

Algorithmica Research AB

Heston Stochastic Volatility Model Demo workspace walkthrough

Last update: 2021-03-02

© 2002-2021 Algorithmica Research AB. All rights reserved.

Table of Contents

The Heston Stochastic Volatility Model	2
The five parameters of the Heston Model.....	3
The role of the different parameters.....	3
v_0	3
θ	3
ρ	3
κ	4
σ	4
The Feller Condition.....	4
Workspace tab “Calibration of the Heston Model”	5
The Calibration Class.....	8
Workspace Tab “Heston Option Pricing”	10
The Option Pricing and Implied Volatility calculation functions.....	11
Workspace Tab “Heston Implied Volatility Smile”	13
Workspace Tab “Option Price Sensitivity to Heston Parameters”	14
Workspace Tab “Black Volatility Sensitivity to Heston Parameters”	15
Workspace Tab “Simulating Heston, one path”	16
The Simulation Class	16
Workspace Tab “Simulating Heston, ten paths”	19
Workspace Tab “Simulated Vanilla Option Prices	20
Workspace Tab “Simulated Barrier Option Prices”	21

The Heston Stochastic Volatility Model

The Heston Stochastic volatility model was introduced by Steven Heston in 1993 in an attempt to provide a more realistic model for assets than the well known Black-Scholes (BS) model.

Simply put, it replaces the constant volatility of the BS model by a stochastic volatility which is modelled by a separate process.

The processes are given by

$$\frac{dS(t)}{S(t)} = \mu(t)dt + \sqrt{V(t)}dW_1(t)$$
$$dV(t) = \kappa(\theta - V(t))dt + \sigma\sqrt{V(t)}dW_2(t),$$

where $W_1(t)$ and $W_2(t)$ are standard Brownian motions with instantaneous correlation given by

$$\langle dW_1, dW_2 \rangle = \rho dt$$

and $\mu(t)$ is the drift term, typically the risk free rate (with any dividend yield subtracted).

The Heston model depends on five free parameters:

The five parameters of the Heston Model

v_0 = The initial value for the variance process $V(t)$, $v_0 = V(0)$.

θ = The long term mean for the variance process $V(t)$.

ρ = The instantaneous correlation between the Brownian motions for the asset and variance process.

κ = The mean reversion speed. How fast the variance process tends to its long term mean.

σ = The volatility of the variance process. Often called vol-vol.

Those parameters have to be calibrated from available market data (in the form of Black volatilities for different maturities and strikes).

The role of the different parameters

v_0

Controls the initial vertical position of the implied Black volatility smile. Larger value means larger initial volatilities.

Allowed values: All positive values are theoretically allowed, even though values too close to zero is not recommended. Typically this value is less than 1.

θ

Controls the vertical position of the long term implied Black volatility smile. Larger value means larger long term volatilities.

Allowed values: All positive values are theoretically allowed, even though values too close to zero is not recommended. Typically this value is less than 1.

ρ

Controls the skew of the smiles. The value 0 gives the most symmetrical smile. Negative and positive values skews the smile in different directions.

Allowed values: All values (strictly) between -1 and 1 theoretically allowed, even though values too close to -1 or 1 are not recommended. In real application this correlation is typically negative.

κ

A larger κ means a flatter smile since the variance process then tends to follow its mean more closely.

Allowed values: All positive values are theoretically allowed, even though values too close to zero is not recommended. Also too large values of this parameter is not recommended. By large we here mean around 10 or higher.

The reason is that the model in general becomes very insensitive to changes in this parameter for larger values.

σ

Controls the kurtosis of the smile, or more simply put the “happiness” of the smile. A larger sigma gives a more pronounced smile. A smaller sigma gives a flatter smile.

Allowed values: All positive values are theoretically allowed, even though values too close to zero is not recommended. Typically this value is less than 3.

Remark:

We may note that the effect of the parameters κ and σ both overlap a bit. We can make the smile flatter by either increasing κ or by decreasing σ . This makes the calibration a bit numerically problematic since a large κ combined with a large σ can produce a very similar looking smile compared to a small κ combined with a small σ .

This is not a huge problem in practice though, since we are not really interested in the Heston parameters as such, but rather how well the implied volatility smile fits the market data.

The Feller Condition

The variance process can hit zero if the so called Feller condition

$$2\kappa\theta > \sigma^2$$

is not satisfied. It is not an absorbing state though, so the process don't stick to zero in that case.

The Feller condition is often fulfilled in real application, but sometimes it is not.

This condition is not a requirement in any way, but in general the simulation of the Heston model will typically become less accurate if this condition is violated by a large margin.

Workspace tab “Calibration of the Heston Model”

This is a screenshot of the contents of the table in first tab in the Workspace. The purpose of this tab is to demonstrate the calibration of the Heston model using some simplified input.

Remark: When using the workspace for the first time. Press recalc twice (after providing the input).

Common parameters	Strikes (% of forward value)	Black Volatility, maturity 1	Black Volatility, maturity 2	Black Volatility, maturity 3	Heston parameter names	Calibrated values of Heston parameters
Calibration Target: option_price	50%	0.500000000	0.400000000	0.370000000	v0	0.233071205
Calibration Weights: vega_weights	60%	0.450000000	0.420000000	0.400000000	theta	0.101285271
Initial Guess v0: 0.2	70%	0.440000000	0.390000000	0.400000000	rho	-0.215320939
Initial Guess theta: 0.2	80%	0.400000000	0.370000000	0.340000000	kappa	3.948729490
Initial Guess rho: 0	90%	0.380000000	0.360000000	0.340000000	sigma	1.295816663
Initial Guess kappa: 1	100%	0.370000000	0.330000000	0.350000000		
Initial Guess sigma: 1	110%	0.350000000	0.320000000	0.330000000		
Time to maturity 1: 0.5	120%	0.360000000	0.320000000	0.320000000		
Time to maturity 2: 1	130%	0.360000000	0.330000000	0.300000000		
Time to maturity 3: 1.5	140%	0.390000000	0.340000000	0.310000000		
Forward value 1: 100	150%	0.400000000	0.350000000	0.310000000		
Forward value 2: 100						
Forward value 3: 100						
Recalc: [Recalc]						
	Implied Black Vol from Heston					
	50%	0.499312448	0.426174046	0.390479138		
	60%	0.463006355	0.400044503	0.371317927		
	70%	0.431308131	0.379829588	0.356090099		
	80%	0.404467801	0.361592081	0.344217871		
	90%	0.383516059	0.348828859	0.335351468		
	100%	0.369725436	0.340209478	0.329137188		
	110%	0.363520831	0.335378419	0.325207164		
	120%	0.363732480	0.333715465	0.323165186		
	130%	0.368254464	0.334463068	0.322615582		
	140%	0.375173420	0.336897786	0.323197509		
	150%	0.383209250	0.340433123	0.324606748		

The calibrated Heston parameters are seen in the two columns on the far right.

Column 2 to 4 contains Black volatilities (expressed as absolute numbers, that is 0.45 for 45% volatility and so on) used as input for the calibration. The upper part of the table of these three columns have editable entries where the user can change the values.

The lower part of the same columns contains the calculated implied Black volatilities using the calibrated parameters in the last column. The point of this is to provide a reference to compare with the upper part of the table. The better match the better the calibrations.

Since the model only have five parameters to calibrate we can of course not expect a very good match to market data here.

The input parameters expected from the user(except the Black volatilities) are on the left side of the table:

Calibration Target :

This specifies an enum:

```
// Enum to describe what target quantity should be matched in the Heston calibration.
// We can compare option prices or implied Black volatilities.
enum calibration_target_option(category : "Finance/Heston-Model") { OPTION_PRICE_option(str : "option_price"), VOLATILITY_option(str : "volatility");
```

Simply put, in our calibration procedure we can choose to compare market and model option prices or implied volatilities.

Calibration weighting :

This specifies an enum:

```
// Enum to describe how we choose the weights.
// If custom weights are chosen, all those those custom weights needs to be provided.
enum calibration_weights_option(category : "Finance/Heston-Model") { CUSTOM_option(str : "custom"),
..... EQUAL_WEIGHTS_option(str : "equal_weights"),
..... EQUAL_TOTAL_WEIGHT_PER_MATURITY_option(str : "equal_total_weight_per_maturity"),
..... VEGA_WEIGHTS_option(str : "vega_weights");
```

Internally, we use a least squares optimizer where we minimize the sum of square errors of model option prices to market option prices (or volatilities if that option is chosen).

We can choose to put weights in front of these square terms. The choices are the following:

CUSTOM = The user can specify a vector of weights for each option freely when adding market data to the calibration class. This is not used in this workspace though, so this option has been excluded here.

EQUAL_WEIGHTS = No weights/all weights equal to 1.

EQUAL_TOTAL_WEIGHT_PER_MATURITY = The sum of all weights for all terms belonging to options for a specific maturity adds up to the same number for all maturities. This can be useful if we have a different number of options for each maturity, but we want each maturity be equally important.

VEGA_WEIGHTS = Here we construct weights based on the option Vega. This is only used when we compare option prices. Using these weights emulates the situation if we would compare volatilities with equal weights instead.

This is typically the desired choice as it gives a very good fit generally.

Technical remark:

More precisely. For each option define

$$v_i = \min\left(\frac{1}{V_i}, \frac{1000}{F_i}\right)$$
$$w_i = \frac{v_i^2}{\sum_j v_j^2}$$

where the sum is over all options (numbered by j) and V_i is the vega of option i and F_i is the forward value of the asset for the maturity of option i.

w_i is now the weight we use for option i.

Initial guess v0, Initial guess theta, Initial guess rho, Initial guess kappa, Initial guess sigma :

The optimizing procedure for the Heston model takes an initial guess for the Heston parameters as starting point. It can happen that the calibration procedure “get lost” if this guess is chosen badly. The above entries are the initial guesses for all the Heston variables.

The above starting values are quite typical and typically there is no need to change them.

A few loose guidelines for choosing “good” values are:

The initial v0, theta, kappa or sigma should not be chosen too close to 0 (the end point of their domain of definition).

Rho must be between -1 and 1, but the starting guess should not be too close to the endpoints.

v0, theta, kappa or sigma should not be chosen too large as starting guess. v0, theta and sigma should preferably be chosen less than 1 and kappa less than say 5.

Time to maturity 1, Time to maturity 2, Time to maturity 3 :

The table has editable cells where the user can enter the Black volatilities used as input for the calibration. We use three columns here, each one representing data for one maturity.

Those three parameters tells the calibration procedure which maturities the black volatilities of the three columns represent.

Forward value 1, Forward value 2, Forward value 3 :

Here the forward values for the corresponding maturities are specified.

The Calibration Class

We now aim to describe the code used behind the scenes to create this tab and to calibrate the Heston model in general.

Create an instance of the calibration class:

```
heston_calibration h = heston_calibration();
```

Now assume we have some market data at the maturity $T=1$. Assume the forward value for the underlying to maturity T is $F=100$. Assume we have strikes at 80, 100 and 120 and the corresponding Black volatilities 45%, 40%, 41%.

This we can add to our calibration class:

```
number T = 1;  
number F = 100;  
vector(number) strikes = [70, 100, 120];  
vector(number) black_vol = [0.45, 0.40, 0.41];  
  
h.add_data_for_maturity(T, F, strikes, black_vol);
```

This is then repeated for all maturities T , where we have market data.

At least two maturities, and preferably more, should be used. This is needed to reflect the time behavior of the model and its parameters.

We also have an alternative version of the above function. We can decide how much weight we should put in front of each squared error term in the least squares procedure. In this way we can explicitly tell the class that some option data are more important than others to match.

Take the same example as before, but want the options to be weighted by 0.5, 1 and 2 respectively.

Then we specify:

```
number T = 1;  
number F = 100;  
vector(number) strikes = [70, 100, 120];  
vector(number) black_vol = [0.45, 0.40, 0.41];  
vector(number) weights = [0.5, 1, 2];  
  
h.add_data_for_maturity(T, F, strikes, black_vol, weights);
```

If such custom weights are used for one maturity it is required to use custom weights for all maturities.

When all data has been added we start the calibration we simply calling the function “calibrate”. There are two versions of this. One which defaults all arguments of the other to suitable values. The signatures of the functions are

```
heston_params calibrate();
heston_params calibrate(calibration_target target,
                        calibration_weights weight_type,
                        vector(number) start_guess,
                        number option_val_rel_tol = 0.0000000001,
                        number alpha_search_abs_tol = 0.000001);
```

The first two arguments are exactly those enums just described for the workspace. If custom weights have been used to fill the class with market data, the input argument **target** must equal CUSTOM. **start_guess** = [v0, theta, rho, kappa, sigma] is a vector containing the start guess of the Heston parameters as described earlier, in the given order.

option_val_rel_tol is a relative error tolerance for the internal option price calculation procedure. This is defaulted if not given.

alpha_search_abs_tol is another internal absolute error tolerance when searching for a dampening parameter alpha used in the option pricing. This is defaulted and should probably be left that way.

So let us choose the initial guesses as v0=0.2, theta = 0.3, rho=0, kappa = 1 and sigma = 0.8. Let us assume we want to compare option prices using vega weighting :

```
vector(number) start_guess = [0.2, 0.3, 0, 1, 0.8];
heston_params result = h.calibrate(OPTION_PRICE, VEGA_WEIGHTS, start_guess);
```

We get a Heston_params object as a result. This contains our calibrated parameters. We can ask for the calibrated parameters by using its member functions :

```
number v0 = result.v0();
number theta = result.theta();
number rho = result.rho();
number kappa = result.kappa();
number sigma = result.sigma();
```

Workspace Tab “Heston Option Pricing”

This is a screenshot of the second tab which is used to demonstrate option pricing in the Heston model.

There are three tables. One table for pricing a single option and calculate its implied Black volatility, one table showing a grid of option prices for different strikes and maturities and one table showing a grid of implied volatilities for different strikes and maturities.

The screenshot displays the Heston Option Pricing workspace with three main tables and a parameter input section.

Single Option Price and Implied Volatility

Common parameters	Heston Option Price	Implied Black Volatility
v_0 0.2	13.556440749	0.341481221

Heston Implied Black Volatilities

Common parameters	Strikes (% of forward value)	Back Volatility, Maturity 1	Back Volatility, Maturity 2	Back Volatility, Maturity 3
v_0 0.2	50%	0.45026920	0.44832373	0.43034362
θ 0.3	60%	0.45730598	0.45344090	0.43131378
ρ 0	70%	0.43081299	0.43965273	0.40437870
κ 1	80%	0.41916934	0.39440217	0.39787091
σ 1.2	90%	0.39772225	0.36119568	0.36368977
Time to Maturity 1 1	100%	0.39244918	0.36027006	0.35254944
Time to Maturity 2 1	110%	0.39777834	0.36162609	0.35447860
Time to Maturity 2 2	120%	0.40941062	0.36978944	0.36979181
Time to Maturity 2 3	130%	0.41583969	0.40131051	0.39913659
Time to Maturity 3 1	140%	0.42762626	0.40773057	0.40146876
Time to Maturity 3 3	150%	0.43862614	0.44604968	0.40746671

Option Prices (Lambert-Contract)

Common parameters	Strikes (% of forward value)	Option Price, Maturity 1	Option Price, Maturity 2	Option Price, Maturity 3
v_0 0.2	50%	51.400987940	54.19232820	57.037024948
θ 0.3	60%	42.833013480	46.942772937	50.786237750
ρ 0	70%	35.68392080	40.541874269	45.26969729
κ 1	80%	28.324375774	34.988948820	40.387939195
σ 1.2	90%	22.648212847	30.238804437	36.204870693
Time to Maturity 1 1	100%	18.943803115	26.214800088	32.80110028
Time to Maturity 1 2	110%	14.40498398	22.82878984	29.438700395
Time to Maturity 1 3	120%	11.572371915	19.98551985	26.712106306
Time to Maturity 2 1	130%	9.378627919	17.92001947	24.338218658
Time to Maturity 2 2	140%	7.67628507	15.576362600	22.287380286
Time to Maturity 2 3	150%	6.346384455	13.870480884	20.451952011

The input to all of the tables is similar. We describe the single option table first. The input entries look like this:

Single Option Price and Implied Volatility

Common parameters

v0 0.2

theta 0.1

rho -0.3

kappa 2

sigma 1

Strike 100

Forward Value 100

Maturity 1

Discount Factor 1

Is Call

Recalc

Recalc

The first five entries are simply the parameters for the Heston model as described before.

The other entries are also self explanatory: We have the strike of the option, the forward value of the underlying, the maturity of the option, the discount factor from maturity and finally we can choose to price a call or put option.

The input for the other two tables is almost the same, but we specify three maturities and the corresponding forward values at those maturities.

The Option Pricing and Implied Volatility calculation functions

Here we describe the functions used for this workspace tab.

Let us say we want to price a call option in the Heston model where $v_0=0.2$, $\theta = 0.3$, $\rho = -0.2$, $\kappa = 1$, $\sigma = 0.7$, $\text{strike} = 100$, $\text{forward_value} = 95$, $\text{maturity} = 1$ (year) and $\text{discount factor} = 0.98$. Then we simply call the pricing function:

```
number v0      = 0.2;
number theta   = 0.3;
number rho     = -0.2;
number kappa   = 1;
number sigma   = 0.7;

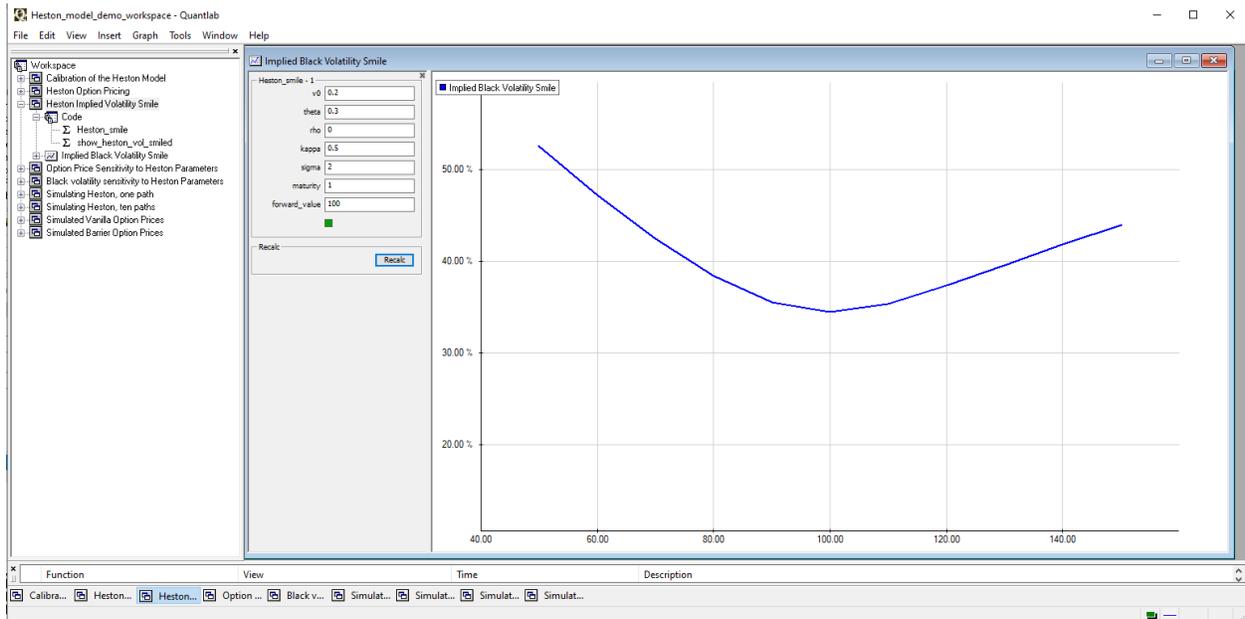
number strike  = 100;
number forward_value = 95;
number maturity = 1;
number df      = 0.98;
logical is_call = true;

number option_price = heston_option_price(v0, theta, rho, kappa, sigma, strike, forward_value, maturity, df, is_call)
```

If we want to calculate the implied Black volatility we call the function doing that (using the same values of the parameters as above):

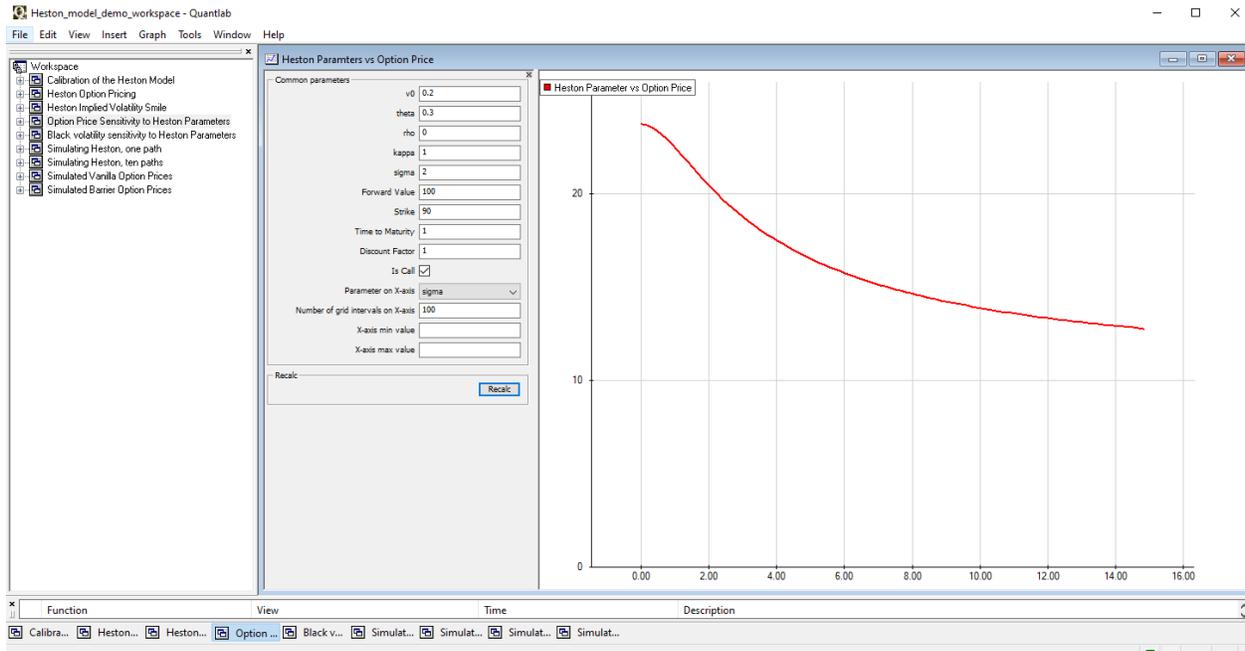
```
number implied_black_vol = heston_implied_black_volatility(v0, theta, rho, kappa, sigma, strike, forward_value, maturity);
```


Workspace Tab “Heston Implied Volatility Smile”



The purpose of this is to show the implied Black volatility smile given the five Heston parameters together with the time to maturity and the forward value of the underlying. The same code as in the last tab is used, so nothing new code-wise here.

Workspace Tab “Option Price Sensitivity to Heston Parameters”



This tab can give a lot of insight in how the different parameters of the Heston model influence the option price. The first input entries in this table is the, by now, well known Heston parameters, the forward value, time to maturity, discount factor and a flag `is_call` for switching between call and put options.

The graph shows the option price as a function of a chosen input parameter. All input parameters will be fixed, as given by the user, except one, which is the one we put on the x-axis.

So if we choose “**Parameter on X-axis**” as sigma. The graph shows how the option price depends on sigma (The value of sigma given by the user in the “sigma” box is then ignored).

We can put all Heston parameters on the X-axis, but also the strike, forward value and time to maturity.

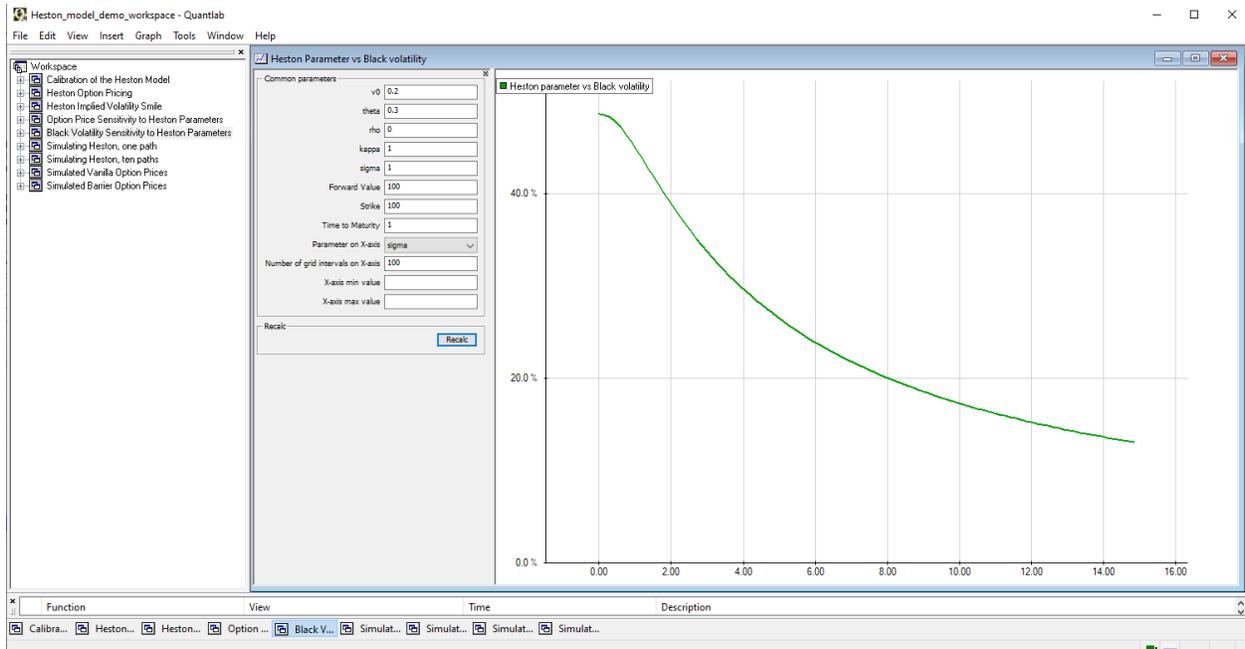
The entry “**Number of grid intervals on X-axis**” determines how many intervals and grid points we will have on the X-axis. Typically around 100 gives a smooth enough graph.

Then we have two optional entries “**X-axis min value**” and “**X-axis max value**”.

Those specify the left and right endpoints in the X-axis. If not given they will be defaulted.

Those are useful if we want to zoom in on a particular part of interest of the graph.

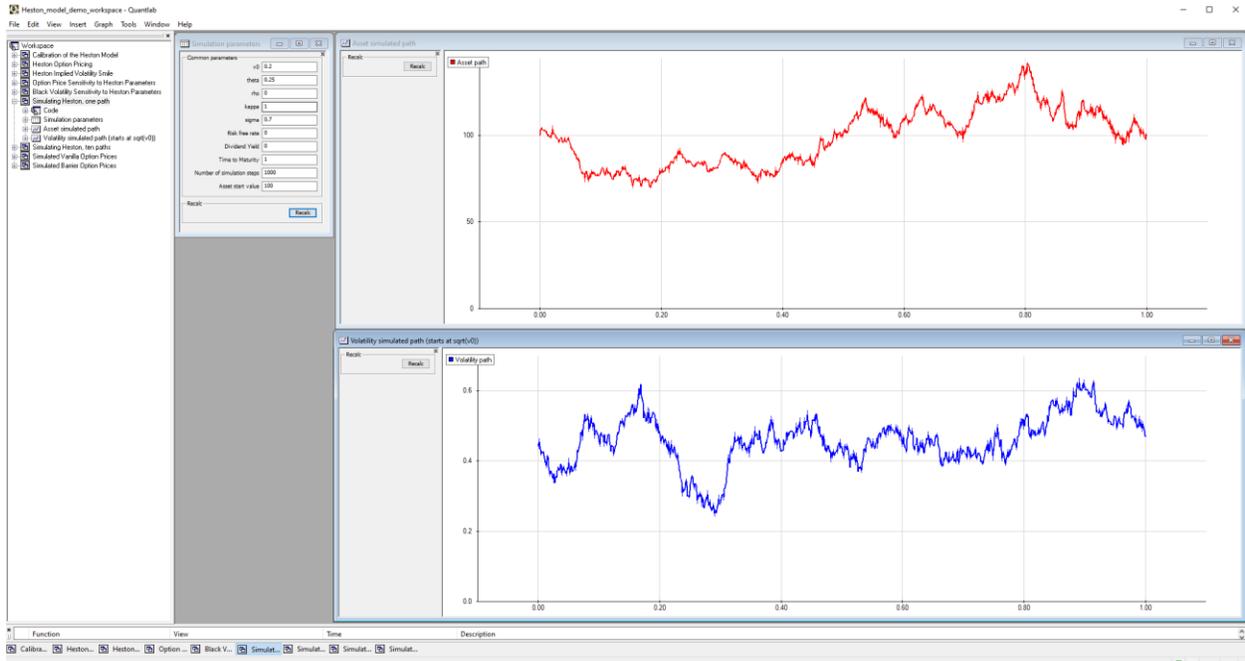
Workspace Tab “Black Volatility Sensitivity to Heston Parameters”



The input to this table is identical to the table last tab: “Option Price Sensitivity to Heston Parameters”.

The only difference is that the graph shows Black implied volatility instead of option price.

Workspace Tab “Simulating Heston, one path”



This tab simulates the Heston model using one sample path and present the result in two graphs. The first graph is the asset itself and the second the volatility process (square root of the variance process).

The input is the usual five Heston parameters, but also a constant risk free rate and dividend yield used for the drift of the asset. We also choose a time to maturity here, meaning where the simulation will stop.

We also need to provide a start value for the asset and the number of simulation time steps to use.

One may note the dependence of the asset on the volatility graphically. When the volatility graph goes down, the asset tend to be less volatile, as expected.

The Simulation Class

Let us say we want to do a 1 year long simulation of an asset with initial value = 100 using Heston parameters $v_0=0.2$, $\theta = 0.3$, $\rho = -0.2$, $\kappa = 1$, $\sigma = 0.7$ using 1000 sample paths and 100 time steps. Let us assume the risk free rate is 0.02 and the dividend yield is 0.01.

The code for creating a simulator and doing the simulation would then look like this.

```

number v0      = 0.2;
number theta   = 0.23;
number rho     = -0.2;
number kappa   = 1;
number sigma   = 0.7;

integer n_paths      = 1000;
number initial_asset_value = 100;

heston_simulation s = heston_simulation(n_paths, initial_asset_value, v0, theta, rho, kappa, sigma, QE);

number simulation_length = 1;
integer n_time_steps     = 100;
number rate              = 0.02;
number div_yield         = 0.01;
number drift             = rate - div_yield;

integer seed = millisecond(now());
rng my_rng = rng(seed);

s.simulate(simulation_length, n_time_steps, drift, my_rng);

```

Note that we needed to create a provide a random number generator to the simulation function.

The last argument QE in the constructor represents the internal simulation model to use. This last argument is an enum describing different simulation models.

```
enum heston_sim_model option(category : "Finance/Heston Model") {TV option(str : "tv"), QE option(str : "qe")};
```

TV stands for “Transformed Volatility” scheme. It is a simple and very fast simulation scheme. However, it seems somewhat biased in some situations.

QE stand for “Quadratic Exponential” Scheme. This is the recommended scheme which is the default.

This scheme performs well, especially when using so called martingale corrections.

The last argument could be skipped and safely be defaulted to the QE scheme.

The simulation class has a few member function to obtain the simulated samples:

```

vector(number) asset_values();
vector(number) log_asset_values();

vector(number) volatility_values();
vector(number) variance_values();

number current_time();

```

asset_values() Returns the simulated values for the asset.

log_asset_values() Returns the logarithm of the simulated values for the asset.

volatility_values() Returns the simulated volatilities (square root of the variances).

variance_values() Returns the simulated variances (squared volatilities).

current_time() Returns where in time the simulation is at the moment.

The point in having the four getter functions instead of two for the asset and its variance process is to avoid doing unnecessary exponentials, logarithms, square roots and squares if it is not necessary. If we call the simulate function many times this can be costly.

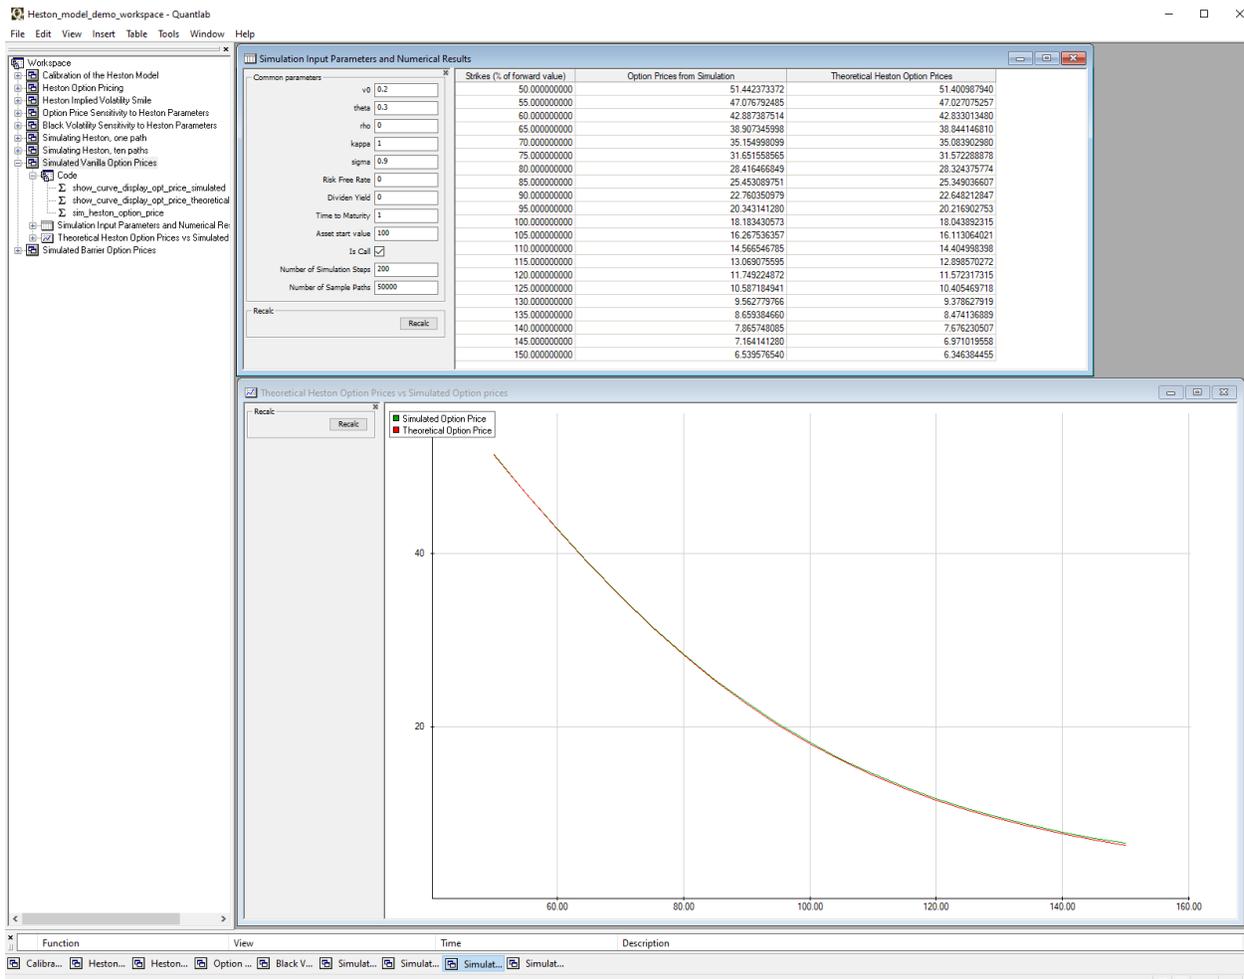
Workspace Tab “Simulating Heston, ten paths”



This tab is identical to the previous tab, except that we simulate ten paths for both the asset and volatility.

The point is to get a bit of understanding how much different paths differ from each other.

Workspace Tab “Simulated Vanilla Option Prices



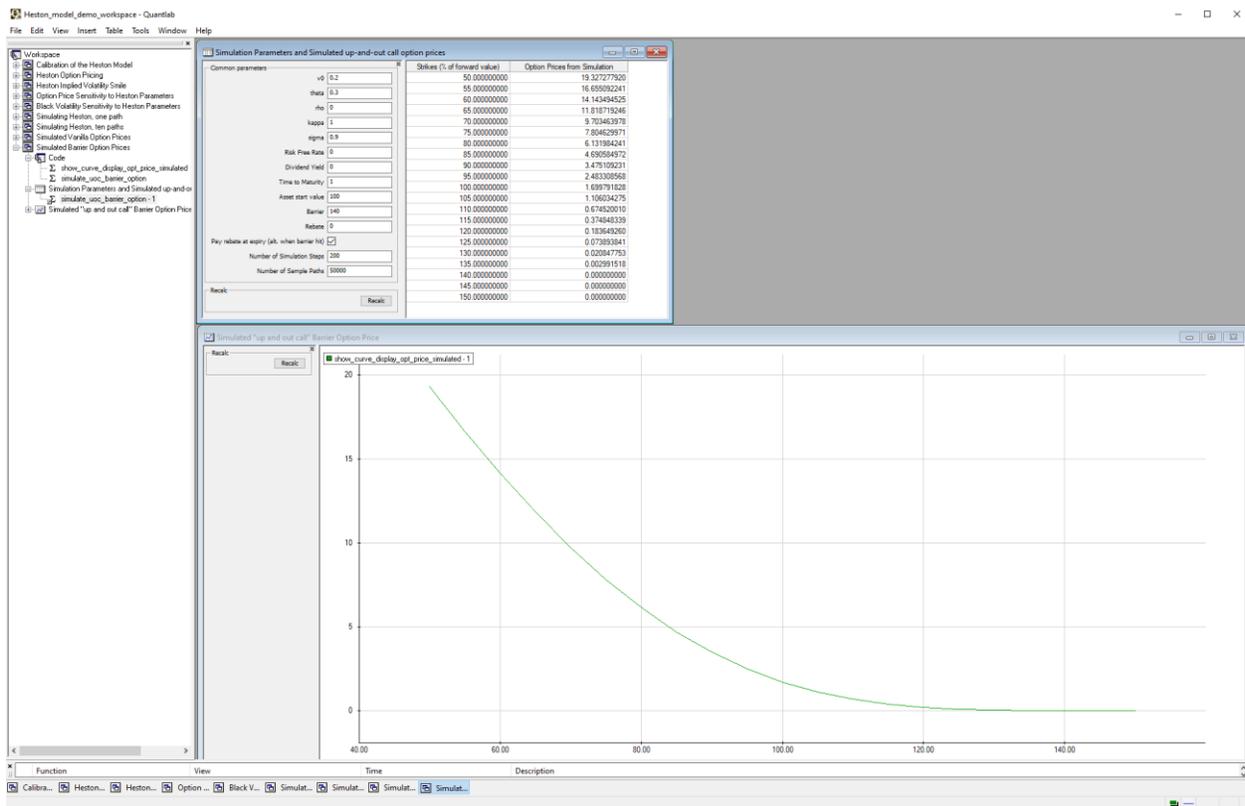
The point of this tab is to price vanilla call/put options in the Heston model and to compare those to the theoretical option prices.

The theoretical option prices are the correct prices in the model. So we can study the quality of the simulation here.

The input is the same as in the “Simulated Heston, one path” and “Simulated Heston, ten paths” tabs. The output is the simulated and theoretical option prices in the Heston model, both in numerical and in the form of graphs.

The simulated option price is computed the usual Monte Carlo way by averaging the payoff over all samples.

Workspace Tab “Simulated Barrier Option Prices”



The point of this tab is to show how the Heston simulation class can be used to price something slightly more complicated, a barrier up-and-out call option with rebate.

This is an ordinary vanilla European call option, but with the extra feature that option expires and becomes worthless if the simulated asset ever goes above the barrier level. If that happens a rebate is paid (which can be 0).

The rebate can be paid either when the barrier is hit or at expiry.

The input here is the usual one, except for the obvious new entries where we can set the barrier, the rebate and if it is paid on hit or expiry.

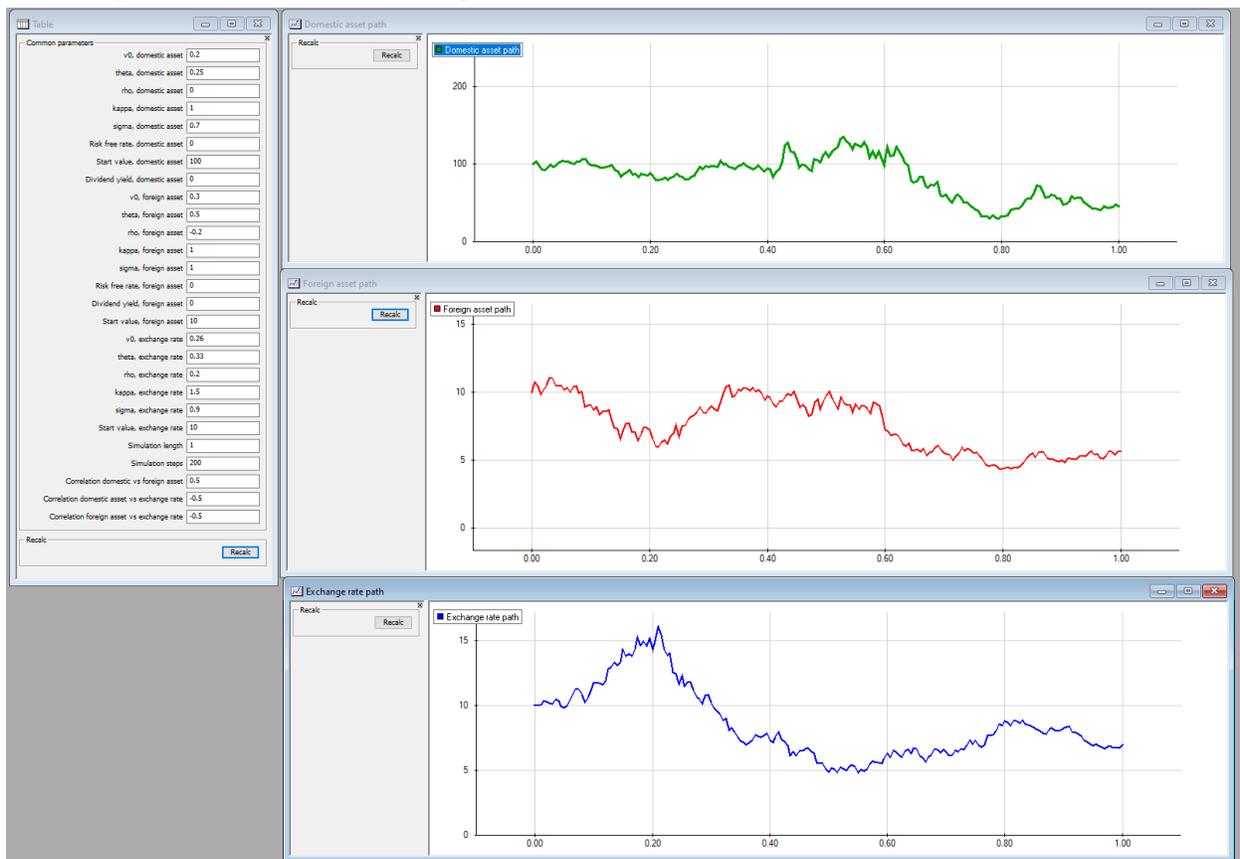
This is an example where we need to split the simulation interval in many sub steps and check if the barrier has been hit in each of those steps.

Here there we have the obvious problem that the barrier might have been hit between steps.

However, one can calculate that probability and adapt the calculations to that.

We don't go into details of the math here.

Workspace Tab “Simulating Multi-Asset Heston”



The point of this tab is to illustrate the Multi-Asset Heston model in a simple example.

We do a simulation of three factors. First we have an asset in the domestic base market, let us say the “SEK” market.

Then we add another “foreign” market, say the “USD” market. When doing so we simulate the exchange rate for USD/SEK in the Heston model.

Then we add an asset on foreign market and simulate this together with the rest.

And finally we also can add correlations between the simulated assets and exchange rate.

The input is of the same type as for the single asset Heston, but repeated, except that we now can give the correlations between the factors too.

So let’s continue to describe the usage of the code behind this.

[The multi-asset Heston simulation class.](#)

The Multi-asset Heston simulation class has a constructor that looks like this

```
multi_asset_heston_simulation(integer n_paths, string domestic_currency, disc_func domestic_yield_curve);
```

Here we create the simulator object and also sets up the domestic market for the simulation.

n_paths is the number of sample paths to use in the simulation.

domestic_currency is the currency of the domestic market, like “SEK”. It is just a label and can be chosen freely, so it does not have to match any existing currency.

domestic_yield_curve is the risk free yield curve on the domestic market represented as a *disc_func* object.

Example:

To create a Multi-asset Heston simulator with “SEK” as domestic market and currency, using 1000 sample paths and using a risk free domestic yield curve with nodes at 1,2 and 3 year with yields 0.01 (1%), 0.011 (1.1%) and 0.006 (0.6%), we would write

```
integer n_paths = 1000;
string domestic_currency = "SEK";
vector(number) maturities = [1,2,3];
vector(number) domestic_yields = [0.01, 0.011, 0.006];

interpolator ip = ip_linear(0,0); /
disc_func domestic_yield_curve = disc_func_interp(maturities, exp(- domestic_yields .* maturities), ip);

multi_asset_heston_simulation sim = multi_asset_heston_simulation(n_paths, domestic_currency, domestic_yield_curve);
```

The interpolator object tells us that we should use linear interpolation and constant extrapolation (derivatives 0) before the first point and after the last point.

Let us now add an asset to the given domestic market. To do this we use the member function with signature

```
void add_asset(string asset_name,
              string asset_currency,
              number initial_value,
              number heston_initial_variance,
              number heston_long_term_variance,
              number heston_correlation,
              number heston_mean_reversion,
              number heston_vol_vol,
              disc_func option(nullable) dividend_yield = null<disc_func>);
```

Example:

We call the asset “TELIA” and its currency should be “SEK”. We choose/construct a dividend yield curve the same way we construct a yield curve. This is not mandatory, though.

We need to provide the assets starting value, which we choose to be 100 SEK and also the five Heston model parameters, which we choose a bit arbitrary in this example.

To do all this, we write:

```

string domestic_asset_name = "TELIA";
string domestic_asset_currency = "SEK";

vector(number) maturities = [1,2,3];
vector(number) domestic_asset_div_yield = [0.005, 0.002, 0.001];
disc_func domestic_asset_div_yield_curve = disc_func_interp(maturities, exp(- domestic_asset_div_yield .* maturities), ip);

number v0 = 0.1;
number theta = 0.3;
number rho = -0.3;
number kappa = 2;
number sigma = 0.4;

number asset_start_value = 100;

sim.add_asset(domestic_asset_name, domestic_asset_currency, asset_start_value, v0, theta, rho, kappa, sigma, domestic_asset_div_yield_curve);

```

So far so good. We can, so far, only add assets to the domestic market, since we have not introduced any foreign markets yet.

Let us now assume we want to add the “USD” market/currency. This can be done in several ways. We have three different ways of adding a foreign market:

1. We can choose to simulate the exchange rate to USD in SEK using a Heston simulation. Then we need to provide the Heston parameters.
2. We can also choose not to simulate the exchange rate, since we might not have any use for it. Then it is assumed that the exchange rate is modelled using a standard Black-Scholes model. Even if the process is not simulated, its volatility and correlations to the assets on the same foreign market are needed to calculate the so called quanto-adjustments for the drift terms of the foreign assets.
3. Alternative 3. Is a lazy version of alternative 2, but here we don't even require the volatility of the exchange rate. This is then assumed to be zero, which will make the quanto adjustments zero as well. Since the quanto-adjustments typically are small compared to the other terms, ignoring them will not have a huge effect on the simulation.

The three member functions to use in all those cases have signatures

```

void add_foreign_market(string foreign_currency,
                       disc_func foreign_yield_curve);

void add_foreign_market(string foreign_currency,
                       disc_func foreign_yield_curve,
                       number fx_black_vol);

void add_foreign_market(string foreign_currency,
                       disc_func foreign_yield_curve,
                       number fx_initial_value,
                       number fx_heston_initial_variance,
                       number fx_heston_long_term_variance,
                       number fx_heston_correlation,
                       number fx_heston_mean_reversion,
                       number fx_heston_vol_vol);

```

The input is pretty straightforward. We need a label/currency for the market, a given risk free yield curve for the market. If we choose to simulate the exchange rate, we need an initial value and the Heston parameters.

So adding a “USD” foreign market with the “USD” exchange rate simulated using Heston (using some arbitrary data) we would write

```
number v0 = 0.1;
number theta = 0.3;
number rho = -0.3;
number kappa = 2;
number sigma = 0.4;

number fx_start_value = 10;

string foreign_currency = "USD";
vector(number) maturities = [1,2,3];
vector(number) foreign_yields = [0.01, 0.015, 0.02];
disc_func foreign_yield_curve = disc_func_interp(maturities, exp(- foreign_yields .* maturities), ip);

sim.add_foreign_market(foreign_currency, foreign_yield_curve, fx_start_value, v0, theta, rho, kappa, sigma);
```

Let us finally assume we want to add an asset to the foreign “USD” market. Note that the market must have been added before the asset can be added. This is done using the same function to add assets to the domestic market.

Example:

```
number v0 = 0.1;
number theta = 0.3;
number rho = -0.3;
number kappa = 2;
number sigma = 0.4;

number foreign_asset_start_value = 10;

string foreign_asset_name = "TESLA";
string foreign_currency = "USD";
vector(number) maturities = [1,2,3];
vector(number) foreign_asset_div_yield = [0.006, 0.003, 0.001];
disc_func foreign_asset_div_yield_curve = disc_func_interp(maturities, exp(- foreign_asset_div_yield .* maturities), ip);

sim.add_asset(foreign_asset_name, foreign_currency, foreign_asset_start_value, v0, theta, rho, kappa, sigma, foreign_asset_div_yield_curve);
```

So we now have three simulated factors in our simulator: The domestic asset “TELIA”, the foreign asset “TESLA” and the exchange rate “USD” (in SEK).

We can choose to add correlations (for the log returns) between any of those. This can be done using the member function with signature

```
void add_correlation(string asset_or_fx_name_1, string asset_or_fx_name_2, number correlation);
```

Example:

```
sim.add_correlation("TELIA", "TESLA", 0.5);
sim.add_correlation("TELIA", "USD", -0.4);
sim.add_correlation("TESLA", "USD", -0.3);
```

When this is done the simulator is ready to start simulating (unless we want to add more factors).

This we do using the simulate member function with signature

```
void simulate(number simulation_length,  
              integer n_simulation_steps,  
              rng     my_rng,  
              logical use_martingale_correction = true);
```

We need to give the length of the simulation (in years), the number of time steps to use, a random number generator and lastly an optional technical parameter which tells the simulator to do a so call martingale correction or not (this is a method to remove model bias due to too large time steps).

When we want to retrieve values from the simulator, we can use the getter functions

```
vector(number) asset_values(string name);  
vector(number) log_asset_values(string name);  
  
vector(number) fx_values(string name);  
vector(number) log_fx_values(string name);  
  
vector(number) volatility_values(string asset_or_fx_name);  
vector(number) variance_values(string asset_or_fx_name);
```

We have separated getters for the simulated exchange rates and the assets for clarity.