset). The probability of the price going up is always $0 < p < 1$, and the probability of it going down is $1 - p$. The authors show that under these conditions, the market is complete: every derivative admits a replicating portfolio.

2.2 Probability theory in Isabelle: existing notions

We briefly present the syntax of the interactive theorem prover Isabelle/HOL; the tool can be downloaded at <https://isabelle.in.tum.de/>, along with tutorials and documentations. Additional material on Isabelle can be found in [19]. This prover is based on higher-order logic; terms are built using types that can be:

- simple types, denoted with the Greek letters α, β, \dots
- types obtained from type constructors, represented in postfix notation (e.g. the type α `set` denotes the type of sets containing elements of type α), or in infix notation (e.g., the type $\alpha \rightarrow \beta$ denotes the type of total functions from α to β).

Functions are curried, and function application is written without parentheses. Anonymous functions are represented with the lambda notation: the function $x \mapsto t$ is denoted by $\lambda x. t$. We will use mathematical notations for standard terms; for example, the set of reals will be denoted by \mathbb{R} .

A large part of the formalization of measure and probability theory in Isabelle was carried out by Hölzl [11] and is now included in Isabelle's distribution. We briefly recap some of the notions that will be used throughout the paper and the way they are formalized in Isabelle. We assume the reader has knowledge of fundamental concepts of measure and probability theory; any missing notions can be found in [7] for example. For the sake of readability, in what follows, a term $F t$ will sometimes be written F_t .

Probability spaces are particular *measure spaces*. A measure space over a set Ω consists of a function μ that associates a nonnegative number or $+\infty$ to some subsets of Ω . The subsets of Ω that can be measured are closed under complement and countable unions and make up a σ -algebra. The σ -algebra generated by a set $C \subseteq 2^\Omega$ is the smallest σ -algebra containing C ; it is denoted in Isabelle by `sigma-sets` ΩC .

The functions μ that measure the elements of a σ -algebra are positive and *sigma additive*²: if $\mathcal{A} \subseteq 2^\Omega$ is a σ -algebra and the sequence $(A_i)_{i \in \mathbb{N}}$ consists of pairwise disjoint elements in \mathcal{A} , then $\mu(\bigcup_{i \in \mathbb{N}} A_i) = \sum_{i \in \mathbb{N}} \mu(A_i)$. In Isabelle, measure spaces are defined as follows (where $\overline{\mathbb{R}}$ denotes $\mathbb{R} \cup \{-\infty, +\infty\}$ and $\mathbb{B} \stackrel{\text{def}}{=} \{\perp, \top\}$):

```

measure-space      ::  $\alpha$  set  $\rightarrow$   $\alpha$  set set  $\rightarrow$  ( $\alpha$  set  $\rightarrow$   $\overline{\mathbb{R}}$ )  $\rightarrow$   $\mathbb{B}$ 
measure-space  $\Omega$   $\mathcal{A}$   $\mu$   $\Leftrightarrow$ 
   $\sigma$ -algebra  $\Omega$   $\mathcal{A}$   $\wedge$  positive  $\mathcal{A}$   $\mu$   $\wedge$  countably-additive  $\mathcal{A}$   $\mu$ 

```

A measure type is defined by fixing the measure of non-measurable sets to 0:

```

typedef  $\alpha$  measure =  $\{(\Omega, \mathcal{A}, \mu) \mid (\forall A \notin \mathcal{A}. \mu A = 0) \wedge \text{measure-space } \Omega \mathcal{A} \mu\}$ 

```

If \mathcal{M} is an element of type α `measure`, then the corresponding space, σ -algebra and measure are respectively denoted by $\Omega_{\mathcal{M}}$, $\mathcal{A}_{\mathcal{M}}$ and $\mu_{\mathcal{M}}$.

The definition of a measure type may seem surprising, especially to mathematicians, because setting the measure of a set not in \mathcal{A} to 0 can be counter-intuitive: there is for example no relationship between elements with a measure 0 and negligible elements on a measure space. The reason for this is that in Isabelle, a function cannot be partial and it is necessary to define the measure function on every subset of $\Omega_{\mathcal{M}}$; the choice of setting these values to 0 may seem arbitrary, but it does not entail any contradiction. In a similar way, in Isabelle, division is extended to 0 by letting $x/0 = 0$. This does not entail any contradiction and any value could have been chosen, but setting this value to 0 permits to have other basic mathematical theorems to hold unconditionally.

We can associate to any σ -algebra $\mathcal{C} \subseteq 2^\Omega$ a measure space with a uniformly null measure, $(\Omega, \mathcal{C}, (\lambda x. 0))$. In Isabelle, this measure space is denoted by `sigma` $\Omega \mathcal{C}$.

² This property is also called *countable additivity* in the literature.

A function between two measurable spaces is *measurable* if the preimage of every measurable set is measurable. In Isabelle, sets of measurable functions are defined as follow:

```
measurable      ::  $\alpha$  measure  $\rightarrow$   $\beta$  measure  $\rightarrow$  ( $\alpha \rightarrow \beta$ ) set
measurable  $\mathcal{M} \mathcal{N} \mu = \{f : \Omega_{\mathcal{M}} \rightarrow \Omega_{\mathcal{N}} \mid \forall A \in \mathcal{A}_{\mathcal{N}}. f^{-1}(A) \cap \Omega_{\mathcal{M}} \in \mathcal{A}_{\mathcal{M}}\}$ 
```

In particular, for a function $f : \Omega_{\mathcal{M}} \rightarrow \Omega_{\mathcal{N}}$, we will consider the smallest measure space in which f is measurable. This measure space is denoted by $\mathcal{M}_{\langle f \rangle}$ and defined by $\mathcal{M}_{\langle f \rangle} \stackrel{\text{def}}{=} (\Omega_{\mathcal{M}}, \mathcal{B}, \mu_{\mathcal{M}})$, where \mathcal{B} is the σ -algebra generated by the set $\{f^{-1}(A) \cap \Omega_{\mathcal{M}} \mid A \in \mathcal{A}_{\mathcal{N}}\}$.

Probability measures are measure spaces on which the measure of Ω is finite and equal to 1. In Isabelle, they are defined in a *locale* [2]; this allows one to delimit a range in which the existence of a measure satisfying the desired assumptions is assumed, instead of having to explicitly add the corresponding hypotheses in every theorem, which would be tedious.

```
locale prob-space = finite-measure + assumes  $\mu_{\mathcal{M}}(\Omega_{\mathcal{M}}) = 1$ 
```

A *random variable* on a probability space \mathcal{M} is a measurable function with domain $\Omega_{\mathcal{M}}$. The average value of a random variable f is called its *expectation*, it is denoted by³ $\mathbb{E}^{\mathcal{M}}[f]$, and defined by $\mathbb{E}^{\mathcal{M}}[f] \stackrel{\text{def}}{=} \int_{\Omega_{\mathcal{M}}} f d\mu_{\mathcal{M}}$. Collections of random variables are called *stochastic processes*. In most cases, stochastic processes are indexed by a totally ordered set, representing time, such as \mathbb{N} or \mathbb{R}^+ .

In what follows, we will consider properties that hold *almost surely* (or *almost everywhere*), i.e., are such that the elements for which they do not hold reside within a set of measure 0:

lemma AE-IFF :

$$(\text{AE}_{\mathcal{M}} x. P x) \Leftrightarrow (\exists N \in \mathcal{A}_{\mathcal{M}}. \mu_{\mathcal{M}}(N) = 0 \wedge \{x \mid \neg P x\} \subseteq N)$$

Given measure spaces \mathcal{M} and \mathcal{N} , we say that \mathcal{N} is a *subalgebra* of \mathcal{M} if $\Omega_{\mathcal{M}} = \Omega_{\mathcal{N}}$ and $\mathcal{A}_{\mathcal{N}} \subseteq \mathcal{A}_{\mathcal{M}}$.

3 Modeling equity markets in discrete time

3.1 General definitions

An equity market is characterized by the set of assets that can be traded and the price at which they are traded⁴. A subset of these assets represents the

³ The superscript may be omitted if there is no confusion.

⁴ This is a simplification as in practice, two prices are associated with each asset: a *bid price*, which represents the price traders are willing to pay to buy the asset, and an *ask price*, which represents the price traders are willing to sell the asset. Bid prices are always lower than ask prices, but on markets on which high volumes of assets are traded, both prices are typically very close.

basic securities that can be traded, these are the *stocks*. Examples of stocks are shares on companies like Google, Apple, Facebook or Amazon, which can be traded on the stock market. The remaining assets are viewed as *derivative products*, the value of which typically depends on that of some stocks. Examples of derivative products are *futures* on Facebook, which are contracts where two actors agree to trade a share of Facebook for a given price at a given time; or *basket options* on Apple and Google, which are options that give the buyer the right, but not the obligation, to buy shares of Apple *and* Google for a given price at a given time⁵. The precise definition of these derivative products is not important at this point, these are assets with a value depending on that of one or several stocks. The price at which an asset can be traded at each time is a random variable, this price is thus represented by a stochastic process; and in this case for which time is discrete, these stochastic processes are indexed by \mathbb{N} . At time n , the random variable associated with an asset thus represents the value of this asset on time interval $[n, n + 1[$. Note that in this general setting, there is no relationship between the price processes of derivative products and that of stocks. As we are concerned with computing fair prices for the former, equity markets are defined in such a way that there always exists at least one derivative product.

```
stk-strict-subs  ::  $\beta$  set  $\rightarrow \mathbb{B}$ 
stk-strict-subs  $S \Leftrightarrow S \neq \text{UNIV}$ 
```

Equity markets are defined as a type, from which the stocks S and price processes P can be obtained:

```
typedef ( $\alpha, \beta$ ) discrete-market =
  {( $S :: \beta$  set,  $P :: \beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}$ ) | stk-strict-subs  $S$ }
```

The type β represents the products that can be traded and the type α represents the random outcome. In Isabelle, the `typedef` command permits to obtain two morphisms for handling the `discrete-market` type: An *abstraction morphism*

```
Abs-discrete-market :: ( $\beta$  set  $\times (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R})$ )  $\rightarrow$ 
  ( $\alpha, \beta$ ) discrete-market
```

and a *representation morphism*

```
Rep-discrete-market :: ( $\alpha, \beta$ ) discrete-market  $\rightarrow$ 
  ( $\beta$  set  $\times (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R})$ )
```

We may thus use the representation morphism to retrieve the set of stocks and the price process from a given market:

```
stocks      :: ( $\alpha, \beta$ ) discrete-market  $\rightarrow \beta$  set
stocks Mkt = fst (Rep-discrete-market Mkt)
```

⁵ Note that buying a basket option on Apple and Google is *not* the same as buying a call on Apple and another one on Google.

`prices` :: (α, β) `discrete-market` $\rightarrow \beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}$
`prices Mkt = snd (Rep-discrete-market Mkt)`

We next consider *quantity processes*. These are used to represent the fact that assets can be bought and sold; in particular, it is possible on financial markets to sell an asset that is not held: when this occurs, we say the seller is *short on the asset* (or *holds a short position on the asset*) and owns a negative amount of the asset. When the holder owns a positive amount of the asset, we say the holder is *long on the asset* (or *holds a long position on the asset*). In several textbooks on financial mathematics, quantity processes and portfolios are represented using vectors. Such a representation permits a compact notation, and using notions such as scalar product, allows one to represent financial notions such as value processes concisely. We chose to formalize quantity processes differently to avoid importing theories that are not necessary, and fixing the set of assets used to construct quantity processes as well as their order.

We assume that any portion of an asset may be traded, thus the quantity withheld is a real number. Quantity processes are formalized as functions that associate a stochastic process to each asset. By convention, for $n > 0$, if q is a quantity process and a is an asset, then $q \ a \ n \ w$ represents the quantity (positive if we are long the asset and negative if we are short the asset) of asset a withheld on the time interval $]n - 1, n]$ for scenario w . The value of a quantity process at time 0 is thus unimportant. Intuitively, the reason for such a convention is that, at time n , a quantity process is meant to only depend on the information available up to time $n - 1$. More formally, in both discrete and continuous-time models, quantity processes of interest will be required to be *predictable processes*, and the convention on quantity processes allows for a uniform presentation. We define the following operators that permit to construct and combine quantity processes⁶.

– No components

`qty-empty` :: $\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}$
`qty-empty = ($\lambda x \ n \ w. 0$)`

– Single component

`qty-single` :: $\beta \rightarrow (\mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}$
`qty-single asset qt-proc = qty-empty(asset := qt-proc)`

– Sum quantities

`qty-sum` :: $(\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow$
 $\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}$
`qty-sum $q_1 \ q_2 = (\lambda x \ n \ w. (q_1 \ x \ n \ w) + (q_2 \ x \ n \ w))$`

⁶ For the definition of `qty-single`, we use the notation $f(a := b)$, which in Isabelle represents an update of function f so that the image of a becomes b .

Table 1 Table of quantity values for Example 1

Time	1	2	3	4
Apl quantity	1	2	3	4
Goog quantity	-1	-2	-3	-4

– Multiply quantities

$$\begin{aligned} \text{qty-mult-comp} &:: (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow (\mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \\ &\quad \beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R} \\ \text{qty-mult-comp } q \text{ prd} &= (\lambda x n w. (q x n w) \cdot (\text{prd } n w)) \end{aligned}$$

– Remove component

$$\begin{aligned} \text{qty-rem-comp} &:: (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow (\mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \\ &\quad \beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R} \\ \text{qty-rem-comp } q \text{ asset} &= q(\text{asset} := (\lambda n w. 0)) \end{aligned}$$

Intuitively, `qty-empty` represents the quantity process in which no asset is bought or sold, and `qty-single` is the process for which a single asset is potentially bought or sold. The other operators permit to respectively add quantity processes, to multiply all of them by another process, and to nullify the quantity of an asset.

Related to the notion of a quantity process is that of its support set, which consists of all the assets that are potentially bought or sold at some point for some scenario. This leads to the definition of portfolios, which are quantity processes that admit a finite support set.

$$\begin{aligned} \text{support-set} &:: (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \beta \text{ set} \\ \text{support-set } q &= \{a \mid \exists n w. q a n w \neq 0\} \\ \text{portfolio} &:: (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{B} \\ \text{portfolio } p &\Leftrightarrow \text{finite } (\text{support-set } p) \end{aligned}$$

In particular, stock portfolios are portfolios for which the support set consists only of stocks.

$$\begin{aligned} \text{stock-portfolio} &:: (\alpha, \beta) \text{ discrete-market} \rightarrow (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{B} \\ \text{stock-portfolio Mkt } p &\Leftrightarrow \text{portfolio } p \wedge \text{support-set } p \subseteq \text{stocks Mkt} \end{aligned}$$

Example 1 Consider a market `Mkt` with stocks including shares on Apple, Facebook and Google: $\{\text{Apl}, \text{Fbk}, \text{Goog}\} \subseteq \text{stocks Mkt}$. We can construct the following portfolio

$$p_1 \stackrel{\text{def}}{=} \text{qty-sum } (\text{qty-single Apl } (\lambda n w. n)) (\text{qty-single Goog } (\lambda n w. -n)).$$

The summary of the quantities for $t = 1, \dots, 4$ is given in Table 1. This is portfolio in which we are long n shares of Apple and short n shares of Google until time n for all scenarios; it has a support set consisting of Apple and Google and is thus a stock portfolio.

We now define *value processes* and *closing value process* for portfolios. Intuitively, the value process of a portfolio at time n represents the total amount of cash that is necessary at time n to invest in the assets of the portfolio until time $n + 1$, and the closing value process of a portfolio at time n represents the total amount of cash received/owed when closing out all positions⁷ at time n . The closing value process of a portfolio at time 0 can be defined arbitrarily; a standard practice consists in setting its value to that of the value process of the portfolio at time 0. Note that if the composition of the portfolio does not change between times $]n - 1, n]$ and $]n, n + 1]$, then the value of the closing value process at time n is the same as that of the value process.

– Value process

```

val-process      :: (α, β) discrete-market → (β → ℕ → α → ℝ) →
                  ℕ → α → ℝ
val-process Mkt p = if ¬(portfolio p) then (λn w. 0) else
  (λn w. ∑a∈support-set p ((prices Mkt) a n w) * (p a (n + 1) w))

```

– Intermediate function for the closing value process

```

tmp-cl-val      :: (α, β) discrete-market →
                  (β → ℕ → α → ℝ) → ℕ → α → ℝ
tmp-cl-val Mkt p 0 = val-process Mkt p 0
tmp-cl-val Mkt p (n + 1) =
  (λw. ∑a∈support-set p ((prices Mkt) a (n + 1) w) * (p a (n + 1) w))

```

– Closing value process

```

cls-val-process :: (α, β) discrete-market →
                  (β → ℕ → α → ℝ) → ℕ → α → ℝ
cls-val-process Mkt p = if ¬(portfolio p) then (λn w. 0) else
  (λn w. tmp-cl-val Mkt p n w)

```

Example 2 Assume the Apple and Google shares have deterministic prices given in Table 2. Then the value process and closing value process of portfolio p_1 defined in Example 1 are given in the same table.

Example 3 Still under the assumption that Apple and Google shares have deterministic prices recalled in Table 3, assume p'_1 is a static portfolio (i.e., one for which the quantity processes are constant) in which we are always long 2 shares of Apple and short 2 shares of Google. Then the value process and closing process of p'_1 is given in the same table.

⁷ Closing out all positions means getting rid of all the assets in a portfolio, i.e., selling those with a long position, and buying back those with a short position.

Table 2 Quantities and portfolio values for Example 2

Time	0	1	2	3
Apl quantity	-	1	2	3
Goog quantity	-	-1	-2	-3
Apl value	100	98	96	98
Goog value	90	92	98	95.5
val-process Mkt p_1	10	12	-6	10
cls-val-process Mkt p_1	10	6	-4	7.5

Table 3 Quantities and portfolio values for Example 3

Time	0	1	2	3
Apl quantity	-	2	2	2
Goog quantity	-	-2	-2	-2
Apl value	100	98	96	98
Goog value	90	92	98	95.5
val-process Mkt p'_1	20	12	-4	5
cls-val-process Mkt p'_1	20	12	-4	5

Table 4 Deterministic price of Facebook share for Example 4

Time	0	1	2	3
Fbk	5	4	4	5

Self-financing portfolios are portfolios in which no cash is invested except possibly at inception. A portfolio is self-financing if its closing value and value at time $n + 1$ are identical; this means that the value of the portfolio may be affected by the evolution of the market but not by the changes in its composition.

$$\begin{aligned} \text{self-financing} &:: (\alpha, \beta) \text{ discrete-market} \rightarrow \\ &(\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{B} \\ \text{self-financing Mkt } p &\Leftrightarrow \\ &\forall n. \text{val-process Mkt } p (n + 1) = \text{cls-val-process Mkt } p (n + 1) \end{aligned}$$

A simple example of a self-financing portfolio is a static portfolio. Since its composition never changes, no cash is ever invested or withdrawn from it. A self-financing portfolio with initial value v_0 can be obtained starting from an arbitrary portfolio, provided the market contains an asset that never admits a price equal to 0, by buying (resp. selling) the required quantity of the asset with the extra (resp. missing) cash.

Example 4 Portfolio p_1 of Example 1 is not self-financing. Assume the stock price of Facebook is deterministic and given in Table 4. Then we can construct

Table 5 Self-financing portfolio for Example 4

Time	0	1	2	3	4
Apl quantity	-	1	2	3	4
Goog quantity	-	-1	-2	-3	-4
Fbk quantity	-	-2	-3.5	-3	-3.5
Apl value	100	98	96	98	-
Goog value	90	92	98	95.5	-
Fbk value	5	4	4	5	-
val-process Mkt p_2	0	-2	-18	-7.5	-
cls-val-process Mkt p_2	0	-2	-18	-7.5	-

a self-financing portfolio p_2 with initial value 0, that has the same quantity processes as p_1 for Apple and Google. The quantities of stocks and the value and closing value processes of p_2 are given in Table 5. For instance, at time 0, the holder buys a share of Apple for 100€ and sells a (borrowed) share of Google for 90€, creating a portfolio for a total cost of 10€. To make the portfolio self-financed with initial value 0, this cost is compensated by selling 2 (borrowed) shares of Facebook at 5€ each, the total cost of the created portfolio is then 0€.

3.2 Modeling time-dependent information

Filtrations are used to represent information accumulated over time. Formally, they are defined as a collection of increasing subalgebras over a totally ordered set with a minimal element \perp –typically \mathbb{N} or \mathbb{R}^+ .

```
class linorder-bot = linorder + bot
filtration      ::  $\alpha$ measure  $\rightarrow ((\iota :: \text{linorder-bot}) \rightarrow \alpha \text{measure}) \rightarrow \mathbb{B}$ 
filtration  $\mathcal{M} \mathcal{F} \Leftrightarrow (\forall t. \text{subalgebra } \mathcal{M} \mathcal{F}_t) \wedge$ 
                     $(\forall s t. s \leq t \Rightarrow \text{subalgebra } \mathcal{F}_t \mathcal{F}_s)$ 
```

In general, when a filtration \mathcal{F} representing available information is provided, we will mainly be interested in stochastic processes that depend on this information. There are two categories of such stochastic processes of interest for our purpose: *adapted stochastic processes*, that at time n are \mathcal{F}_n -measurable; and *predictable stochastic processes*, that at time $n > 0$ are \mathcal{F}_{n-1} -measurable. The definition of adapted stochastic processes in the more general case is a straightforward generalization of that in the discrete case, which is the one that is formalized below. We also introduce abbreviations for stochastic processes with a range in a Borel measure space.

– Adapted stochastic processes

```
adapt-sp      ::  $(\iota \rightarrow \alpha \text{measure}) \rightarrow (\iota \rightarrow \alpha \rightarrow \beta) \rightarrow$ 
                 $\beta \text{measure} \rightarrow \mathbb{B}$ 
adapt-sp  $\mathcal{F} X \mathcal{N} \Leftrightarrow \forall t. X_t \in \text{measurable } \mathcal{F}_t \mathcal{N}$ 
```

abbreviation borel-adapt-sp $\mathcal{F} X \equiv \text{adapt-sp } \mathcal{F} X \text{ borel}$

– Predictable stochastic processes

predict-sp $:: (\mathbb{N} \rightarrow \alpha \text{ measure}) \rightarrow (\mathbb{N} \rightarrow \alpha \rightarrow \beta) \rightarrow$
 $\beta \text{ measure} \rightarrow \mathbb{B}$
predict-sp $\mathcal{F} X \mathcal{N} \Leftrightarrow X_0 \in \text{measurable } \mathcal{F}_0 \mathcal{N} \wedge$
 $\forall n. X_{n+1} \in \text{measurable } \mathcal{F}_n \mathcal{N}$

abbreviation borel-predict-sp $\mathcal{F} X \equiv \text{predict-sp } \mathcal{F} X \text{ borel}$

In our context, filtrations are meant to represent the currently available information. A standard filtration used in financial mathematics is the one defined as follows: for all $n \geq 0$, \mathcal{F}_n is the smallest subalgebra of \mathcal{M} in which for any stock s and time $k \leq n$, the price process (**prices Mkt**) $s k$ is Borel-measurable. It is straightforward to verify that \mathcal{F} is indeed a filtration. In particular, at time 0, there is no information available, thus the measure space $\mathcal{F}_0 = \mathcal{F}_\perp$ is trivial. Filtrations satisfying such a requirement are called *initially trivial filtrations*.

init-triv-filt $:: \alpha \text{ measure} \rightarrow (\iota \rightarrow \alpha \text{ measure}) \rightarrow \mathbb{B}$
init-triv-filt $\mathcal{M} \mathcal{F} \Leftrightarrow \text{filtration } \mathcal{M} \mathcal{F} \wedge \text{sets } \mathcal{F}_\perp = \{\emptyset, \Omega_{\mathcal{M}}\}$

We define a locale for discrete equity markets by fixing a market and considering a probability space equipped with an arbitrary filtration that is initially trivial.

locale init-triv-prob-space = prob-space +
fixes $\mathcal{F} :: \mathbb{N} \rightarrow (\alpha \text{ measure})$
assumes init-triv-filt \mathcal{F}
locale disc-equity-market = init-triv-prob-space +
fixes $\text{Mkt} :: (\alpha, \beta) \text{ discrete-market}$

Most of the assets that will be considered in this locale are those that have an adapted price process w.r.t. the given filtration. Quantity processes in which only assets with an adapted price process are bought or sold are called *support-adapted* quantity processes.

support-adapt $:: (\alpha, \beta) \text{ discrete-market} \rightarrow$
 $(\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{B}$
support-adapt $\text{Mkt qt-proc} \Leftrightarrow \forall a \in \text{support-set qt-proc.}$
borel-adapt-sp $\mathcal{F} (\text{prices Mkt } a)$

The only portfolios that it is reasonable to consider are those with a composition that depends only on the information available at the current time. More precisely, these are the portfolios for which the amounts that are bought or sold of each asset on the time interval $]n, n + 1]$ is known at time n . This means that the quantity of each asset is a predictable process; such portfolios are called *trading strategies*.

```

trading-strat  :: ( $\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}$ )  $\rightarrow \mathbb{B}$ 
trading-strat p  $\Leftrightarrow$  portfolio p  $\wedge$ 
                    ( $\forall a \in$  support-set p. borel-predict-sp  $\mathcal{F}$  (p a))

```

In particular, the value process of a support-adapted trading strategy is itself an adapted process:

```

lemma TRADING-STRATEGY-ADAPTED
  assumes trading-strat p
  and support-adapt Mkt p
  shows borel-adapt-sp  $\mathcal{F}$  (val-process Mkt p)

```

Since the filtration \mathcal{F} is assumed to be initially trivial, such a strategy necessarily admits a constant value at inception:

```

lemma TRADING-STRATEGY-INIT
  assumes trading-strat p
  and support-adapt Mkt p
  shows  $\exists c. \forall w \in \Omega_{\mathcal{M}}. \text{val-process Mkt } p \ 0 \ w = c$ 

```

We denote by `init-value p` the constant value equal to the value process of a support-adapted trading strategy at time 0.

4 The notion of a fair price

4.1 Definitions

We define the notion of a fair price, which is meant to represent the price at which a derivative available on the market should be bought or sold. Intuitively, a fair price for an asset is one that does not allow a buyer or seller of the asset to make a risk-free profit with this transaction. Making a risk-free profit is called an *arbitrage*. We begin by formally defining this notion. An arbitrage is a self-financing trading strategy with a zero initial value, that at some point in time is almost surely positive and with a strictly positive probability of making a gain. Arbitrage opportunities exist in real financial markets, however they tend to be quickly exploited, which makes them vanish. In fact, there is an entire category of traders on markets with the goal of detecting and exploiting arbitrages as quickly as possible. Arbitrage opportunities are thus viewed by researchers in financial mathematics as glitches in the financial systems that appear and disappear almost immediately. In our context, the no-arbitrage assumption is used as a theoretical tool to define the notion of a fair price; intuitively it ensures that products that are essentially equivalent have the same price. Almost all pricing results in financial mathematics are based on a no-arbitrage assumption.

arbitrage-process $:: (\alpha, \beta) \text{ discrete-market} \rightarrow$
 $(\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{B}$

arbitrage-process Mkt $p \Leftrightarrow$
 $(\exists m \in \mathbb{N}.$
 $(\text{trading-strat } p) \wedge (\text{self-financing } p) \wedge$
 $(\forall w \in \Omega_{\mathcal{M}}. (\text{val-process } p) 0 w = 0) \wedge$
 $(\text{AE}_{\mathcal{M}} w. (\text{cls-val-process } p) m w \geq 0) \wedge$
 $(\mathcal{P}(\{w \in \Omega_{\mathcal{M}} \mid (\text{cls-val-process } p) m w > 0\}) > 0))$

Next we define the notion of a *price structure* for a derivative. Derivatives are characterized by their maturity (i.e., their expiry date) and the payoff they deliver at maturity; a price structure is a stochastic process with a constant initial value that coincides with the payoff of the derivative almost everywhere at maturity. The initial value of a price structure will represent the price of the derivative under consideration.

price-struct $:: (\alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{N} \rightarrow \mathbb{R} \rightarrow (\mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{B}$
price-struct $\kappa T \pi \text{ pr} \Leftrightarrow (\forall w \in \Omega_{\mathcal{M}}. \text{pr } 0 w = \pi) \wedge$
 $(\text{AE}_{\mathcal{M}} w. \text{pr } T w = \kappa w) \wedge$
 $(\text{pr } T \in \text{borel-measurable } \mathcal{F}_T)$

In order to formalize the notion of a fair price for a derivative, we need to formalize the fact that buying or selling the derivative at price π does not lead to any arbitrage opportunity. More precisely, it should not be possible to obtain an arbitrage process using only stocks from the market and an asset with a price process identical to a price structure of the derivative, with an initial value π . In order to guarantee the existence of such an asset, we define the notion of coincidence between two markets.

coincides $:: (\alpha, \beta) \text{ discrete-market} \rightarrow$
 $(\alpha, \beta) \text{ discrete-market} \rightarrow \beta \text{ set} \rightarrow \mathbb{B}$
coincides Mkt Mkt' $A \Leftrightarrow \text{stocks Mkt} = \text{stocks Mkt}' \wedge$
 $\forall a. a \in A \Rightarrow \text{prices Mkt } a = \text{prices Mkt}' a$

fair-price $:: (\alpha, \beta) \text{ discrete-market} \rightarrow \mathbb{R} \rightarrow (\alpha \rightarrow \mathbb{R}) \rightarrow$
 $\mathbb{N} \rightarrow \mathbb{B}$

fair-price Mkt $\pi \kappa T \Leftrightarrow (\exists \text{pr. price-struct } \kappa T \pi \text{ pr} \wedge$
 $(\forall a \text{ Mkt}' p. a \notin \text{stocks Mkt} \Rightarrow$
 $(\text{coincides Mkt Mkt}') \wedge$
 $(\text{prices Mkt}' a = \text{pr}) \wedge$
 $(\text{portfolio } p) \wedge$
 $(\text{support-set } p \subseteq \text{stocks Mkt} \cup \{a\}) \Rightarrow$
 $\neg \text{arbitrage-process Mkt}' p))$

4.2 Replicating portfolios

We prove the central result that, for a market satisfying a so-called *viability* condition defined below, under the hypothesis that a *replicating portfolio* exists

for a given derivative, the latter admits a fair price that is unique. A replicating portfolio for a given derivative is a self-financing trading strategy that consists of stocks only, and that at maturity, has a value identical to the payoff of the derivative almost everywhere. If such a portfolio exists, then the derivative is *attainable*, and if every derivative available on a market is attainable, then the market is *complete*:

$$\begin{aligned} \text{replic-pf} &:: (\beta \rightarrow \mathbb{N} \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow (\alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{N} \rightarrow \mathbb{B} \\ \text{replic-pf } p \kappa T &\Leftrightarrow (\text{stock-portfolio Mkt } p) \wedge \\ &\quad (\text{self-financing } p) \wedge (\text{trading-strat } p) \wedge \\ &\quad (\text{AE}_{\mathcal{M}} w. \text{cls-val-process Mkt } p T w = \kappa w) \end{aligned}$$

$$\begin{aligned} \text{attainable} &:: (\alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{N} \rightarrow \mathbb{B} \\ \text{attainable } \kappa T &\Leftrightarrow (\exists p. \text{replic-pf } p \kappa T) \end{aligned}$$

$$\begin{aligned} \text{complete-market} &:: \mathbb{B} \\ \text{complete-market} &\Leftrightarrow \forall T. \forall \kappa \in (\text{borel-measurable } \mathcal{F}_T). \text{attainable } \kappa T \end{aligned}$$

The existence of a replicating portfolio by itself is not sufficient to guarantee the existence of a fair price: indeed, if for example it is already possible to construct an arbitrage process on the market using only stocks, then there clearly cannot be any fair price for any derivative product. It is thus necessary to forbid arbitrage opportunities using only stocks from the market. This is captured by the notion of a *viable market*.

$$\begin{aligned} \text{viable-market} &:: (\alpha, \beta) \text{discrete-market} \rightarrow \mathbb{B} \\ \text{viable-market Mkt} &\Leftrightarrow \forall p. \text{stock-portfolio } p \Rightarrow \\ &\quad \neg \text{arbitrage-process Mkt } p \end{aligned}$$

We first show that, if the market is viable, then every derivative admitting a replicating portfolio has a fair price that is the initial value of the replicating portfolio.

lemma REPLICATING-FAIR-PRICE
assumes viable-market Mkt
and replic-pf $p \kappa T$
and support-adapt Mkt p
shows fair-price Mkt (init-value p) κT

The proof of this result goes as follows. If p is a replicating portfolio for a derivative κ , then clearly the closing value process of p is a price process for κ . Assume the initial value of this process, π , is not a fair price. Then there must exist an arbitrage constructed using only stocks and an asset x whose price coincides with the closing value of p . But then it is possible to construct another arbitrage using only stocks, by replacing the quantities of x by identical quantities of the replicating portfolio. This is an arbitrage consisting of stocks only, and we obtain a contradiction since the market is assumed to be viable.

We also provide a proof of the uniqueness of a fair price for attainable derivatives based on the existence of a stock on the market with a strictly positive price process. The proof could also be carried out assuming the existence of a stock on the market with a strictly negative price process, but that does not really make sense from a financial point of view: the owner of such a share would receive a negative amount of money by selling the share, i.e., would be paying to be rid of the share. We also assume that the price processes of all stocks in the market are adapted to the filtration under consideration.

```

locale disc-mkt-pos-stock = disc-equity-market +
  fixes pos-stock ::  $\beta$ 
  assumes pos-stock  $\in$  stocks Mkt
  and  $\forall n w$ . prices Mkt pos-stock  $n w > 0$ 
  and  $\forall a \in$  stocks Mkt. borel-adapt-sp  $\mathcal{F}$  (prices Mkt  $a$ )

```

lemma REPLICATING-FAIR-PRICE-UNIQUE

```

assumes replic-pf  $p \kappa T$ 
and fair-price Mkt  $\pi \kappa T$ 
shows  $\pi =$  (init-value  $p$ )

```

The principle of the proof is as follows. Assume for example that $\pi > (\text{init-value } p)$, where p is a replicating portfolio for derivative κ . Then the derivative can be sold for price π , the amount $(\text{init-value } p)$ can be invested to construct the replicating portfolio p and the amount $\pi - (\text{init-value } p)$ invested in the strictly positive stock (typically a money market account⁸). The initial value of the corresponding portfolio is 0, and at maturity, the closing value of p can be used to pay the derivative's payoff. The amount invested in the positive stock is untouched, hence an arbitrage was constructed, contradicting the fact that π is a fair price. The proof when $\pi < (\text{init-value } p)$ is similar, this time the derivative is bought to create an arbitrage.

5 Risk-neutral probability spaces

The results of Section 4.2 show that when a replicating portfolio exists for a given derivative, the fair price for this derivative is unique and equal to the initial value of the portfolio. In this section we prove that this initial value can be computed without explicitly constructing any replicating portfolio under the hypothesis of the existence of a *risk-neutral probability space*.

5.1 Interest rates and discounted values

We begin by defining the notion of interest rates. The existence of an interest rate is modeled by assuming that the market contains a stock with a deterministic return. The price process of this stock is parameterized by an interest

⁸ A money market account represents a deposit account on which any amount of cash can be deposited/withdrawn at each time.

rate r . In this setting, the interest rate is constant, although there exist more general models in which the interest rate can be time-dependent, and even stochastic. For simplicity, it is common to assume that the initial value of the asset is 1.

```

disc-rfr-proc          :: ℝ → ℕ → α → ℝ
disc-rfr-proc r 0 w    = 1
disc-rfr-proc r (n + 1) w = (1 + r).(disc-rfr-proc n w)

```

We call an asset a *risk-free* if a is such that $\text{prices Mkt } a = \text{disc-rfr-proc } r$ for some rate r , and define a locale for a market containing a risk-free asset.

```

locale risk-free-stock-market = disc-equity-market +
  fixes rf-asset :: β
  and r :: ℝ
  assumes - 1 < r
  and rf-asset ∈ stocks Mkt
  and prices Mkt rf-asset = disc-rfr-proc r

```

Having a risk-free asset as a stock in a market makes it possible to deposit (by buying the asset) or borrow (by shorting the asset) cash on this market.

Example 5 Assume there is a risk-free asset with an annual rate of 2% on the market. This means that buying 100€ worth of the asset today entails receiving 102€ by selling the asset in one year. Assuming the time lapse between times n and $n+1$ is a day and there are 252 business days in one year, the daily rate r in the definition of `disc-rfr-proc` then satisfies the equation $(1 + r)^{252} = 1.02$, so we have $r \approx 7.85 \cdot 10^{-5}$.

Remark 6 Observe that if the market is viable, then all risk-free assets must have the same rate⁹. Indeed, if there exist two risk-free assets with interest rates $r_1 < r_2$ then an arbitrage can be constructed: it suffices to buy 1 share of the second asset and sell 1 share of the first one. The initial investment is $1 - 1 = 0$, and at any time $n > 0$ the closing value of the portfolio is $(1 + r_2)^n - (1 + r_1)^n > 0$.

We also define the *discounted value* of a stochastic process. This notion is related to that of the present value of a future cash-flow, given an interest rate.

```

discount-factor        :: ℝ → ℕ → α → ℝ
discount-factor r n w = inverse (disc-rfr-proc r n w)

discounted-value      :: ℝ → (ℕ → α → ℝ) → ℕ → α → ℝ
discounted-value r X = λ n w. (discount-factor r n w).(X_n w)

```

⁹ Recall that the model of an equity market does not model foreign-exchanges with several currencies, although more sophisticated models for this setting do exist. The latter are closer to reality, since they permit to account for, e.g., the fact that national banks may have different risk-free rates

Example 7 Assume we have a viable market that contains a risk-free asset with a rate of 2% per year, and that the price of a share of Apple today is 95€. Consider a forward contract for buying a share of Apple stock at a strike price of 98€ in two years. The fair price for this contract is obtained by computing the discounted value of the strike and subtracting it from the current price of a share today. Here the discounted value of the strike is $98 \cdot (1 + 0.02)^{-2} \approx 94.19$, hence the fair price of this contract is 0.81€. Indeed, this amount of money can be used to construct a replicating portfolio as follows.

1. Borrow 94.19€ today.
2. Use the cash, along with the 0.81€ received at the sale of the contract to buy a share of Apple stock today.
3. Wait for two years.
4. Sell the share of Apple stock to the buyer of the forward contract for 98€.
5. Use this to reimburse the cash that was borrowed at the start and is now worth $94.19 \cdot (1 + 0.02)^2 \approx 98$ €.

5.2 Conditional expectations and martingales

From a financial point of view, assets carry different levels of risk. The more risky an asset, the higher the return buyers will be expecting when investing in the asset. This additional return is called the *market price of risk*. A risk-neutral probability space is meant to represent a world in which investors do not expect an increased return for a more risky asset: they are neutral to risk and expect the returns of all assets to be identical.

The expected returns of assets are modeled using the notion of *conditional expectations*. A conditional expectation is meant to represent the best approximation of a random variable given the currently available information. Formally, a conditional expectation of a random variable X given a measure space \mathcal{N} that is a subalgebra of \mathcal{M} is any random variable $X_{\mathcal{N}}$ that is \mathcal{N} -measurable, and such that for any set $N \in \mathcal{N}$,

$$\int_N X_{\mathcal{N}} d\mu_{\mathcal{M}} = \int_N X d\mu_{\mathcal{M}}.$$

A conditional expectation of X given \mathcal{N} always exists as long as X is integrable, and is almost surely unique, meaning that two conditional expectations of X given \mathcal{N} are identical almost everywhere. In what follows, we will therefore refer to *the* conditional expectation of X given \mathcal{N} , and denote it by $\mathbb{E}^{\mathcal{M}}[X | \mathcal{N}]$ (or simply by $\mathbb{E}[X | \mathcal{N}]$ if there is no possible confusion). Conditional expectations are already formalized in Isabelle. One property of conditional expectations that is particularly useful in our setting is that, for an initially trivial filtration, the conditional expectation of a random variable

at time 0 coincides with its expectation:

lemma TRIVIAL-SUBALG-COND-EXPECT-EQ
assumes subalgebra $\mathcal{M} \mathcal{N}$
and sets $\mathcal{N} = \{\emptyset, \Omega_{\mathcal{M}}\}$
and integrable $\mathcal{M} X$
shows $\forall x \in \Omega_{\mathcal{M}}. \mathbb{E}[X | \mathcal{N}](x) = \mathbb{E}[X]$

Conditional expectations are used to define *martingales*. Martingales are an essential tool in risk-neutral pricing, and we refer the reader to [21, Chapter 2] for a gentle introduction to this notion. Given a filtration \mathcal{F} , martingales are stochastic processes $(X_t)_t$ such that for all $t \leq s$, X_t is the best estimation of X_s given the information \mathcal{F}_t . In other words, X_t and $\mathbb{E}[X_s | \mathcal{F}_t]$ are equal almost everywhere.

martingale $:: \alpha \text{ measure} \rightarrow (\iota \rightarrow \alpha \text{ measure}) \rightarrow (\iota \rightarrow \alpha \rightarrow \mathbb{R}) \rightarrow \mathbb{B}$
martingale $\mathcal{M} \mathcal{F} X \Leftrightarrow$
 (**filtration** $\mathcal{M} \mathcal{F}$) \wedge (**borel-adapt-sp** $\mathcal{F} X$) \wedge
 ($\forall t. \text{integrable } \mathcal{M} X_t$) \wedge ($\forall t s. t \leq s \Rightarrow (\text{AE}_{\mathcal{M}} w. X_t w = \mathbb{E}[X_s | \mathcal{F}_t] w)$)

Because the risk-free asset we defined has a deterministic price process with a constant return rate, it is straightforward to verify that the discounted value of this price process is constant, and is trivially a martingale. In a risk-neutral probability space, the martingale property holds for *all* the stocks of the market:

risk-neutral-prob-space $:: \alpha \text{ measure} \rightarrow \mathbb{B}$
risk-neutral-prob-space $\mathcal{N} \Leftrightarrow \text{prob-space } \mathcal{N} \wedge$
 $\forall a \in (\text{stocks Mkt}). \text{martingale } \mathcal{N} \mathcal{F} (\text{discounted-value } r (\text{prices Mkt } a))$

5.3 Filtration-equivalence

If there were no relationship whatsoever between a risk-neutral probability space and the actual probability space, the former would not be of much use. In general, both spaces are assumed to be *equivalent*, meaning that they agree on the events that have a zero probability. Most textbooks rely on the notion of equivalence, but it turns out that when there is a filtration associated with a probability space, this notion can be relaxed into that of *filtration-equivalence*, which is sufficient for our purpose.

filt-equiv $:: (\iota \rightarrow \alpha \text{ measure}) \rightarrow \alpha \text{ measure} \rightarrow \alpha \text{ measure} \rightarrow \mathbb{B}$
filt-equiv $\mathcal{F} \mathcal{M} \mathcal{N} \Leftrightarrow \text{filtration } \mathcal{M} \mathcal{F} \wedge \mathcal{A}_{\mathcal{M}} = \mathcal{A}_{\mathcal{N}} \wedge$
 $\forall i A. A \in \mathcal{A}_{\mathcal{F}_i} \Rightarrow (\mu_{\mathcal{M}}(A) = 0 \Leftrightarrow \mu_{\mathcal{N}}(A) = 0)$

An advantage of filtration-equivalence which is used in this formalization is that it can be much simpler to prove that two probability spaces are filtration-equivalent, especially when the underlying space $\Omega_{\mathcal{M}}$ is infinite. When probability spaces are filtration-equivalent, almost everywhere properties propagate

from one space to the other. In particular, a replicating portfolio for a derivative in one given probability space will necessarily be a replicating portfolio for the derivative in a filtration-equivalent probability space, even if the probabilities assigned to different events may differ.

lemma `FILT-EQUIV-BOREL-AE-EQ`
assumes `filt-equiv \mathcal{F} \mathcal{M} \mathcal{N}`
and `$f \in \text{borel-measurable } \mathcal{F}_i$`
and `$g \in \text{borel-measurable } \mathcal{F}_i$`
and `$\text{AE}_{\mathcal{M}} w. f w = g w$`
shows `$\text{AE}_{\mathcal{N}} w. f w = g w$`

The following result states that, in a filtration-equivalent risk-neutral probability space, the closing value of any self-financing trading strategy must be a martingale. The only necessary hypothesis to obtain this result is an integrability condition on the assets of this trading strategy.

lemma `SELF-FIN-TRAD-STRAT-MART`
assumes `filt-equiv \mathcal{F} \mathcal{M} \mathcal{N}`
and `risk-neutral-prob-space \mathcal{N}`
and `trading-strat p`
and `self-financing Mkt p`
and `stock-portfolio Mkt p`
and `$\forall n. \forall a \in \text{support-set } p. \text{integrable } \mathcal{N}$`
`($\lambda w. (\text{prices Mkt } a n w)(p a (n + 1) w)$)`
and `$\forall n. \forall a \in \text{support-set } p. \text{integrable } \mathcal{N}$`
`($\lambda w. (\text{prices Mkt } a (n + 1) w)(p a (n + 1) w)$)`
shows `martingale \mathcal{N} \mathcal{F} (discounted-value r (cls-val-process Mkt p))`

We obtain the following result, which in a viable market, provides an effective way of computing the fair price of an attainable derivative. When a filtration-equivalent risk-neutral probability space exists, this fair price can be computed by considering the discounted value of the derivative payoff, and returning its expectation in the risk-neutral probability space.

lemma `REPLICATING-EXPECTATION`
assumes `filt-equiv \mathcal{F} \mathcal{M} \mathcal{N}`
and `risk-neutral-prob-space \mathcal{N}`
and `$\kappa \in \text{borel-measurable } \mathcal{F}_T$`
and `replic-pf p κ T`
and `$\forall n. \forall a \in \text{support-set } p. \text{integrable } \mathcal{N}$`
`($\lambda w. (\text{prices Mkt } a n w)(p a (n + 1) w)$)`
and `$\forall n. \forall a \in \text{support-set } p. \text{integrable } \mathcal{N}$`
`($\lambda w. (\text{prices Mkt } a (n + 1) w)(p a (n + 1) w)$)`
and `viable-market Mkt`
and `$\mathcal{A}_{\mathcal{F}_0} = \{\{\}, \Omega_{\mathcal{M}}\}$`
shows `fair-price Mkt $\mathbb{E}^{\mathcal{N}}$ [discounted-value r κ T] κ T`

The proof of this result goes as follows. If portfolio p is a replicating portfolio for derivative κ under probability space \mathcal{M} , then it is also a replicating portfolio for κ under the probability space \mathcal{N} , since \mathcal{N} and \mathcal{M} are filtration-equivalent. By Lemma `REPLICATING-FAIR-PRICE-UNIQUE`, the initial value of p is the unique fair price of κ . By Lemma `SELF-FIN-TRAD-STRAT-MART` the discounted closing value process of p is an \mathcal{N} -martingale, thus, by Lemma `TRIVIAL-SUBALG-COND-EXPECT-EQ` we have

$$\begin{aligned} \text{init-value } p &= \mathbb{E}^{\mathcal{N}}[\text{discounted-value } r \text{ (cls-val-process Mkt } p) T \mid \mathcal{F}_0] \\ &= \mathbb{E}^{\mathcal{N}}[\text{discounted-value } r \text{ (cls-val-process Mkt } p) T] \\ &= \mathbb{E}^{\mathcal{N}}[\text{discounted-value } r \kappa T]. \end{aligned}$$

6 Fair prices in the Cox-Ross-Rubinstein model

The Cox-Ross-Rubinstein model (or CRR model) is a discrete-time model consisting of a market in which there are two stocks, a risk-free asset and a risky one. At every time n , the risky asset can only move upward or downward with respective probabilities p and $1 - p$. This means that the evolution of the risky asset price can be modeled by tossing at each time n a coin that lands on its head with a probability p , and having the price move upward at time $n + 1$ exactly when the coin lands on its head. The evolution of this price is thus controlled by sequences of coin tosses. In most introductory textbooks on the CRR model, these sequences are finite as the results are presented for a given derivative with a finite maturity. We choose to consider infinite sequences—or streams—of coin tosses for the sake of generality. Since at time n no event other than the outcome of the coin toss is required, this outcome can be represented by a Bernoulli distribution of parameter p . In Isabelle, because discrete probability distributions and probability mass functions are isomorphic, the type of probability mass functions are defined as a subtype of measures [13], along with an injective representation function `measure-pmf :: α pmf \rightarrow α measure`. The Bernoulli distribution is thus defined as `measure-pmf (bernoulli-pmf p)`. The measure space for infinite sequences of independent coin tosses is isomorphic to the infinite product of Bernoulli distributions with the same parameter. In Isabelle, this measure space is defined using the function `stream-space :: α measure \rightarrow (α stream) measure`. The measure space thus defined is the smallest one in which the function `nth :: α stream \rightarrow $\mathbb{N} \rightarrow \alpha$` such that `(nth s n)` is the n th element of stream s is measurable [12]. The measure spaces we consider are defined as follows:

$$\begin{aligned} \text{bernoulli-stream} &:: \mathbb{R} \rightarrow (\mathbb{B} \text{ stream}) \text{ measure} \\ \text{bernoulli-stream } p &= \text{stream-space (measure-pmf (bernoulli-pmf } p)) \end{aligned}$$

We define a locale in which we impose that $0 \leq p \leq 1$, and thus obtain a probability space:

$$\begin{aligned} \text{locale infinite-coin-toss-space} &= \\ \text{fixes } p \text{ and } M & \\ \text{assumes } 0 \leq p \leq 1 \text{ and } M &= \text{bernoulli-stream } p \end{aligned}$$

The fact that an infinite coin toss space is a probability space is expressed using the **sublocale** keyword:

sublocale infinite-coin-toss-space \subseteq **prob-space**

The maximal amount of information that should be available at time n is the outcome of the first n coin tosses, and we define a filtration \mathcal{F}^{nat} accordingly: intuitively, in this filtration, two streams of coin tosses with the same first n outcomes cannot occur in distinct sets that are measurable in \mathcal{F}^{nat} . In our setting, each restricted measure space $\mathcal{F}_n^{\text{nat}}$ can be defined as generated by an arbitrary measurable function which maps all streams that agree on the first n coin tosses to the same element. We thus considered the sequence of so-called *pseudo-projection functions* $(\pi_n^\top)_{n \in \mathbb{N}}$, where:

$$\begin{aligned} \pi_n^\top : \quad \Omega_{\mathcal{M}} &\rightarrow \Omega_{\mathcal{M}} \\ (w_1, \dots, w_n, w_{n+1}, \dots) &\mapsto (w_1, \dots, w_n, \top, \top, \dots) \end{aligned}$$

These functions are measurable and permit to define a sequence of restricted measure spaces which is indeed a filtration:

$$\begin{aligned} \mathcal{F}^{\text{nat}} &:: \mathbb{N} \rightarrow (\mathbb{B} \text{ stream}) \text{ measure} \\ \mathcal{F}_n^{\text{nat}} &= \mathcal{M}_{\langle \pi_n^\top \rangle} \end{aligned}$$

We can thus define a locale for the infinite coin toss space along with this filtration:

locale infinite-cts-filtration = **infinite-coin-toss-space** +
fixes \mathcal{F} **assumes** $\mathcal{F} = \mathcal{F}^{\text{nat}}$

Because the information available at time n is the outcome of the first n coin tosses, $\mathcal{F}_n^{\text{nat}}$ is initially trivial:

sublocale infinite-cts-filtration \subseteq **init-triv-prob-space**

Any other considered filtration on this probability space will be a sub-filtration of the natural filtration.

An important feature of the natural filtration is that the expectation of any $\mathcal{F}_n^{\text{nat}}$ -measurable function is very similar to that of a function on a finite probability space: for $w = w_1, \dots, w_n, \dots$ and $i \in \mathbb{N}$, if $\nu_i(w) \stackrel{\text{def}}{=} \mathbf{if } w_i \text{ then } p \text{ else } 1 - p$, then we have

lemma EXPECT-PROB-COMP
assumes $f \in \text{borel-measurable } \mathcal{F}_n^{\text{nat}}$
shows $\mathbb{E}[f] = \sum_{w \in \pi_n^\top(\Omega_{\mathcal{M}})} (\prod_{i=1}^n \nu_i(w_i)) \cdot f(w)$

In the CRR model, the price of the risky asset is modeled by a *geometric random walk* with parameters specifying the upward and downward movements as well as the price of the asset at time 0:

geom-rand-walk $:: \mathbb{R} \rightarrow \mathbb{R} \rightarrow \mathbb{R} \rightarrow$
 $(\mathbb{N} \rightarrow (\mathbb{B} \text{ stream}) \rightarrow \mathbb{R})$
 $(\text{geom-rand-walk } u \ d \ v) \ 0 \ w = v$
 $(\text{geom-rand-walk } u \ d \ v) \ (n + 1) \ w = (\mathbf{if } w_n \ \mathbf{then } u \ \mathbf{else } d) \cdot$
 $(\text{geom-rand-walk } u \ d \ v) \ n \ w$




Fig. 1 Example of a geometric random walk

Example 8 Figure 1 depicts the first values of the geometric random walk process (`geom-rand-walk 1.2 0.8 10`). This is a process with a deterministic initial value 10. Intuitively, at time 1, if the outcome of a coin toss is a head, then this process has a value of 12, and if the outcome is a tail, then it has a value of 8. If at time 2 the first two outcomes are a head then a tail, then the value is 9.6, etc.

The geometric random walk process is an adapted process, in the infinite coin toss space equipped with its natural filtration, since its value at time n depends only on the outcome of the first n coin tosses:

lemma GEOM-RAND-WALK-BOREL-ADAPTED :
`borel-adapt-sp (geom-rand-walk u d v)`

We define a locale in which there is a stochastic process that is a geometric random walk:

```

locale prob-grw = infinite-coin-toss-space +
fixes geom-proc and u and d and v
assumes geom-proc = geom-rand-walk u d v

```

The locale for the market in the CRR model is defined as follows:

```

locale CRR-hyps = prob-grw + risk-free-stock-market +
fixes S
assumes stocks Mkt = {S, rf-asset}
and prices Mkt S = geom-proc
and 0 < v and 0 < d < u
and 0 < p < 1

```

In particular, we require that $0 < p < 1$ and $u \neq d$, so that S is indeed a risky asset. Note that we do not postulate that $d < 1$ or $u > 1$.

The filtration associated with this probability space is meant to represent the fact that the information available at time n is the price evolution of the risky asset up to time n . We thus define a function that associates a filtration to a stochastic process X , such that at time n , the corresponding measure space is the smallest subalgebra for which X_k is measurable for all $k \leq n$.

```
stoch-proc-filt      ::  $\alpha$  measure  $\rightarrow (\mathbb{N} \rightarrow \alpha \rightarrow \beta) \rightarrow \beta$  measure  $\rightarrow$ 
                     $\mathbb{N} \rightarrow \alpha$  measure
stoch-proc-filt  $\mathcal{M} X \mathcal{N} n = \text{sigma } \Omega_{\mathcal{M}} \cup_{k \leq n} \{X_k^{-1}(A) \cap \Omega_{\mathcal{M}} \mid A \in \mathcal{A}_{\mathcal{N}}\}$ 
```

In the locale below, we denote by \mathcal{G} the filtration such that at time n , \mathcal{G}_n is the smallest subalgebra for which prices $\text{Mkt } S k$ is Borel-measurable for all $k \leq n$.

```
locale CRR-market = CRR-hyps+
fixes  $\mathcal{G}$ 
assumes  $\mathcal{G} = \text{stoch-proc-filt } \mathcal{M} \text{ geom-proc borel}$ 
```

In order to compute fair prices, the CRR market is required to be viable. We have the following result:

```
lemma VIABLE-IFF
shows viable-market Mkt  $\Leftrightarrow (d < 1 + r < u)$ 
```

Proof (Outline) The left-to-right implication is straightforward to prove. For example, if the risky asset always has a return greater than the risk-free rate (i.e., $1 + r \leq d$), then an arbitrage can be obtained as follows. At time 0, borrow v , the initial value of the risky asset, and use the cash to buy one share of the risky asset. This results in a portfolio with initial value 0. At time 1, the closing value of the portfolio is either $dv - (1 + r)v$ or $uv - (1 + r)v$; in both cases this value is positive, and it is strictly positive with probability $p > 0$. The market can therefore not be viable.

Assume that $d < 1 + r < u$ and that there exists a stock portfolio p that is also an arbitrage process. Then the initial value of p is 0 and there is a time m at which its closing value is almost surely positive, and strictly positive with a nonzero probability. Consider an element $y \in \Omega_{\mathcal{M}}$ such that $\text{cls-val-process } p m y > 0$, and define

$$n_0 \stackrel{\text{def}}{=} \min \{n \leq m \mid \text{cls-val-process } p n y > 0\}.$$

Intuitively, once there is a scenario at which the closing value of a portfolio is negative at a given time, there necessarily exists a scenario at which its closing value at the following time is also strictly negative. Because by definition $\text{cls-val-process } p (n_0 - 1) y \leq 0$, there must exist an element $x' \in \Omega_{\mathcal{M}}$ such that $\text{cls-val-process } p m x' < 0$. Now if $z \in \Omega_{\mathcal{M}}$ has the same prefix as x' up to time m , then $\text{cls-val-process } p m z = \text{cls-val-process } p m x'$. The

set of elements in $\Omega_{\mathcal{M}}$ with the same prefix as x' up to time m has a nonzero measure, we therefore deduce that

$$\mathcal{P}(\{w \in \Omega_{\mathcal{M}} \mid (\text{cls-val-process } p) \text{ } m \text{ } w < 0\}) > 0,$$

contradicting the fact that the closing value of p is almost surely positive.

We may thus define a locale for a viable CRR market:

locale CRR-market-viable = CRR-market+
assumes viable-market Mkt

Next, we provide a necessary and sufficient condition for the existence of a risk-neutral bernoulli stream space that is filtration-equivalent to \mathcal{M} .

lemma RISK-NEUTRAL-IFF

assumes $\mathcal{N} = \text{bernoulli-stream } q$

and $0 < q < 1$

shows risk-neutral-prob-space \mathcal{G} Mkt r $\mathcal{N} \Leftrightarrow q = \frac{1+r-d}{u-d}$

We finally prove that every derivative is attainable in the CRR model:

lemma CRR-MARKET-COMPLETE :

shows complete-market

The result is proven by constructing a replicating portfolio for any \mathcal{G}_T -measurable payoff $\kappa : \alpha \rightarrow \mathbb{R}$ and exercise time T . Note that the fact that function κ is \mathcal{G}_T -measurable ensures that the payoff only depends on information available up to time T . The principle of the construction of the portfolio is explained in details on an example in Section 7.

We obtain the final result:

lemma CRR-MARKET-FAIR-PRICE :

assumes $\kappa \in \text{borel-measurable } \mathcal{G}_T$

and $\mathcal{N} = \text{bernoulli-stream } \frac{1+r-d}{u-d}$

shows fair-price Mkt

$$\left[\sum_{w \in \pi_T^{\top}(\Omega_{\mathcal{M}})} \left(\prod_{i=1}^T \nu_i(w_i) \cdot (\text{discounted-value } r \text{ } \kappa \text{ } w) \right) \right] \kappa T$$

7 A complete example

The entry <https://www.isa-afp.org/entries/DiscretePricing.html> contains functions for computing the fair prices of a few standard options. For example, it contains a function that permits to compute the fair price of a *lookback option* (see, e.g., [10]). A lookback option is characterized by a maturity T , and at this maturity, pays $\max_{0 \leq i \leq T} S_i - S_T$. In other words, instead of having a payoff that only depends on the value of the risky asset at maturity, a lookback option has a payoff that depends on *all* the values of the risky asset

Table 6 Possible outcomes for the lookback option.

Outcomes	TT	T↓	↓T	↓↓
Probability	0.330625	0.244375	0.244375	0.180625
Payoff	0	2.4	0.4	3.6
Disc. payoff	0	2.262	0.377	3.393

Table 7 Possible outcomes for the lookback option when the first coin toss is a head.

Outcomes	TT	T↓
Probability	0.575	0.425
Payoff	0	2.4
Disc. payoff	0	2.33

until maturity. It is called a *path-dependent* option. In Isabelle, the payoff of such an option with maturity T is represented by `lbk-option T`.

The function `lbk-price` is based on Lemma CRR-MARKET-FAIR-PRICE. Consider a Cox-Ross-Rubinstein market parameterized by upward movement u , downward movement d and initial value v . If r denotes the risk-free rate, then we have the following lemma:

lemma LBK-PRICE :
shows fair-price Mkt
 $(\text{lbk-price } u \ d \ v \ r \ T) \ (\text{lbk-option } T) \ T$

Assume that in the Cox-Ross-Rubinstein market, the risky asset has an initial value 10, an upward movement $u = 1.2$ and a downward movement $d = 0.8$. The risk-free rate is $r = 3\%$ (see Figure 2). Consider a lookback option with maturity $T = 2$. Its payoff is depicted on the right-hand side of the figure. If the outcomes of the first two coin tosses are heads, then the maximal value of the risky asset is its value at time 2, so that this option does not pay anything. If the outcomes are first a head then a tail, then the value of the risky asset at time 2 is 9.6€ and the option pays off 2.4€. Note that if the first two coin tosses are a tail then a head, then the value of the risky asset at time 2 is also 9.6€, but the option only pays off 0.4€.

The fair price of this option is computed in Isabelle using the command

```
value lbk-price 1.2 0.8 10 0.03 2
```

The output is $\frac{13345}{10609}$, which means that the fair price of this option is approximately 1.2579€.

In what follows, we construct a replicating portfolio for this option to give an intuition of the way the completeness of the CRR model is proved. This portfolio will be constructed by going backward in time.

First assume the outcome of the first coin toss is a head. In this scenario, we construct a portfolio that starts at time 1. The fair price of the option is

Table 8 Possible outcomes for the lookback option when the first coin toss is a tail.

Outcomes	$\perp\top$	$\perp\perp$
Probability	0.575	0.425
Payoff	0.4	3.6
Disc. payoff	0.38835	3.49515

given using Table 7, from which we deduce that the fair price of the option (and the initial value of the portfolio under construction starting at time 1) is approximately 0.9903€. The quantity invested in the risky asset is given by

$$\Delta_{\top} \stackrel{\text{def}}{=} \frac{\kappa_{\top\top} - \kappa_{\top\perp}}{S_{\top\top} - S_{\top\perp}} = \frac{0 - 2.4}{14.4 - 9.6} = -0.5,$$

where $\kappa_{w_1 w_2}$ and $S_{w_1 w_2}$ respectively denote the payoff of the derivative and the value of the risky asset at time 2, depending on the outcomes of the first two coin tosses w_1 and w_2 . The quantity Δ_{\top} can be viewed as the discrete version of the derivative of the payoff w.r.t. the risky asset. This is no coincidence: in continuous-time models, the derivative of an option's value w.r.t. an underlying security is called the *delta* of the option, and buying/selling this quantity of the underlying security permits to obtain what is called a *delta-neutral portfolio*, the value of which is globally unaffected by price movements of the underlying security. In discrete- and continuous-time models, the quantity delta is key in constructing replicating portfolios, and constructing delta-neutral portfolios is known as *delta-hedging*.

This means that half a share of the risky asset is sold (a short sell) for 6€. Since the initial value of the portfolio is 0.9903€, this cash, along with the one obtained by selling the risky asset for a total of 6.9903€, is invested in the risk-free rate. At time 2, the cash invested in the risk-free rate is recovered and worth $6.9903 * 1.03 = 7.2$ €; the half-share of the risky asset is bought back.

- If the outcome of the second coin toss is a head, then the risky asset is worth 14.4€, so 7.2€ are necessary to buy half the share back. The value of the portfolio is 0€.
- If the outcome of the second coin toss is a tail, then the risky asset is worth 9.6€, so 4.8€ are necessary to buy half the share back. The value of the portfolio is 2.4€.

Now assume the outcome of the first coin toss is a tail. The fair price of the option at time 1 in this scenario is given using the Table 8, from which we deduce that the fair price of the option under this scenario is approximately 1.7087€. The quantity invested in the risky asset is

$$\Delta_{\perp} \stackrel{\text{def}}{=} \frac{\kappa_{\perp\top} - \kappa_{\perp\perp}}{S_{\perp\top} - S_{\perp\perp}} = \frac{0.4 - 3.6}{9.6 - 6.4} = -1.$$

One share of the risky asset is sold for 8€, and the proceeds of this sale, along with the initial value of the portfolio are invested in the risk-free asset. At time




Fig. 2 Lookback option settings and payoff. R denotes the risk-free asset and S the risky one.

2, the amount thus invested is worth 10€, and the share of the risky asset is bought back.

- If the outcome of the second coin toss is a head, then the risky asset is worth 9.6€, so the value of the portfolio is 0.4€.
- If the outcome of the second coin toss is a tail, then the risky asset is worth 6.4€, so the value of the portfolio is 3.6€.

We now construct a portfolio with initial value 1.2579€, and worth 0.9903€ if the outcome of the first coin toss is a head, and 1.7087€ if the outcome is a tail. The quantity invested in the risky asset is given by

$$\Delta_{\top} \stackrel{\text{def}}{=} \frac{0.9903 - 1.7087}{12 - 8} = -0.1796.$$

This quantity of the risky asset is sold for 10€, and the proceeds are invested in the risk-free asset, along with the initial value of the portfolio. At time 1, the cash thus invested is worth 3.145517€. The quantity of risky asset that was shorted is bought back.

- If the outcome of the first coin toss is a head, then the risky asset is worth 12€ and buying back the quantity that was shorted costs 2.1552€, so the value of the portfolio is 0.9903€.
- If the outcome of the first coin toss is a tail, then the risky asset is worth 8€ and buying back the quantity that was shorted costs 1.4368€, so the value of the portfolio is 1.7087€.

To recap, the seller of the lookback option sells it for 1.2579€, and constructs a replicating portfolio as follows.

1. The seller receives 1.796€ by short selling the risky asset and invests the 3.0539€ in the risk-free asset until time 1.
2. If at time 1 the outcome of the first coin toss is a head, then the seller uses the closing value of the portfolio, 0.9903€, to short half a share of the risky asset and invest 6.9903€ in the risk-free asset. Otherwise, the seller uses the closing value of the portfolio, 1.7087€, to short one share of the risky asset and invest 9.7087€ in the risk-free asset.
3. At time 2, quantity of risky asset that was shorted is bought back and the cash invested in the risk-free asset is withdrawn; the closing value of the portfolio is exactly equal to the payoff of the lookback option.

This construction can be generalized to arbitrary $\mathcal{F}_T^{\text{nat}}$ -measurable functions. At any time $t < T$, the composition of the portfolio is determined in such a way that its closing value at time $t + 1$ matches the value already computed at time $t + 1$, for both outcomes of the next coin toss. This yields a system of two linear equations, one for each possible outcome of the toss coin, with two variables (the amounts of risk-free and risky assets, respectively). Lemma VIABLE-IFF on Page 28 imposes additional conditions on u, d, r that ensure that the system admits a unique solution.

8 Discussion

We have formalized a framework for proving financial results in Isabelle. It permits a formal definition of fair prices in Isabelle, and presents one of the main pricing results in finance: under a risk-neutral probability, the fair price of an attainable derivative is equal to the expectation of its discounted payoff. This formalization is quite extensive, as many financial notions had to be introduced, and it was used to prove that every derivative admits a fair price in the Cox-Ross-Rubinstein model of an equity market, by proving the completeness of this market. The proof is constructive and we also provide pricing functions for a few standard derivatives. As far as future work is concerned, we intend to work on the pricing in the Cox-Ross-Rubinstein model of American options, that can be exercised at any time by the buyer until the maturity –and not simply at maturity, as for European options. Pricing such options will require the definition of additional notions, such as *sub* and *supermartingales*, and our aim will be to implement a completely certified pricer for such options. We also intend to pursue our formalization effort of mathematical finance and extend our results to a continuous-time setting. This is an ambitious and interesting task, and we hope this first formalization will encourage other researchers, from computer science and financial mathematics to extend these results in Isabelle.

References

1. J. Avigad, J. Hölzl, and L. Serafin. A formally verified proof of the Central Limit Theorem. *J. Autom. Reasoning*, 59(4):389–423, 2017.

2. C. Ballarín. Locales: A module system for mathematical theories. *J. Autom. Reasoning*, 52(2):123–153, 2014.
3. N. H. Bingham and R. Kiesel. *Risk-neutral valuation : pricing and hedging of financial derivatives*. Springer-Verlag London, 2004.
4. F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.
5. J. C. Blanchette, J. Hölzl, A. Lochbihler, L. Panny, A. Popescu, and D. Traytel. Truly modular (co)datatypes for Isabelle/HOL. In G. Klein and R. Gamboa, editors, *Interactive Theorem Proving - 5th International Conference, ITP 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014. Proceedings*, volume 8558 of *Lecture Notes in Computer Science*, pages 93–110. Springer, 2014.
6. J. C. Cox, S. A. Ross, and M. Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, 1979.
7. R. Durrett. *Probability : theory and examples*. The Wadsworth & Brooks/Cole statistics/probability series. Wadsworth Inc. Duxbury Press, Belmont, California, 1991.
8. M. Eberl. Randomized Social Choice Theory. https://www.isa-afp.org/entries/Randomised_Social_Choice.html, May 2016.
9. M. Echenim and N. Peltier. The binomial pricing model in finance: A formalization in Isabelle. In L. de Moura, editor, *Automated Deduction - CADE 26 - 26th International Conference on Automated Deduction, Gothenburg, Sweden, August 6-11, 2017, Proceedings*, volume 10395 of *Lecture Notes in Computer Science*, pages 546–562. Springer, 2017.
10. S. G. Kou. Chapter 8 - Discrete barrier and lookback options. *Handbooks in Operations Research and Management Science*, 15, 12 2007.
11. J. Hölzl. *Construction and Stochastic Applications of Measure Spaces in Higher-Order Logic*. PhD thesis, Institut für Informatik, Technische Universität München, October 2012.
12. J. Hölzl. Markov chains and Markov decision processes in Isabelle/HOL. *Journal of Automated Reasoning*, pages 1–43, 2016. <http://dx.doi.org/10.1007/s10817-016-9401-5>.
13. J. Hölzl, A. Lochbihler, and D. Traytel. A Formalized Hierarchy of Probabilistic System Types - Proof Pearl. In *Proc. of ITP*, pages 203–220, 2015.
14. J. Hull. *Options, Futures and Other Derivatives*. Pearson/Prentice Hall, 2009.
15. N. J. Cutland and A. Roux. *Derivative Pricing in Discrete Time*. Springer-Verlag London, 01 2013.
16. C. Kaliszyk and J. Parsert. Formal microeconomic foundations and the first welfare theorem. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2018, pages 91–101, New York, NY, USA, 2018. ACM.
17. R. Merton. The theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4:141–183, 1973.
18. T. Nipkow. Social choice theory in HOL. *Journal of Automated Reasoning*, 43(3):289–304, Oct 2009.
19. T. Nipkow and G. Klein. *Concrete Semantics: With Isabelle/HOL*. Springer Publishing Company, Incorporated, 2014.
20. T. Nipkow, M. Wenzel, and L. C. Paulson. *Isabelle/HOL: A Proof Assistant for Higher-order Logic*. Springer-Verlag, Berlin, Heidelberg, 2002.
21. S. E. Shreve. *Stochastic Calculus for Finance I: The Binomial Asset Pricing Model*. Springer Finance, 2003.
22. R. Vestergaard. A constructive approach to sequential Nash equilibria. *Inf. Process. Lett.*, 97(2):46–51, 2006.
23. F. Wiedijk. Formalizing Arrow’s theorem. *Sadhana*, 34(1):193–220, Feb 2009.