

Aidyia Limited - CONFIDENTIAL

Aidyia Methodology Overview

Cassio Pennachin & Ben Goertzel - October 2014

This document gives a very brief , nontechnical summary of the process that went into generating the US Equities and Global backtest results that Aidyia produced in October 2014.

1. Overview of Aidyia's Predictive Methodology

The novel predictive methodology we have developed involves a series of stages:

- **Input features**, including
 - Price/volume based features, including standard financial indicators and more advanced mathematical indicators, produced based on daily financial data (and in some cases more granular data)
 - News based features, based on English and Chinese language newsfeeds
 - Fundamental and macroeconomic features
- **Predictive algorithms**, including 5 different algorithms used in the backtests, and 2 in development.
 - Each of these algorithms, to be applied to a certain universe of financial instruments, is first trained on historical data regarding these instruments. Based on this training, it learns a set of "predictive models" for each of the instruments in the universe on that day.
 - A predictive model for a certain instrument, applied on a certain day, views some subset of the input features generated relative to the instrument in the past up till that day, and then makes a prediction as to the direction and magnitude of price movement of the instrument N days in the future (where N is specified at the start of the training process)
 - Some of our algorithms are entirely proprietary to Aidyia; others are heavily customized versions of open-source machine learning tools.
 - We use the term "model class" to refer to the set of models coming out of a particular algorithm and predicting movements in a particular direction with a particular lookahead N
- **Signal weighting**: Methods for assigning numerical weights to predictions, based on diverse data regarding each prediction.
- **Model class weighting**: Methods for assigning numerical weights to the different model classes on a given day, based on the historical performance characteristics (and other aspects) of each model class. These are among the inputs used by our methods to assign weights to individual predictions (as one important property of a prediction is which model class it came from)

The predictive algorithms are the core of our proprietary trading methodology; they are our "secret sauce." However, our predictive algorithms alone, applied in a naive way without embedding in an appropriate domain-specific framework, would not yield adequate trading performance. Alongside our research on predictive algorithms, we have put very significant R&D effort into what comes before (the input features) and after (signal weighting, model class weighting and signal combination).

What comes from this series of stages is a set of predictive signals, for each instrument in the universe in question, on each day. Then, in Aidyia's overall trading framework, these predictions are fed into a portfolio management system which handles capital allocation, risk management, and trading system integration.

2. High Level Development and Testing Narrative

Now we will review the overall high level development and testing process of which this backtesting has been the most recent part. The overall process is important, because careful handling of data is extremely critical in this kind of work to avoid falling prey to the various overfitting-related errors that may otherwise occur.

Aidyia Limited was founded in late 2011, and since that time most of the firm's R&D efforts went into creating and refining algorithms for predicting the prices of financial instruments. Until Summer 2014, this work was focused more specifically on creating algorithms aimed at:

- Forecasting the prices of relatively high-liquidity stocks on the Hong Kong stock exchange (liquidity criteria were defined to restrict the stock universe used for experimentation, resulting in around 200 adequately liquid stocks at any point in time.)
- Forecasting these stocks' prices 1, 5, 10 or 20 days in the future.
- Achieving adequate predictive performance as validated by experimentation on historical data from 2007-2011.

Data from 2012-2013 was held out until October 2014, for use as true out of sample validation data. This data was not used at all in the experimentation or testing process, until the "true out of sample validation" test was done in late October 2014. All the testing and development was done, until this point, as if the world ended on Dec. 31, 2011.

After a period of experimentation, we decided to restrict focus to forecasting 5 or 20 days in the future. 1 day was eliminated because the results were generally not good enough (which may be largely an artifact of the large transaction costs on the Hong Kong market, on which we were doing all our testing at that point). 10 days was eliminated purely for simplicity, because running our model learning algorithms is computationally costly, so by restricting attention to 5 and 20 days we could explore more variations of our algorithms using the limited compute time at our disposal.

The Hong Kong stock market has relatively high transactions costs due both to fees imposed on each trade by law (stamp duty tax), and relatively high friction. This posed difficulty in creating a system with adequate returns on the 2007-2011 Hong Kong equity data we were experimenting with. Therefore, in Summer 2014, we decided to shift focus and try our software at the tasks of predicting US equity and global macro asset prices instead. Our US equity universe was defined as the members of the S&P 500, and the global portfolio is a collection of about 50 instruments, covering commodities, interest rates, equity indexes, and currency pairs.

Some of our software's input features were not appropriate for these other markets, e.g. news-based features and most fundamental features. We simply disabled these features for our tests on these other tradable universes. (Obviously this is not optimal, and we suspect better results could be obtained by modifying these features appropriately for these other universes, rather than simply disabling them. However, that would have been time-consuming, so we chose a more expedient route for our first pass of experimentation.)

On the US equity and global macro universes, we ran our predictive modeling algorithms on data from 2002-2011. Results were good, with basic predictive accuracy around the same level as we had been finding on the Hong Kong equity data, but with profitability looking much easier to come by due to the much lower transaction costs in these other markets.

We then designed a simple backtesting framework, including code accounting for *signal weighting and combination* (combining the signals from multiple predictive models produced by our multiple learning algorithms), *capital allocation*, *portfolio management* and *trade exit management*. Our assumptions in these regards have been relatively simplistic, but customized in some regards to the particulars of our

overall prediction framework (particularly to the fact that we are combining predictions produced by multiple heterogeneous prediction algorithms, possessing quite different properties).

We then ran backtests on the US equity and global macro data from 2007-2011. This exposed some things in our backtesting framework in need of adjustment. Once these adjustments were done, we ran backtests on these universes running all the way from 2003-2011. These results looking promising, and so we finally (in October 2014) decided the time had come to run our code on data from the long-held-out "true out of sample" time period of 2012-2013. These results were also reasonably good, thus providing significant validation that we had not overfit our methodology to the time period before 2011. (Of course, the applicability of our HK equity tuned system to the US equity and global macro universes, also serves as significant validation that we had not overfit our methodology to the HK equity universe either.)

Based on these results, it seems fair to provisionally conclude that we have found a trading methodology with reasonable effectiveness across different asset classes and time periods.

To recap the high level experimentation and validation steps that have just been reviewed above in narrative form, what we did was, in order:

1. **Nov 2011-June 2014:** Develop predictive algorithms via testing on HK equity data from 2007-2011.
2. **July-Sep 2014:**
 - a. Validate these predictive algorithms via testing on US equity and global macro data from 2002-2011.
 - b. Develop a trading methodology (model combination, capital allocation, portfolio and trade management) via testing on US equity and global macro data from 2007-2011.
 - c. Validate the trading methodology and predictive algorithms via testing on US equity and global macro data from 2003-2006.
3. **Oct. 2014:** Validate the predictive algorithms and trading methodology via testing on US equity and global macro data from 2012-2013 (data that had been kept totally "virgin" with respect to Aidya's work, across all asset classes, until we reached this final step).