Notes on Foundations of Financial Mathematics – Draft

Alec Kercheval
Florida State University

Greg Anderson
MSCI-Barra, Inc.

Draft Version: September 21, 2004

Chapter 1

The Discrete Case

1.1 Setup

The basic idea is that things work out too nicely in the regular binomial tree case to really see what's going on. So we're going to carry out the analysis in more generality.

1.1.1 The Tree

Consider a finite probability space (Ω, μ). The measure μ is not really essential for us beyond the requirement that every $\omega \in \Omega$ should have positive probability.

We also want a filtration $\mathcal{F} = \{\mathcal{F}_t \text{ for } t = 0 \ldots T\}$ on Ω such that $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and $\mathcal{F}_T$ is the power set of Ω.

For example, take any tree with any finite branching you like at times $t = 0 \ldots T$. Then each path starting at time 0 and ending at time T is a point ω in the sample space, and every path should have positive probability (i.e. no spurious branching). The filtration here is "history up to time t". The equivalence relation "ω_0 and ω_1 agree up to and including time t" gives equivalence classes which are the atoms of $\mathcal{F}_t$. Given a path ω, we denote by $\omega(t)$ the atom of $\mathcal{F}_t$ containing ω, the set of paths agreeing with ω up to and including time t.

Conversely, given any finite probability space and a filtration as above, one can construct a tree such that the space of paths of the tree, along with the natural filtration induced by the tree as above, is isomorphic to the given space and filtration. Therefore we may use the tree terminology without any loss of generality.

1.1.2 The Stock Process

Now consider also a stock process $S = \{S_t \text{ for } t = 0 \ldots T\}$. Here S_t is a positive random variable on Ω which is $\mathcal{F}_t$-measurable, i.e. so that $S_t(\omega_0) = S_t(\omega_1)$ whenever ω_0 and ω_1 agree up to time t. This just means you decorate the nodes of the tree with positive stock prices.

A convenient level of generality is to let the S_t's be $\mathbb{R}^{k+1}$-valued random variables. So now

$$S_t = (S_t^0, \ldots, S_t^k)$$

is a list of positive prices of $k+1$ securities, and we agree that the first one, S_t^0 is called a bond. Remarkably, in the discrete case, the bond need not be distinguished in any way from the other securities, and in particluar need not be riskless. Each component has units of \$.

1.1.3 Trading Strategies and Attainable Claims

A trading strategy ϕ is a previsible $\mathbb{R}^{k+1}$-valued process,

$$\phi = \{\phi_t \text{ for } t = 1 \ldots T\} = \{(\phi_t^0, \ldots, \phi_t^k) \text{ for } t = 1 \ldots T\},\ \phi_t \text{ is } \mathcal{F}_{t-1}\text{-measurable.}$$

The interpretation is that from time $t-1$ to time t, your dimensionless vector of holdings of the $k+1$ securities is ϕ_t. How exactly you choose to rebalance can depend on the entire history of security prices up to that point. It's important to get the subscripts straight — note that you start off with the portfolio ϕ_1.

Having fixed a stock process earlier, associated to a trading strategy ϕ is another ($\mathbb{R}$-valued) process $V(\phi)$ giving the value of the indicated portfolio. For t between 1 and T, this is defined by

$$V_t(\phi) = \phi_t \cdot S_t$$

and for $t = 0$ we have $V_0 = \phi_1 \cdot S_0$. That dot, by the way, is the usual dot product in $\mathbb{R}^{k+1}$. Note that V_t is $\mathcal{F}_t$-measurable since S_t is, so that at time t you find out what your portfolio is worth that you picked at time $t-1$. Pay attention to the units of $V_t(\phi)$ — they are dollars at time t, just like S_t.

The only sort of strategy we care about is the self-financing kind. This means that at time $t-1$ you are holding ϕ_{t-1}, you want to rebalance to ϕ_t, and your buys should be exactly financed by your sales. In other words, the rebalanced portfolio should have the same value as the original one, or

$$\phi_{t-1} \cdot S_{t-1} = \phi_t \cdot S_{t-1}.$$

This is worth a few more words. For a general trading strategy, self-financing or not, the change in value of the indicated portfolio is

$$\begin{aligned} \Delta V_t &= V_t - V_{t-1} \\ &= \phi_t \cdot S_t - \phi_t \cdot S_{t-1} + \phi_t \cdot S_{t-1} - \phi_{t-1} \cdot S_{t-1} \\ &= \phi_t \cdot \Delta S_t + S_{t-1} \cdot \Delta \phi_t \end{aligned}$$

which is just the discrete version of the product rule for differentiation. The first term is the change in the value of the portfolio due to changes in security prices, and the second is the change due to rebalancing alone. The self-financing condition is precisely that the second term vanishes.

A *claim* is simply a (real-valued) random variable X defined on Ω. If stocks follow the particular path ω, think of the claim paying off $X(\omega)$ dollars at time T.

A claim is *attainable* if there is a self-financing trading strategy which replicates it, i.e. if there is a self-financing ϕ such that $V_T(\phi) = X$. The attainable claims are the ones that have "arbitrage enforced prices", as we describe below.

1.1.4 Arbitrage

Loosely speaking, an arbitrage is a way of making a riskless profit. One must be very careful in making this statement precise, since there is in fact a subtle gradation of types of arbitrage.

Definition 1 *In increasing degree of severity:*

- *A* (weak) arbitrage *is a self-financing strategy ϕ with $V_0(\phi) = 0$ and $V_T(\phi) \geq 0$ and $P(V_T(\phi) > 0) > 0$.*
- *A* strong arbitrage *is a self-financing strategy ϕ with $V_0(\phi) = 0$ and $V_T(\phi) > 0$.*
- *A market admits* inconsistent pricing *if there is a self-financing strategy ϕ with $V_T(\phi) \equiv 0$ and $V_0(\phi) \neq 0$.*

A market is viable *if it admits no arbitrage.*

For our purposes, the most important condition is the first one, the weak arbitrage. Unless otherwise noted, aribtrage always means weak arbitrage. Think of it as a free lottery ticket — it costs nothing to set up, will not lose money, and might make money. Note that only the equivalence class of the measure μ has made an appearance here, in the form of asserting that a certain set has positive measure.

By contrast, a strong arbitrage is a free lunch — it costs nothing to set up and will definitely make money. Evidently a strong arbitrage is an arbitrage. Some texts refer to the equivalent condition of the existence of a *dominant strategy*, i.e. if there are self-financing strategies ϕ and ψ such that $V_0(\phi) = V_0(\psi)$ and $V_T(\phi) > V_T(\psi)$ then ϕ dominates ψ.

For discrete models, the existence of a strong arbitrage is also equivalent to the existence of a self-financing strategy ϕ such that $V_0(\phi) < 0$ and $V_T(\phi) \geq 0$. Some texts define strong arbitrage this way.

The condition of inconsistent pricing is the most pernicious. It implies the existence of a strong arbitrage, since one may suppose $V_0(\phi) < 0$, then modify the strategy by holding $-V_0(\phi)/S_0^0$ shares of S_0 until maturity. If inconsistent pricing prevails, then for any attainable claim X and any initial price p, there is a self-financing strategy ϕ with $V_0(\phi) = p$ and $V_T(\phi) = X$. In such a market the concept of the price of a claim does not exist. Some texts refer to this case by saying the Law of One Price fails.

The relations among these types of arbitrage described above may be summarized as follows.

Proposition 1 *A viable market does not admit strong arbitrage. A market that does not admit strong arbitrage also does not admit inconsistent pricing.*

That is, in a viable market none of these kinds of arbitrage exist. It follows easily from the absence of inconsistent pricing that any attainable claim must have a unique price at time zero (the Law of One Price), namely $V_0(\phi)$, the cost of setting up the replicating portfolio. This is sometimes called the "arbitrage enforced price" of the claim.

The existence of arbitrage is a local property, in the sense that if there is an arbitrage strategy over the full time horizon from 0 to T, then in fact one can find a single state at a certain time such that an arbitrage exists in the next time step. More precisely, we say a market *admits a local arbitrage* if there exists a time t, a path ω_0, and a holdings vector h such that $h \cdot S_t(\omega_0) = 0$, $h \cdot S_{t+1}(\omega) \geq 0$ whenever $\omega \in \omega_0(t)$, and $P(h \cdot S_{t+1}(\omega) > 0 \,|\, \omega_0(t)) > 0$. As a practical matter, if one is presented with a tree and security prices on it, one may determine if an arbitrage exists by inspecting each node individually for a local arbitrage.

Proposition 2 *A market admits arbitrage if and only if admits a local arbitrage.*

Proof: First suppose there is a local arbitrage at time t. To create an arbitrage strategy, first do nothing until time t. Then, if state ω_0 obtains, trade into holding vector h (at zero cost). At time $t+1$, cash out any positive holdings into the numeraire and hold. The result at time T will be nonnegative with at least one positive state.

Conversely, suppose there exists an arbitrage strategy ϕ. Consider the set Λ of times $s \leq T$ such that $V_s(\phi)$ is nonnegative for all states and positive in at least one state. $T \in \Lambda$, so $\Lambda \neq \emptyset$. Let t be the least element of Λ. Note $t > 0$.

We now argue that at $t-1$ there is a local arbitrage. By choice of t, there are only two possibilities: either $V_{t-1}(\phi)$ is negative in some state ω_0, or else it is zero in all states.

If the former, then ω_0 is the desired state, since one needs only to add a positive position in the numeraire to bring the value in that state up to zero and thereby also making the value at time t positive.

If the latter, then choose a time $t-1$ state which leads to one of the positive values of $V_t(\phi)$. ■

For strong arbitrage and inconsistent pricing, local and global are not equivalent. The proof of the following is left as an excercise.

Proposition 3 *A market admits strong arbitrage (respectively, inconsistent pricing) only if admits a local strong arbitrage (respectively, local inconsistent pricing), but not conversely.*

1.1.5 The Discrete Arbitrage Pricing Theorem

In this section we try to illuminate the the importance of martingale measures, and their relation to the no arbitrage condition and to market completeness.

Definition 2 *If $(\Omega, \mu, \mathcal{F})$ is a measure space and S is a random process, we say that $(S, \mu, \mathcal{F})$ is a* martingale*, or μ-*martingale*, if for all s, t with $0 \leq s \leq t \leq T$,*

$$S_s = E_\mu(S_t | \mathcal{F}_s).$$

The discussion is best separated into three parts. The first explains why we are interested in martingales at all. (Note that the proof does not require ν equivalent to μ, or even positivity of ν.) We give two proofs.

Theorem 1 *Suppose ν is any measure such that $(S/S^0, \nu, \mathcal{F})$ is a martingale. For an attainable claim X with replicating strategy ϕ and $0 \leq t \leq T$, we have*

$$V_t(\phi) = E_\nu\left(X \frac{S_t^0}{S_T^0} \middle| \mathcal{F}_t \right).$$

First Proof: We simply calculate the expectation conditionally. To begin with, $E_\nu\left(X \frac{S_t^0}{S_T^0} \middle| \mathcal{F}_t \right) = E_\nu\left(V_T(\phi) \frac{S_t^0}{S_T^0} \middle| \mathcal{F}_t \right)$. Now for $\tau > t$ we have

$$\begin{aligned}
E_\nu\left(V_\tau(\phi) \frac{S_t^0}{S_\tau^0} \middle| \mathcal{F}_t \right) &= E\left(E\left(V_\tau(\phi) \frac{S_t^0}{S_\tau^0} \middle| \mathcal{F}_{\tau-1} \right) \middle| \mathcal{F}_t \right) \\
&= E\left(E\left(\phi_\tau \cdot S_\tau \frac{S_t^0}{S_\tau^0} \middle| \mathcal{F}_{\tau-1} \right) \middle| \mathcal{F}_t \right) \\
&= E\left(S_t^0 \phi_\tau \cdot E\left(\frac{S_\tau}{S_\tau^0} \middle| \mathcal{F}_{\tau-1} \right) \middle| \mathcal{F}_t \right) \\
&= E\left(S_t^0 \phi_\tau \cdot \frac{S_{\tau-1}}{S_{\tau-1}^0} \middle| \mathcal{F}_t \right) \\
&= E\left(\phi_{\tau-1} \cdot S_{\tau-1} \frac{S_t^0}{S_{\tau-1}^0} \middle| \mathcal{F}_t \right) \\
&= E\left(V_{\tau-1}(\phi) \frac{S_t^0}{S_{\tau-1}^0} \middle| \mathcal{F}_t \right)
\end{aligned}$$

by properties of conditional expectation, previsibility of ϕ, definition of martingale, and self-financing of ϕ, in that order. By induction, we find $E\left(X \frac{S_t^0}{S_T^0} \middle| \mathcal{F}_t \right) = E\left(V_t(\phi) \frac{S_t^0}{S_t^0} \middle| \mathcal{F}_t \right) = V_t(\phi)$. ∎

Second Proof: Consider the discounted value process

$$\tilde{V}_t = \frac{V_t(\phi)}{S_t^0} = \phi_t \cdot \tilde{S}_t$$

written here in terms of the discounted stock process $\tilde{S}_t = S_t/S^0_t$. That ϕ is self-financing means $S_t \cdot \Delta\phi_{t+1} = 0$, so we also have $\tilde{S}_t \cdot \Delta\phi_{t+1} = 0$. For the discounted value process, this means

$$\Delta\tilde{V}_{t+1} = \tilde{V}_{t+1} - \tilde{V}_t = \phi_{t+1} \cdot \Delta\tilde{S}_{t+1} + \tilde{S}_t \cdot \Delta\phi_{t+1} = \phi_{t+1} \cdot \Delta\tilde{S}_{t+1}.$$

We remark that we just showed

a strategy is self-financing with respect to the stock process if and only if it is self-financing with respect to the discounted stock process.

Previsibility of ϕ immediately yields

$$E\left(\Delta\tilde{V}_{t+1}\middle| \mathcal{F}_t\right) = \phi_{t+1} \cdot E\left(\Delta\tilde{S}_{t+1}\middle| \mathcal{F}_t\right) = 0.$$

This says that $\tilde{V}_t$ is a martingale, so in particular

$$\frac{V_t(\phi)}{S^0_t} = E\left(\frac{V_T(\phi)}{S^0_T}\middle| \mathcal{F}_t\right).$$

■

We may conclude that in a viable market the unique price of any attainable claim is equal to the expected value of the discounted claim using any martingale measure. In particular, all martingale measures price the attainable claims equally.

The next theorem tells us what it means to say that our market is viable in terms of martingale measures. To keep the proof concise, we will postpone commentary and narration until the next section on forwards.

Theorem 2
1. *There exists a probability measure ν equivalent to μ such that $(S/S^0, \nu, \mathcal{F})$ is a martingale if and only if there are no arbitrage opportunities.*
2. *There exists a probability measure ν, possibly inequivalent, such that $(S/S^0, \nu, \mathcal{F})$ is a martingale if and only if there are no strong arbitrage opportunities.*

Proof: One direction –necessity– is easy. If ν is a probability measure, let ϕ be any self-financing strategy with $V_T(\phi) > 0$. Then $E_\nu(V_T(\phi)/S^0_T) > 0$, so by theorem 1 $V_0(\phi) > 0$. Hence there can be no strong arbitrage.

If ν is equivalent to μ, and ϕ is a self-financing strategy such that $V_T(\phi) \geq 0$ and $\mu(V_T(\phi) > 0) > 0$, then again $E_\nu(V_T(\phi)/S^0_T) > 0$, so again by theorem 1 $V_0(\phi) > 0$. Hence there can be no arbitrage.

The other direction is harder. In each case we must prove the existence of an appropriate martingale measure.

First, note that the space of claims is just $\mathbb{R}^\Omega$, i.e. a claim is a list of payoffs, one per path. We make the following definitions:

$$\begin{aligned}
\mathcal{X}^+ &= \{X \in \mathbb{R}^\Omega \text{ such that } X \geq 0 \text{ and } \mu(X > 0) > 0\} \\
\mathcal{X}^{++} &= \{X \in \mathbb{R}^\Omega \text{ such that } X > 0\} \\
\mathcal{X}^0 &= \{X \in \mathbb{R}^\Omega \text{ such that } X = V_T(\phi) \text{ for } \phi \text{ self-financing with } V_0(\phi) = 0\}
\end{aligned}$$

We take up the proofs of sufficiency one at a time.

(1) The definition of no arbitrage is that $\mathcal{X}^+$ and $\mathcal{X}^0$ are disjoint. Note that $\mathcal{X}^+$ is exactly the closed positive orthant of $\mathbb{R}^\Omega$ minus the origin, and that $\mathcal{X}^0$ is a closed linear subspace. Hence we may apply Proposition 1(a) of the Appendix to conclude that there exists a vector $\vec{v} \in \mathcal{X}^{++}$ which is orthogonal to $\mathcal{X}^0$. This defines a linear functional λ on $\mathbb{R}^\Omega$ which is zero on $\mathcal{X}^0$ and positive on $\mathcal{X}^+$. We may normalize so that $\lambda(1) = 1$.

Such a linear functional uniquely determines a measure, which we denote by the same symbol λ, by means of the formula

$$\lambda(X) = \int_\Omega X \, d\lambda.$$

Positivity of λ means that, as a measure, λ is positive on each path in Ω, and hence is equivalent to μ.

Because λ vanishes on $\mathcal{X}^0$, it follows that

$$E_\lambda \left(\left. \frac{S_T^k}{S_t^k} \right| \mathcal{F}_t \right) = E_\lambda \left(\left. \frac{S_T^0}{S_t^0} \right| \mathcal{F}_t \right)$$

for all k. This results from following the strategy "wait until time t and, if the history is right, borrow \$1 of S^0 to buy \$1 of S^k." [**Exercise: fill in the details.**]

Because S_t^k is $\mathcal{F}_t$-measurable and positive, an equivalent form of this property is

$$0 = E_\lambda \left(\left. S_T^k - \frac{S_t^k}{S_t^0} S_T^0 \right| \mathcal{F}_t \right).$$

Notice that up to this point, we have made no use of the choice of S^0 as numeraire. Now that selection becomes important as we define ν by

$$\frac{d\nu}{d\lambda} = \frac{S_T^0}{\lambda(S_T^0)}.$$

(The denominator $\lambda(S_T^0)$ is simply a normalization factor to ensure that ν has total mass 1.)

Verifying the claimed martingale property is a piece of cake. Let Y be any $\mathcal{F}_t$-measurable function, and just check [**Excercise**] that

$$E_\nu \left(Y \left(\frac{S_T^k}{S_T^0} - \frac{S_t^k}{S_t^0} \right) \right) = \frac{E_\lambda \left(Y \left(S_T^k - \frac{S_t^k}{S_t^0} S_T^0 \right) \right)}{E_\lambda(S_T^0)} = 0.$$

(2) The definition of no strong arbitrage is that $\mathcal{X}^{++}$ is disjoint from $\mathcal{X}^0$. By appendix, there is λ which is positive on $\mathcal{X}^{++}$ and zero on $\mathcal{X}^0$. It looks like dot product with $\vec{v}$, which is positive with some possibly zero entries. This means the corresponding measure λ is non-negative but possibly inequivalent to μ. However, the remainder of the argument proceeds exactly as before.

∎

Evidently the real hero of the proof is the measure λ, which does not depend on a choice of numeraire, rather than the more popular ν. We know that ν-expectation of a discounted (with respect to S^0) claim gives the present value of the claim, but what can we say about λ? We will show that λ provides *forward prices* of claims consistent with no arbitrage. We will call λ the *forward measure*, and provide a full discussion in the section on forwards to follow.

The last part of the big theorem addresses the two sources of ambiguity in our pricing machinery, namely that our measure may not be unique, and we can only price attainable claims.

Definition 3 *We say our market is* complete *if every claim is attainable.*

Theorem 3 *Assume the market admits no arbitrage. Then there exists exactly one measure ν equivalent to μ such that $(S/S^0, \nu, \mathcal{F}_t)$ is a martingale if and only if the market is complete.*

Proof: Let $\mathcal{A}$ denote the set of all attainable claims, a linear subspace of $\mathbb{R}^\Omega$.

Suppose not every claim is attainable. Then $\dim\mathcal{A} < |\Omega|$. (Note this is a finite dimensional proof.) Since $\dim\mathcal{X}^0 < \dim\mathcal{A}$, the space $\mathcal{X}^0$ must have codimension greater than one. By the proof of the previous theorem, this means λ, and hence ν is not unique. [Exercise: uniqueness of ν implies uniqueness of λ.]

Conversely, suppose every claim is attainable: $\mathcal{A} = \mathbb{R}^\Omega$. Let ν_1, ν_2 be two martingale measures for S/S^0, and let X be any claim. Since $X \in \mathcal{A}$, X has a replicating strategy ϕ. By Theorem 1,

$$V_t(\phi) = S_t^0 E_{\nu_1}\left(\frac{X}{S_T^0}|\mathcal{F}_t\right) = S_t^0 E_{\nu_2}\left(\frac{X}{S_T^0}|\mathcal{F}_t\right)$$

for all t. Setting $t = 0$, we deduce that

$$E_{\nu_1}(Y) = E_{\nu_2}(Y)$$

for all claims Y. Hence $\nu_1 = \nu_2$.

■

It is quite possible for all claims to be attainable even though there is no martingale measure at all. In this case, however, there must be arbitrage opportunities by Theorem 2.

For non-attainable claims, when there is no arbitrage, there are at least arbitrage enforced bounds on the price. Namely, an upper bound is the infimum over all dominating attainable claims, and a lower bound is the supremum over all dominated attainable claims.

We can spell this out a little further. Theorem 3 says completeness is the same as uniqueness of the equivalent martingale measure. What if the market is not complete? Let Σ denote the set of all possible equivalent martingale measures, and let ν_1 and ν_2 be two different elements of Σ. For any attainable claim, lack of arbitrage determines a unique price at any time t. By theorem 1, this is equal to the conditional expectation of the discounted payoff with respect to either ν_1 or ν_2.

However, for a non-attainable claim, these conditional expectations need not agree. The choice of a measure in Σ corresponds to the choice of a consistent (no-arbitrage) way to price all claims, even the non-attainable ones.

For a given claim, the collection of all conditional expectations of discounted payoffs, as the measure ranges over Σ, is the set of all "fair" prices of the claim in the sense of no-arbitrage. [**Exercise: prove this.**]

1.1.6 Forwards

Many treatments of this subject miss the significance of the forward measure λ in contrast with the numeraire-dependent equivalent martingale measure ν. This is because many treatments make the extraneous assumption that the asset S^0 is deterministic. Indeed, under the weaker assumption that the numeraire S^0 is a bond maturing at time T (i.e. only that S^0_T is constant), we would have $\frac{d\nu}{d\lambda} = 1$, so $\lambda = \nu$. So the "simplifying" bond obscures the distinction between these generally different measures.

We call a claim paying 1 at time T in all states of the world a "T-bond". The forward measure is sometimes described as the equivalent martingale measure we get when we choose the T-bond for our numeraire. However, this presumes that the T-bond is attainable. In an incomplete market where the T-bond is not attainable, we still have forward measures but they are not unique.

The forward measure does depend on the time T to maturity, so is often called the "T-forward measure" λ_T.

Definition 4 *Given a claim X, time t with $0 \leq t < T$, and $\mathcal{F}_t$-measurable function K, the forward contract on X struck at K at time t is the claim $X - K$. The time t forward price of X is the strike K for which this forward contract has value 0 at time t.*

Intuitively, the forward price is the fair price, agreed upon at time t, to exchange for delivery of X at time T. It is somewhat disingenuous to speak of "the" forward price of an attainable claim X unless the T-bond is also attainable, since otherwise the value of K for which $X - K \cdot 1$ has present value zero depends on the choice of ν, hence λ_T. Nevertheless, ν provides *some* set of present values of all claims, even the non-replicable ones, and within this pricing framework, λ_T gives "the" forward prices of claims. In this case we call K above the λ_T-forward price.

Henceforth we bow to convention and assume that the T-bond is attainable in our market model.

Theorem 4 *In a viable market with an attainable T-bond, the time t λ_T-forward price of a claim X is $K = E_\lambda(X|\mathcal{F}_t)$. Moreover, if X is attainable, then K is independent of the choice of λ_T.*

Proof: Choose a numeraire S^0, and consider a measure ν making $(S/S^0, \nu, \mathcal{F})$ a martingale. (Implicit here is also the choice of the corresponding λ_T by $\frac{d\nu}{d\lambda} = \frac{S^0_T}{\lambda(S^0_T)}$.) Denote λ_T by *lambda*. For a claim X the forward price K should satisfy

$$\begin{aligned} 0 &= E_\nu\left((X-K)\frac{S^0_t}{S^0_T} \middle| \mathcal{F}_t \right) \\ &= \frac{E_\lambda\left(S^0_t(X-K) \middle| \mathcal{F}_t \right)}{E_\lambda(S^0_T|\mathcal{F}_t)} \\ &= \frac{S^t_0\left(E_\lambda(X|\mathcal{F}_t) - K\right)}{E_\lambda(S^0_T|\mathcal{F}_t)} \end{aligned}$$

and the theorem follows. We have made use of the fact that $\frac{d\nu}{d\lambda} = S^0_T/\lambda(S^0_T)$ and the general fact about conditional expectations

$$E_\nu(X|\mathcal{F}_t) = \frac{E_\lambda\left(X\frac{d\nu}{d\lambda} \middle| \mathcal{F}_t \right)}{E_\lambda(\frac{d\nu}{d\lambda}|\mathcal{F}_t)}.$$

For the statement about uniqueness of K, note that

$$K = \frac{E_\nu\left(X\frac{S^0_t}{S^0_T} \middle| \mathcal{F}_t \right)}{E_\nu\left(\frac{S^0_t}{S^0_T} \middle| \mathcal{F}_t \right)}$$

and the right hand side is unambiguous so long as X and 1 are replicable. ∎

An immediate consequence is that forward prices are independent of the choice of numeraire (since λ is). How can this be compatible with, say, the discussion in chapter 3 of Hull? There, it is found that the forward price K is $Se^{r(T-t)}$, where r is the risk free rate and S is the spot price, and this formula seems to have the discounting (and so the bond as numeraire) built inextricably into it. An explanation comes from our observation that $E_\lambda(\frac{S^k_T}{S^k_t}|\mathcal{F}_t)$ is independent of k. Calling this expectation $1/d(t,T)$, we find that the time t forward price of S^k is

$$E_\lambda(S^k_T|\mathcal{F}_t) = S^k_t/d(t,T),$$

i.e. spot over the same discount factor that applies to all securities in the given state of the world. In the event that one of the securities is a riskless bond (i.e. S^0_t is constant for each t), then this discount factor would be the usual S^0_t/S^0_T, but in general $d(t,T)$ varies from node to node at time t.

Now that we understand what λ is, it is instructive to revisit the proof that led to its discovery. The assumption of no arbitrage is that there be no free lottery tickets. Evidently the set $\mathcal{X}^+$ of non-negative, not identically zero claims is exactly the set of lottery tickets. The subspace $\mathcal{X}^0$ is the space of claims whose (arbitrage enforced) time 0 forward value is 0. At time 0, the risk-neutral investor will agree to receive at time T whatever gift or whipping a specified X in $\mathcal{X}^0$ might decree. The arbitrage enforcement of this value of 0 manifests in the following way. If an investor were willing to agree to pay \$1 at time T in exchange for X, a counterparty would immediately set up, for free, the replicating portfolio and maintain the strategy until, at time T, he presents the X, which cost him nothing but patience, for the agreed-upon \$1.

The useful fact about λ that all securities have the same forward return is an nice example. If S^0 had lower forward return than S^1, an arbitrageur would go long the forward contract on S^0, short \$1 of S_0, short the forward contract on S^1, and long \$1 of S^1. Then at time T he delivers the S^1 at the agreed-upon higher return (times \$1) than he pays for the S^0 he has agreed to buy (and return from the short sale at time 0). With this story in mind, we may also remark that the existence of risk-neutral discount factors $d(t,T)$ and $d(s,T)$ implies the existence of $d(s,t)$ for $s < t$. **Is that really true?**

1.1.7 Futures

Having just expounded the importance and numeraire indpendence of forward contracts, we digress briefly to contrast with the very fickle futures contracts. This topic is different from ones handled so far, not least because a futures contract entails cash flows not just at the horizon T, but at all intervening times as well. The next few paragraphs, which outline a naive but reasonable first attempt to define the problem, are intended to illustrate some of the pitfalls. [An alternate discussion follows.]

First, we should attempt to define what mean by futures. For forwards, we first defined "forward contract with strike K," then "forward price" as the strike giving present value zero to that contract. Futures work a bit differently.

Fix a claim X. Then the futures price process $F_t(X)$ is a certain sequence of random variables which we will attempt to determine more explicitly below. Now a futures contract on X may be entered into at any moment t at no charge and closed at any subsequent moment at no charge. At the end of each period during which the contract is open, say time $t+1$, the long position sees a cash flow of $F_{t+1}(X) - F_t(X)$ dollars. (The contract is *marked to market* after each period by having this amount added to or deducted from the contract holder's margin account.) If the contract is still open at time T, the long position actually receives X, so we must have $F_T(X) = X$. **Ahem?**

The defining property of the futures price F_t is that the expected value of discounted cash flows from the contract is zero.

The reader will note that we have not said anything about arbitrage. In fact we have been less than forthright by choosing a numeraire ("...discounted cash flows ...") without mentioning it. The first nasty shock about futures prices is that $F_t(X)$ depends on the choice of numeraire (in stark contrast to forward prices). In view of the arbitrage pricing theorem, one might have hoped that the defining property of $F_t(X)$ is numeraire-invariant and that we were just being sloppy, but this is not the case.

For the immediate purposes of this discussion only, we will need our numeraire to be previsible.

Definition 5 *A money market account is a security S whose value S_t is $\mathcal{F}_{t-1}$-measurable.*

Note that a money market account still allows for random interest rates. The primary effect of assuming the existence of a money market account is that the 1-period forward rates of return $E_\lambda(\frac{S^k_{t+1}}{S^k_t}|\mathcal{F}_t)$ are determined by the money market account and are independent of the choice of λ. **(Is that true? Is there a consistency necessary between bonds and money market?)**

Theorem 5 *Assume that S^0 is a money market account and that ν is an equivalent measure such that $(S/S^0, \nu, \mathcal{F})$ is a martingale. For a claim X, the futures prices process $F_t(X)$ (with respect to the numeraire S^0) is $E_\nu(X|\mathcal{F}_t)$.*

Proof: Consider the trading strategy of waiting until time t, then at each step τ going long ϕ_τ futures contracts, closing the contracts from the step before and investing the proceeds (or loss) in units of S^0. We'll pick a clever set of ϕ_τ's once we get a feel for the outcome. This strategy is self-financing since it costs nothing to enter or leave futures contracts. At time T we will hold a certain number of units of S^0, namely

$$\frac{\phi_t}{S^0_{t+1}}(F_{t+1} - F_t) + \frac{\phi_{t+1}}{S^0_{t+2}}(F_{t+2} - F_{t+1}) + \cdots = \sum_{\tau=t+1}^{T} \frac{\phi_{\tau-1}}{S^0_\tau} \Delta F_\tau.$$

To be clear on this, at, say, time $t+1$ we receive a payoff on ϕ_t contracts, which is $\phi_t \Delta F_{t+1}$ dollars, or equivalently $\frac{\phi_t}{S^0_{t+1}} \Delta F_{t+1}$ units of S^0. Our clever choice of $\phi_{\tau-1} = S^0_\tau$ is permissible precisely because S^0 is a money market account.

Now we have that ϕ is self-financing, $V_t(\phi) = 0$ and $V_T(\phi) = (F_T(X) - F_t(X))S^0_T$. Applying the arbitrage pricing theorem and $F_T(X) = X$ yields,

$$0 = E_\nu \left((X - F_t) S^0_T \frac{S^0_t}{S^0_T} \middle| \mathcal{F}_t \right)$$

and the claim follows immediately. ■

For our discrete time setting, it was crucial for the proof that the numeraire be a money market account. From a real world standpoint, this is probably not such a bad thing. It does, however, raise the question of what the correct statement is for an arbitrary numeraire. In continuous time, remarkably enough, the proof goes through for arbitrary numeraire (see Elliot & Kopp p. 216-7), essentially because the distinction between S^0_t and S^0_{t+1} obstructing the discrete version goes away.

Here is an alternate discussion.

We continue to assume our market admits no arbitrage, but we allow it to be incomplete. A futures contract on a claim X, maturing at time T, will be characterized by a process F_t defined below, intended to represent the "futures price" of X at any time t. This process will depend not only on the claim X and the maturity T, but also on the numeraire (which we must take to be a money market account), and also on the choice of equivalent martingale measure ν. (Of course, if the market is complete, there is only one such measure.)

Theorem 6 *Suppose S^0 is a money market account and ν is a measure equivalent to μ such that $(S/S^0, \nu, \mathcal{F})$ is a martingale. Let X be a claim. Then there exists a unique $\mathcal{F}$-adapted process F_t such that*

1. *$F_T = X$, and*
2. *for all s, t such that $0 \leq s < t \leq T$, $E_\nu((F_t - F_{t-1})\frac{S^0_s}{S^0_t}|\mathcal{F}_s) = 0$.*

The second condition means that at any time s, the ν-value of the future cash flows of the contract is zero.

Proof:

Definition 6 *The process F_t of the previous theorem is called the T-futures prices of X at time t relative to S^0 and ν.*

Theorem 7 *Suppose S^0 is a money market account and ν is a measure equivalent to μ such that $(S/S^0, \nu, \mathcal{F})$ is a martingale. Let X be a claim, and let F_t be the T-futures prices of X at time t relative to S^0 and ν.*

Then $F_t = E_\nu(X|\mathcal{F}_t)$.

Proof:

Remember that if our market is complete, then the ν-value of a claim is simply the unique no-arbitrage price of the claim. Also, from the proof of theorem 2, we can see that if the money market account is assumed deterministic, then $\nu = \lambda$, and so the futures price is simply equal to our friend the forward price.

1.2 Morals and Examples

We include a few examples to illustrate the general framework and show how the various hypotheses interact in the context of a stock lattice.

1.2.1 What No Arbitrage Means

Warning: this section is not quite right.

No arbitrage, in the tree context, means that at each node, failure of a security to be an independent source of noise determines the price of that security. In particular, no arbitrage is a local condition. Let's try to make that more precise.

Choose a node $F_0 \in \mathcal{F}_t$, and choose some ordering $(F_1, \ldots, F_n)$ of the nodes which branch from it. So F_0 is the disjoint union of the other F_i, which are themselves in $\mathcal{F}_{t+1}$. Consider now the matrix $[\mathcal{S}^0 \ldots \mathcal{S}^k]$ whose i^{th} column is the possible outcomes of the i^{th} sectурity,

$$\mathcal{S}^i = (S^i_{t+1}(F_1), \ldots, S^i_{t+1}(F_n))$$

and the vector of current stock prices $\mathcal{S} = (S^0_t(F_0), \ldots, S^k_t(F_0))$. The assumption of no arbitrage at this node is that

$$V_{t+1}(\phi) = [\mathcal{S}^0 \ldots \mathcal{S}^k]\vec{\phi} = 0 \Rightarrow V_t(\phi) = \mathcal{S} \cdot \vec{\phi} = 0.$$

If a linear relation exists among the $\mathcal{S}^i$, then a portfolio consisting one of the S^i at time $t+1$ can be replicated with a portfolio of the others. So $\mathcal{P}_0$ and $\mathcal{P}_1$ definitely have the same value at time $t+1$. If the same relation fails to persist among the securities at the earlier time t node, let's assume $V_t(\mathcal{P}_0) = \alpha V_t(\mathcal{P}_0)$ for some $\alpha > 1$. Then our trading strategy is, in the event the given node is reached, form the zero value portfolio $\alpha\mathcal{P}_1 - \mathcal{P}_0$. Then at time t rebalance to hold $(\alpha - 1)\mathcal{P}_1$, and this portfolio will remain positive in value.

Suppose there is an arbitrage ϕ. Then $V_0(\phi)$ is identically 0, but $V_T(\phi)$ is not. So there is a smallest $t > 0$ for which $V_t(\phi)$ is not identically 0. Since it is $\mathcal{F}_t$-measurable, there is a node $F_t \in \mathcal{F}_t$ on which it is constant and positive. Then the node one step before, namely the unique F_{t-1} in $\mathcal{F}_{t-1}$ containing F_t is vulnerable to the node-wise arbitrage discussed above. So "no arbitrage" is a local condition.

1.2.2 More Securities Than Branches

If there are $k+1$ securities and some node has less than $k+1$ branches, typically that node yields an arbitrage and there is no martingale measure. Usually all claims are replicable in this case.

To see how this happens, consider the simplest case of a single binomial step and three securities. So at time 0, the security vector is $\vec{S}_0 = (R_0, S_0, T_0)$, and at time 1, it is either $\vec{S}_u = (R_u, S_u, T_u)$ or $\vec{S}_d = (R_d, S_d, T_d)$. Assuming some independence, there will be α and β such that

$$T_u = \alpha R_u + \beta S_u$$

$$T_d = \alpha R_d + \beta S_d.$$

Except in the unlikely event that $T_0 = \alpha R_0 + \beta S_0$, there is an arbitrage and no martingale measure.

1.2.3 More Branches Than Securities

If there are $k+1$ securities and some node has more than $k+1$ branches, typically there is no arbitrage, there is a non-unique martingale measure, and not all claims are replicable.

The simplest example is a single trinomial branch with two securities. We won't bother with the details.

1.2.4 Unique Martingale Measure

Finally, we ask what conditions must be imposed upon a stock lattice for it to admit a unique martingale measure and replicability of all claims. As the preceding examples make clear, this must be a regular k-nomial tree with k securities representing k independent sources of noise at each node. In particular, the two security binomial tree, which is usually presented as a typical case conveying the correct flavor of the general case, is very special indeed and not at all representative of general disrete time stock processes.

Question: Can B&R's bookie parable be seen as $k + 1$ securities (horses) with $k + 1$ outcomes (possible winners) at $T = 1$?