Stochastic Calculus & Finance: The Heath-Jarrow-Morton Model

Dhruv Ajmera

The University of Texas at Austin

July 2026

What will be talked about?

Today, we will be talking about this stochastic differential equation:

What will be talked about?

Today, we will be talking about this stochastic differential equation:

$$df(t,T) = \alpha(t,T)\,dt + \sigma(t,T)\,dW(t),$$

where $f(t,T)$ is the instantaneous rate for time $T > t$ that can be locked in at time t.

What will be talked about?

Today, we will be talking about this stochastic differential equation:

$$df(t,T) = \alpha(t,T)\,dt + \sigma(t,T)\,dW(t),$$

where $f(t,T)$ is the instantaneous rate for time $T > t$ that can be locked in at time t. This is the Heath-Jarrow-Morton model, meant to model forward rates.

Below is the rough ordering of topics that will be covered:

Overview

Below is the rough ordering of topics that will be covered:

- Measures
- Sigma-Algebras, Random Variables, and Information
- Random Walks and Brownian Motion
- Putting it all Together

The Canonical Example

To ground the theory that will be introduced, we will always return to the example of the so-called coin toss space:

The Canonical Example

To ground the theory that will be introduced, we will always return to the example of the so-called coin toss space:

$$\Omega = \{-1, 1\}^{\mathbb{N}}$$

The Canonical Example

To ground the theory that will be introduced, we will always return to the example of the so-called coin toss space:

$$\Omega = \{-1, 1\}^{\mathbb{N}}$$

For example, one element of this space would be

$$(1, 1, 1, \dots),$$

which represents a countably infinite chain of heads.

Measures: Assigning Size to Events

In probability, a measure tells us how large or likely a set of outcomes is.

Measures: Assigning Size to Events

In probability, a measure tells us how large or likely a set of outcomes is.

An event is a subset of Ω. For example:

$$A = \{\omega \in \Omega : \omega_1 = 1\}$$

is the event that the first toss is heads.

Measures: Assigning Size to Events

In probability, a measure tells us how large or likely a set of outcomes is.

An event is a subset of Ω. For example:

$$A = \{\omega \in \Omega : \omega_1 = 1\}$$

is the event that the first toss is heads.

In particular, a probability measure $\mathbb{P}$, requires that $\mathbb{P}(\Omega) = 1$.

Why Do We Need a Sigma-Algebra?

Since not every subset of Ω is necessarily meaningful for our model, we choose a collection of observable events, called a sigma-algebra, $\mathcal{F}$.

Why Do We Need a Sigma-Algebra?

Since not every subset of Ω is necessarily meaningful for our model, we choose a collection of observable events, called a sigma-algebra, $\mathcal{F}$.

Definition 2

Formally, let Ω be a nonempty set, and $\mathcal{F}$ a collection of subsets of Ω. $\mathcal{F}$ is a sigma-algebra if:

1. $\emptyset \in \mathcal{F}$
2. $\forall A \subseteq \Omega,$ s.t. $A \in \mathcal{F} \implies A^{\complement} \in \mathcal{F}$
3. $A_1, A_2, \cdots \in \mathcal{F} \implies \bigcup_{n=0}^{\infty} A_n \in \mathcal{F}$

Sigma Algebras (cont.)

Intuitively, $\mathcal{F}$ is designed to contain all events whose probabilities make sense in our model.

Sigma Algebras (cont.)

Intuitively, $\mathcal{F}$ is designed to contain all events whose probabilities make sense in our model.

Together, the triple

$$(\Omega, \mathcal{F}, \mathbb{P})$$

is called a probability space, and it is the backbone of probability theory.

A random variable is a function from the outcome space to the real numbers:

$$X : \Omega \to \mathbb{R}.$$

Random Variables

A random variable is a function from the outcome space to the real numbers:

$$X : \Omega \to \mathbb{R}.$$

For example, define

$$X_i(\omega) = \omega_i.$$

This extracts the result of the ith coin toss.

Random Variables

A random variable is a function from the outcome space to the real numbers:

$$X : \Omega \to \mathbb{R}.$$

For example, define

$$X_i(\omega) = \omega_i.$$

This extracts the result of the ith coin toss.

So if

$$\omega = (1, -1, 1, 1, \dots),$$

then

$$X_1(\omega) = 1, \quad X_2(\omega) = -1.$$

Information Over Time

After n coin tosses, we know only the first n coordinates:

$$X_1, X_2, \ldots, X_n.$$

Information Over Time

After n coin tosses, we know only the first n coordinates:

$$X_1, X_2, \ldots, X_n.$$

The sigma-algebra generated by this information is written

$$\mathcal{F}_n = \sigma(X_1, \ldots, X_n).$$

Information Over Time

After n coin tosses, we know only the first n coordinates:

$$X_1, X_2, \ldots, X_n.$$

The sigma-algebra generated by this information is written

$$\mathcal{F}_n = \sigma(X_1, \ldots, X_n).$$

The sequence

$$\mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \cdots$$

is called a filtration.

Information Over Time

After n coin tosses, we know only the first n coordinates:

$$X_1, X_2, \ldots, X_n.$$

The sigma-algebra generated by this information is written

$$\mathcal{F}_n = \sigma(X_1, \ldots, X_n).$$

The sequence

$$\mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \cdots$$

is called a filtration. A filtration is the mathematical model of "what we know so far."

Information Over Time (cont.)

Definition 3

Formally, a filtration is a time-indexed sequence of σ-algebras $\mathcal{F}(t)$, $t \in [0, T]$ on a nonempty Ω, such that, for $s \leq t$,

$$\mathcal{F}(s) \subseteq \mathcal{F}(t)$$

Why This Matters for Finance

In finance, prices and interest rates evolve as new information enters the market.

Why This Matters for Finance

In finance, prices and interest rates evolve as new information enters the market.

The filtration $\mathcal{F}_t$ represents all market information available up to time t.

Why This Matters for Finance

In finance, prices and interest rates evolve as new information enters the market.

The filtration $\mathcal{F}_t$ represents all market information available up to time t.

A model should not depend on future information. So quantities like

$$f(t,T), \quad \sigma(t,T), \quad W(t)$$

must be based only on information available at time t.

Random Walk: Accumulating Randomness

From the coin toss space, define

$$X_i(\omega) = \omega_i,$$

where each X_i is either -1 or 1 with probability $1/2$.

Random Walk: Accumulating Randomness

From the coin toss space, define

$$X_i(\omega) = \omega_i,$$

where each X_i is either -1 or 1 with probability $1/2$.

The symmetric random walk is defined as

$$S_n = X_1 + X_2 + \cdots + X_n.$$

Random Walk: Accumulating Randomness

From the coin toss space, define

$$X_i(\omega) = \omega_i,$$

where each X_i is either -1 or 1 with probability $1/2$.

The symmetric random walk is defined as

$$S_n = X_1 + X_2 + \cdots + X_n.$$

So S_n records the cumulative position after n random steps.

From Random Walk to Brownian Motion

A random walk moves in discrete time:

$$S_0, S_1, S_2, \ldots$$

From Random Walk to Brownian Motion

A random walk moves in discrete time:

$$S_0, S_1, S_2, \dots$$

However, financial models usually evolve in continuous time.

From Random Walk to Brownian Motion

A random walk moves in discrete time:

$$S_0, S_1, S_2, \ldots$$

However, financial models usually evolve in continuous time.

Thus, we must adapt this process to make use of it for our modeling purposes.

The Transition to Continuous

First, we define the scaled symmetric random walk.

First, we define the scaled symmetric random walk.

Definition 5

Fix some $n \in \mathbb{Z}^+$. Then, the scaled symmetric random walk is defined as

$$W^{(n)}(t) = \frac{1}{\sqrt{n}} S_{nt},$$

where S_{nt} is a regular symmetric random walk, provided $nt \in \mathbb{Z}$.

The Transition to Continuous (cont.)

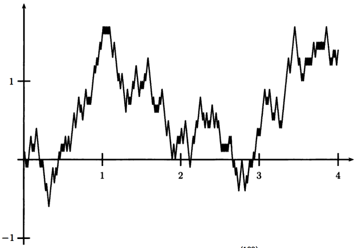

Fig. 3.2.2. A sample path of $W^{(100)}$.

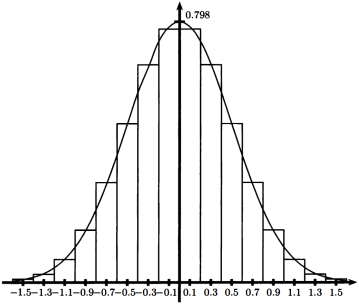

Fig. 3.2.3. Distribution of $W^{(100)}(0.25)$ and normal curve $y = \frac{2}{\sqrt{2\pi}}e^{-2x^2}$.

The Transition to Continuous (cont.)

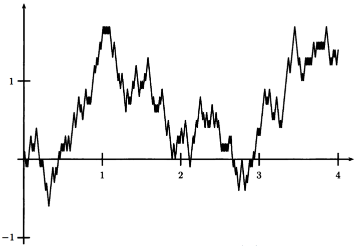

Fig. 3.2.2. A sample path of $W^{(100)}$.

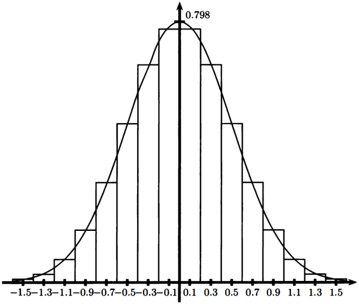

Fig. 3.2.3. Distribution of $W^{(100)}(0.25)$ and normal curve $y = \frac{2}{\sqrt{2\pi}} e^{-2x^2}$.

As we make the increments and time steps smaller, the path becomes increasingly jagged, with increments approaching a normal distribution.

Obtaining the Brownian Motion

We then take the limiting object of these scaled random walks (proven rigorously by taking the limit as $n \to \infty$ and using CLT.)

Obtaining the Brownian Motion

We then take the limiting object of these scaled random walks (proven rigorously by taking the limit as $n \to \infty$ and using CLT.)

By doing so, we obtain Brownian Motion.

Obtaining the Brownian Motion

We then take the limiting object of these scaled random walks (proven rigorously by taking the limit as $n \to \infty$ and using CLT.)

By doing so, we obtain Brownian Motion.

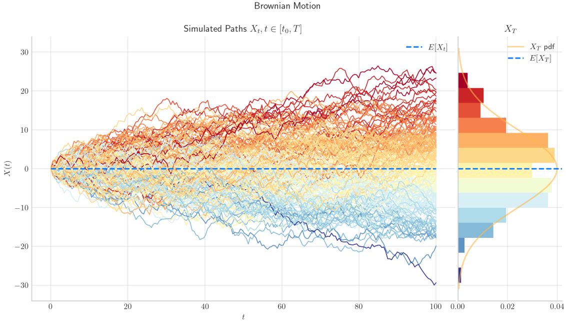

Source: quantgirluk.github.io/Understanding-Quantitative-Finance

Definition 6

Formally, a continuous function $W(t)$ dependent on $\omega \in \Omega$ is a Brownian Motion if, for all $0 = t_0 < \cdots < t_m$, the increments

$$W(t_1) = W(t_1) - W(t_0), \ldots, W(t_m) - W(t_{m-1})$$

are independent and normally distributed. In addition,

$$E[W(t_{i+1}) - W(t_i)] = 0,$$

$$Var[W(t_{i+1}) - W(t_i)] = t_{i+1} - t_i.$$

Now, we can come back to the HJM Model.

$$df(t,T) = \alpha(t,T)\,dt + \sigma(t,T)\,dW(t),$$

Now, we can come back to the HJM Model.

$$df(t,T) = \alpha(t,T)\,dt + \sigma(t,T)\,dW(t),$$

Under this model, the forward rate is driven by a Brownian Motion, with rates being determined by a drift term α and a volatility term σ.

Now, we can come back to the HJM Model.

$$df(t,T) = \alpha(t,T)\,dt + \sigma(t,T)\,dW(t),$$

Under this model, the forward rate is driven by a Brownian Motion, with rates being determined by a drift term α and a volatility term σ.

Although expressing the solution to f is possible in terms of an ordinary Lebesgue integral and an Ito Integral (for the Brownian Motion), in practice, numerical methods are often needed for evaluation.

Practical Implementation

The model takes in historical observations as inputs.

Practical Implementation

The model takes in historical observations as inputs.

Specifically, we need observations of the forward rate for various different times in the past for various different relative maturities. In addition, we also require incremental observations for some time offset δ.

Practical Implementation

The model takes in historical observations as inputs.

Specifically, we need observations of the forward rate for various different times in the past for various different relative maturities. In addition, we also require incremental observations for some time offset δ.

Technically, this only gives us enough information to estimate σ, but through a technique called change of measure, we do not require α.

A Coded Example

Below is a coded implementation of the HJM model on a dummy initial forward rate curve

$$f(0,T) = 0.05e^{-0.3T}.$$

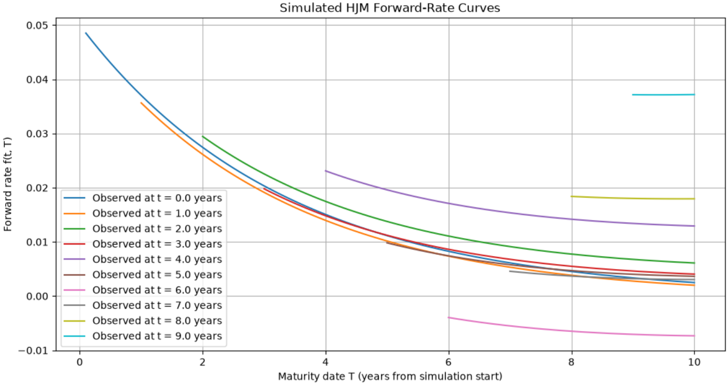

References

[1] Steven E. Shreve. *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer Finance, Springer, 2004.