

MetaTrader 4 Backtesting & Optimization

To get the most out of your expert advisor, you'll need to optimize and backtest your strategy using MetaTrader's Strategy Tester. While forward testing on a demo account is essential, backtesting allows you to simulate trading over a long period of time in just minutes. And with the optimization feature, you can find out which settings performed best over a selected historical chart period.

There is considerable debate over the accuracy of MetaTrader's strategy tester. At best, backtesting offers only a close approximation of how trades would be executed in real-time. But it's the only tool available to rapidly test any strategy over a wide range of trading situations, and one that you should learn how to use well.

Open the Strategy Tester in MetaTrader by clicking the appropriate button on the toolbar or by selecting Strategy Tester from the View menu.

History Data

Before backtesting or optimizing, it's important to make sure that your history data is complete and accurate, especially if you're using 'Every tick' as your testing model. If you see 'mismatched chart' errors in your Journal log or if your modeling quality is less than 90%, your history data is insufficient to generate accurate ticks.

Open the History Center from the Tools menu or by pressing F2 on your keyboard. Double click the chart pair in the left column that you plan to backtest for. A list of time periods will appear below.

Start by double clicking on 1 Minute (M1) to load the history data for that period. The backtester uses M1 data to generate ticks, so it is important that your M1 data is complete.

Click the Download button to download the full history data for that pair and period. You may receive a Download Warning message – hit OK to continue. The history center will commence downloading. Repeat this process for every other time period and pair that you plan to test with.

You can also import history data into the History center, using accurate historical price data from Alpari. The [Alpari databank](#) has M1 data for over a dozen popular currency pairs. Download and unzip the history data to your hard drive. Click the Import button in the History Center window, and then click Browse to select the history data file (.hst).

You can also convert this accurate M1 data to other timeframes using the Period Converter script.

Period Converter

I decided to write up this little tutorial, because backtesting different systems comes up very often in threads on this forum. There seems to be a lot of confusion about reliability issues and how to go about achieving the most accurate possible results. I am not a programming or trading guru, but I believe I can provide a helpful little FAQ on backtesting using MT4.

Good backtesting is important when considering a system-trading approach, because you want to have some idea of the feasibility of your idea before you go live with it [at least I do]. If you're backtesting with a 50% model quality, eh... you can't really be sure what's going on. If you have a 90% modeling quality, you can have more confidence on how your system actually would have performed.

Contents:

- Section 1: Is MT4 Backtesting Reliable?
- Section 2: Downloading/Importing/Converting 1M Data
- Section 3: Configuring the Backtester
- Section 4: Other Issues

Section 1: Is MT4 Backtesting Reliable?

This question often gets pretty heated and people even get to the point of flaming each other about it. Backtesting in MT4 can be reliable, but its reliability is contingent upon the data you are backtesting on. Demo account data that is streamed in through a demo account broker has gaps, holes, and is basically not suitable for testing.

When backtesting, you want to use the EVERY TICK MODEL and have accurate 1M data to get the most accurate test possible. The 1M data is important, because the EVERY TICK MODEL uses whatever the smallest available timeframe available is and “fakes” the movement of price within the smallest available bars. Having 1M data allows for the fractal interpolation within bars to occur only within the very narrow range of 1M bars.

The easiest [and only] solution to this is to use good 1M data. The most complete data you can get [at least for free] is from Alpari's Databank. They have data in MT native format, on the 1M timeframe back through mid 2004. However, setting up the data for use requires some doing.

History Center: GBPUSD, M1

Database: 36860 records

Time	Open	High	Low	Close	Volume
2008.03.28 23:00	1.9937	1.9937	1.9937	1.9937	1
2008.03.28 22:59	1.9932	1.9936	1.9932	1.9936	7
2008.03.28 22:58	1.9935	1.9935	1.9932	1.9933	14
2008.03.28 22:57	1.9934	1.9936	1.9933	1.9934	14
2008.03.28 22:56	1.9935	1.9935	1.9931	1.9933	12
2008.03.28 22:55	1.9937	1.9937	1.9935	1.9936	6
2008.03.28 22:54	1.9931	1.9936	1.9928	1.9936	11
2008.03.28 22:53	1.9931	1.9931	1.9930	1.9930	6
2008.03.28 22:52	1.9933	1.9933	1.9932	1.9932	2
2008.03.28 22:51	1.9934	1.9934	1.9932	1.9932	7
2008.03.28 22:50	1.9933	1.9936	1.9932	1.9935	9
2008.03.28 22:49	1.9926	1.9933	1.9926	1.9932	13
2008.03.28 22:48	1.9926	1.9929	1.9925	1.9925	9
2008.03.28 22:46	1.9923	1.9925	1.9923	1.9925	3
2008.03.28 22:45	1.9924	1.9925	1.9924	1.9924	3
2008.03.28 22:44	1.9924	1.9925	1.9923	1.9925	6
2008.03.28 22:43	1.9926	1.9928	1.9924	1.9925	14

Section 2: Downloading/Importing/Converting 1M Data

(1) You need to modify MT4 to allow for more bars. Go into the Tools Menu, then go to Options [or just hit C+O]. Go into the charts tab and put in 9999999999999 for bars in history. MT4 will default to whatever it's maximum is.

[Note: The reason MT4 has a limited bar count to begin with is because more bars (particularly when used in backtesting models) means MT4 is going to eat up more HD space.]

(2) Download the 1M data from Alpari's Databank in whatever currency[ies] you're going to test on.

(3) Import the data into MT4 using the History Center. Go to Tools => History Center [or push F2]. Make sure you import it in proper currency and in the M1 timeframe. You don't want EURUSD data being imported into USDCAD for instance.

(4) Convert the data using the period converter script included in MT4 [you only have 1M bars right now]. You have to open offline charts to do this.

-Go to the File Menu, then Open Offline, select the 1M data of the currency you need to convert. A chart will pop up with that data.

-Then drag & drop the period_converter script onto the offline chart. The ExtPeriodMultiplier int that you can modify is the multiplier you are applying to the chart. So making it 5, will convert 1M data into 5M data.

-For simplicity's sake, you need to run the period converter with the following integers to get all the backtesting timeframes: 5,15,30,60,240, and 1440.

[NOTE: you can also convert 1M data to timeframes not native to MT4 if you want to do some indicator analysis or something on another timeframe.]

Congratulations, you have now imported and converted data into MT4. Now, for the sake of illustrating one of my earlier points, open up a currency you have imported data on. Look at the difference in the bars from the downloaded data as opposed to data streamed in from a Demo broker [So, if you downloaded 1M data from July 04 to August 05, look at the chart at August 05's end and September 05's beginning]. You will notice that the bars (on every time frame if you have converted them properly) from your downloaded time period will be more complete.

Section 3: Configuring the Backtester

Now that you've successfully imported complete data, there are a few more things you need to do to run a reliable backtest.

(1) Check the recalculate option the next time you run a backtest, because you need the backtester to utilize your shiny new happy data (which it won't do unless you tell it). Anytime you import new data, you need to recalculate (I recalculate every few tests just to feel safe, maybe it's a reflection of internal confidence problems, but that's for another FAQ).

(2) Check the use date option and set the date range only over a time period where you have good reliable data. This way you're only backtesting the good stuff. It will be reflected in the modeling quality percentage.

(3) Make sure the model is set to EVERY TICK. If you're not, all this hard work we just did was for nothing. I addressed why we do this earlier in the FAQ.

Section 4: Other Issues

MT4 is a work in progress, sometimes there are strange bugs that crop up in backtesting. However, usually when you think you have a bug on your hands, there is something wrong with your code. I can't emphasize enough how important debugging is. If you have problems, check

your code first because it's probably the problem. If you really think you have a legit bug on your hands, post it to the MT4 forums.

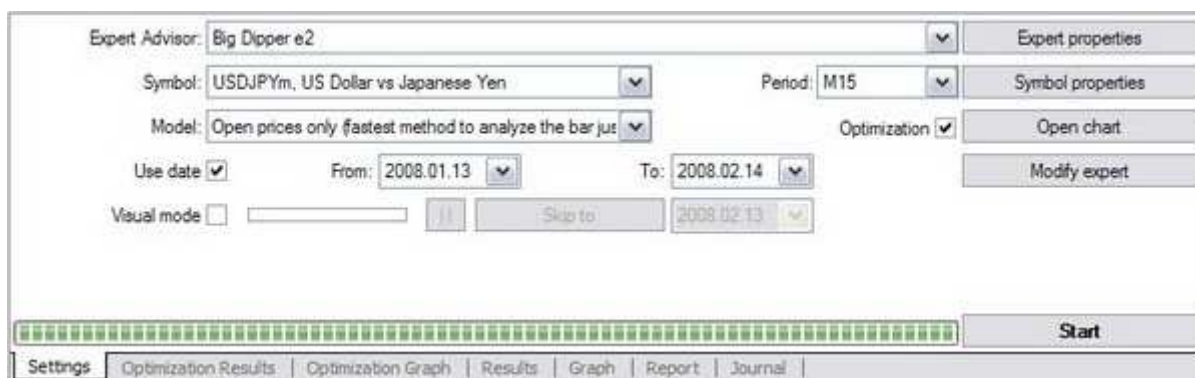
Because you are not actually backtesting on every tick that happened [you are dealing with an interpolation on 1M data], it is still not a perfect reproduction of what actually happened in the markets. Because of this, 1M and 5M scalping EAs that get in and out of trades really quickly will run into some problems just because of this limitation. The longer timeframe you are trading on, the less likely your testing is to be hampered by this.

Well, that's all I can think of now. I read this over, I think I made everything clear and have the steps outlined correctly. If you notice a mistake, let me know, and I'll correct it in my next version of the MT4 Backtesting FAQ.

Optimization

The optimization feature of MetaTrader 4 allows you to test thousands of combinations of expert advisor settings to find the most profitable settings for the selected chart, period and date range. Indicator-based strategies will need to be optimized for maximum profitability. However, almost all EAs will benefit from optimization – even those that trade on tick data, provided you have complete M1 history data (see above).

While the optimizer will return the most profitable settings for the selected date range, this is no guarantee that these settings will be profitable in the future. Market conditions change often, so it is important to regularly re-optimize your expert advisor for best results. Read on for more information on how to properly optimize your expert advisor for profitability.

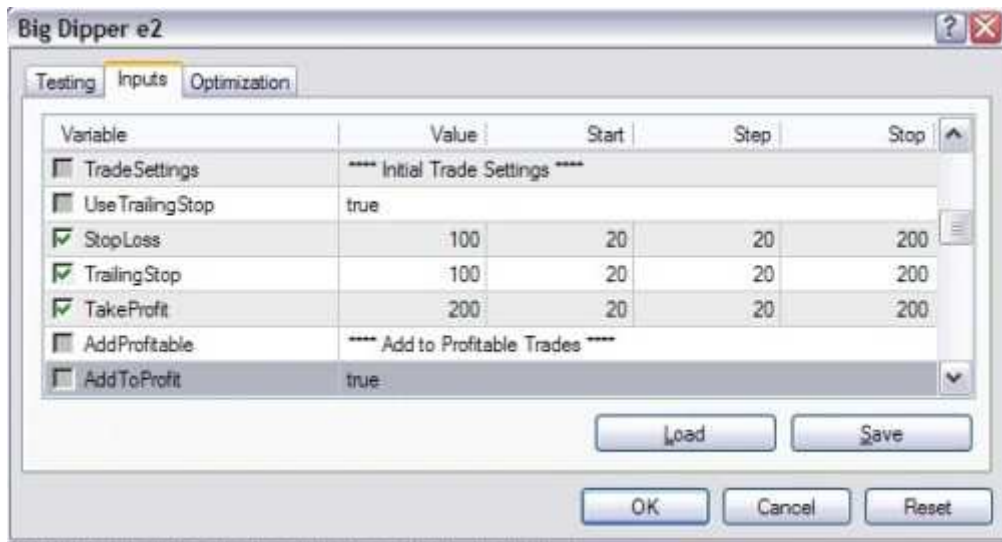


To optimize your expert advisor, first select it from the Expert Advisor drop-down box. Select the currency pair from the Symbol box and chart period from the Period box. For Model, you'll generally want to select "Open Prices Only," unless you are optimizing an EA that runs on tick data. In that case, select "Every Tick." Check the Use Date option and select a range of dates to optimize for.

Lastly, make sure that Optimization is checked.

Click the Expert Properties button to open your expert advisor settings.

Under the Inputs tab is where you'll enter the range of values to optimize for. The Start column will be the lowest value for a given setting, while the Stop column will be the highest. The Step column is the amount that the optimizer will "step through" from the Start to the Stop setting.



In the image above we are optimizing SL, TS and TP settings for an expert advisor. The Start value is 20, the Step is 20, and the Stop is 200. The optimizer will test every combination of values from 20, 40, 60 and so on up to 200. Use a start, step and stop value that is appropriate for the setting you are optimizing. Even values (5, 10, etc.) are good.

The checkbox to the far left must be selected for that setting to be optimized. Any settings that aren't checked will use the number in the Value column when optimizing.

Under the Testing tab, you can adjust the Initial Deposit to something a bit more realistic. Leave the other settings at their defaults.

When you're ready to begin optimizing, hit the Start button at the bottom right of the Strategy Tester window. Depending on the period, the date range, the testing model and the number of settings to be optimized it can take anywhere from a few minutes to several hours. If it's taking too long, consider shortening the date range, optimizing fewer settings, or using a larger step value.

Pass	Profit	Total trades	Profit factor	Expected Payoff	Drawdown \$	Drawdown %
60	214.21	3		71.40	174.60	60.01
73	190.85	2		95.43	174.60	83.14
59	178.20	5	3.80	35.64	167.40	78.94
74	124.49	3		41.50	130.55	38.83
47	121.14	2		60.57	86.26	32.35
117	118.00	2		59.00	26.73	12.90
49	94.32	4	2.08	23.58	103.36	48.32
76	94.06	4	2.63	23.51	90.67	43.41
58	93.69	4	9.48	23.42	68.40	33.58

Once the optimization is finished, open the Optimization Results tab and double-click the Profit column to sort the results. Double-click any of the results to load it into the tester. Hit the Start button again to backtest with the selected settings.

Backtesting

By now, it should be obvious how the backtester works. Select your Expert Advisor, Symbol, Period and Model, check the Use Date box and select a date range.

Select Visual Mode only if you want a visual walkthrough of the backtesting. Leave Optimization unchecked.

Hit the Expert Properties button and enter your settings in the Value column under the Inputs tab. You can also load or save settings using the buttons in the bottom right. The Start, Step and Stop columns are ignored, as are the checkboxes.

Close the Expert Properties dialog and press Start to begin testing. It will take anywhere from a few seconds to several minutes depending on your settings. Once testing has finished, open the Report tab on the bottom to see your results.

Bars in test	3181	Ticks modelled	5359	Modeling quality	n/a
Mismatched charts errors	0				
Initial deposit	200.00				
Total net profit	214.21	Gross profit	214.21	Gross loss	0.00
Profit factor		Expected payoff	71.40		
Absolute drawdown	83.65	Maximal drawdown	174.60 (60.01%)	Relative drawdown	60.01% (174.60)
Total trades	3	Short positions (won %)	0 (0.00%)	Long positions (won %)	3 (100.00%)
		Profit trades (% of total)	3 (100.00%)	Loss trades (% of total)	0 (0.00%)
	Largest	profit trade	123.26	loss trade	0.00
	Average	profit trade	71.40	loss trade	0.00
	Maximum	consecutive wins (profit in money)	3 (214.21)	consecutive losses (loss in money)	0 (0.00)
	Maximal	consecutive profit (count of wins)	214.21 (3)	consecutive loss (count of losses)	0.00 (0)
	Average	consecutive wins	3	consecutive losses	0

Settings | Optimization Results | Optimization Graph | Results | Graph | **Report** | Journal

A few statistics to take note of:

- Total net profit – The Gross profit minus the Gross loss.
- Profit factor – The ratio of gross profit to gross loss. Higher is better, anything above 2 is acceptable.
- Absolute drawdown – The drawdown of your initial deposit. High drawdowns increase the likelihood that your account will be blown out.
- Profit trades – Your overall win percentage.
- Modeling quality – Only important if your testing model is Every Tick. If so, this should be at 90%. If not, follow the instructions above to update your history with accurate M1 data.

The Results tab at the bottom of the strategy tester will give you the details on opened and closed orders, including trailing stop, take profit and stop loss. Click the Open Chart button to get a visual representation of your results.

When testing your new EA, examine these closely to ensure that your strategy is working as intended.

Optimizing for Best Results

When optimizing your EA, you should optimize as few parameters as possible. Optimizing too many parameters leads to curve-fitting, which produces great results in optimization, but performs poorly in real time trading.

For indicator-based strategies, the period settings for your indicators will have the greatest impact on your results. Risk-based settings such as lot size, stop loss and take profit should be set according to your predetermined risk profile. Other settings can be left at their defaults.

The surest way to evaluate real-time performance based on optimized results is to do a walk-forward analysis. This involves optimizing your EA over a reasonably large time period (the optimization window), and then testing your optimized parameters over a shorter time period following the optimization window (the walk-forward window).

Be sure to choose an optimization window that is large enough to include a sufficient number of trades, as well as a variety of market conditions. Optimizing a daily chart on 1 months data is not enough to generate sufficient testing data. On the other hand, optimizing a 1 minute chart over several years is probably overkill.

The walk-forward window should be roughly 25% of the size of the optimization window. A parameter set optimized over a 1 year window, for example, will only be good for a few months before performance starts to degrade. Your results may vary, however.

By repeating the process of optimizing and forward testing over many consecutive testing windows and across several chart pairs, we can get a good idea of how our EA really performs. Short of real-time trading, which would take months, if not years, this is the best way to evaluate an EA's performance.

Let's say we have a [trend](#)-following strategy that uses 2 MACD indicators. Trades are only placed when both indicators agree, so this strategy is not always in the market. We are optimizing the FastEMA and SlowEMA only. Our lot size is set to 2% of our equity, and our stop loss is set to 5% of our equity. The starting balance is 1000.

We are trading this strategy on the [GBPUSD](#) on an H1 chart period, and we have determined that an optimization window of 12 months will give us a large enough sample to work with. This gives us an average of approximately 100 trades over the optimization window.

We optimize the strategy from 12.01.2006 to 11.31.2007, using "Open prices only" as our testing model. (Using "Every tick" as the testing model can take a long time, so optimize as few parameters as possible if you do).

Sort the results in the Optimization Results window by profit, and look for a good parameter set to test with. Don't automatically pick the result with the largest profit. If a result is significantly more profitable than the rest, then it is a statistical outlier and not representative of the performance of your EA. Look for settings that are profitable across multiple result sets, and choose a representative result to test.

This is also a good time to examine the percentage of the results that have a profit or a loss. Sort the results by profit and scroll down through the list. A robust strategy should have a high percentage of profitable result sets. Also take note of the profit factor and the drawdown.

If the profit factor is less than 1.5, and too close to 1.0, you're losing too much of your gross profit to losses. The drawdown percentage should ideally be around 10-15%, and no more than 20%. If your drawdown consistently exceeds this, then you are either taking on too much risk (large lot size or stop loss), or the strategy is potentially unprofitable.

For our strategy above, we've chosen the top performing result set with a profit of 117.75. Now we're going to test it over the walk-forward window. Since we used an optimization window of 12 months, we'll use a walk-forward window of 4 months to test our optimized results.

We test our optimized results (make sure Optimization is unchecked!) from 12.01.2007 to 3.31.2008.

Our result for this period is a profit of 104.27 in just 4 months. This is a good result.

But we're not done yet. Now we must repeat our optimization and testing process many times and keep track of our results to get a better idea of the forward performance of our EA. Since we used a testing period of 4 months, we'll step forward our optimization window to 4.01.2007 – 3.31.2008.

Then we'll test our optimized result in a walk-forward window from 4.01.2008 – 7.31.2008.

Our second round gave us an optimized profit of 415.02, with a walk-forward loss of -296.30. In this case, the optimized parameters did not perform as well in the walk-forward. Here are the results of two years of walk-forward testing for this strategy:

Optimized Walk-forward

(12 months)	(4 months)
167.74	302.73
321.88	-115.73
277.31	89.96
396.92	-130.13
117.75	104.20
415.02	-296.30

As you can see, the results are mixed. The strategy made money in 3 out of 6 walk-forward periods. But our total profit for our walk-forward testing is -45.27. Clearly this strategy needs some work. Several more years of testing may provide a more statistically significant sample, but for our purposes this is insufficient.

If your results are more favorable than these, you can calculate the walk-forward efficiency. This is an annualized ratio between the optimized results and the walk-forward results.

First, add up all of the optimized results. Then divide that by the total number of years in your optimization windows. In the example above, we have 6 optimization windows of 12 months each, so we would divide by 6. This is your optimized annual return.

Then, add up the walk-forward results and divide that by the total number of years in your walk-forward windows. This is your walk-forward annual return. Divide this by the optimized annual return and you should have a ratio or percentage. This is your walk-forward efficiency.

A walk-forward efficiency of 50% (0.5) or greater is desirable. If it's less than this, your system is likely unprofitable.

If you plan on trading an EA on several pairs, you will want to repeat this process for each pair you plan on trading.