

**NEWS ANALYTICS TOOL KIT
USER MANUAL**

April 2013



<http://www.optirisk-systems.com>



<http://www.carisma.brunel.ac.uk>

[This product has been superseded by Sentiment Enhanced Signals: SES]

NAT Tool User Manual

Prepared by
Cristiano Arbex Valle & Gautam Mitra
OptiRisk Systems

Copyright © 2013 OptiRisk Systems

DO NOT DUPLICATE WITHOUT PERMISSION
All brand names, product names are trademarks or registered
trademarks
of their respective holders.

The material presented in this manual is subject to change without prior notice and is intended for general information only. The views of the authors expressed in this document do not represent the views and/or opinions of OptiRisk Systems.



OptiRisk Systems
One Oxford Road
Uxbridge, Middlesex, UB9 4DA
United Kingdom
www.optirisk-systems.com
+44 (0) 1895 256484

Table of Contents

| | |
|--|----|
| 1. Scope and Purpose | 4 |
| 2. Overview and outline | 4 |
| 3. Information flow and data sources | 5 |
| 4. User Interface: Control Panel | 6 |
| 4.1. Period Panel | 7 |
| 4.2. News Data Source | 8 |
| 4.3. Companies Filters Panel | 8 |
| 4.4. Export Data Panel | 10 |
| 4.5. Search Results (Companies) Panel | 11 |
| 4.6. Selected Companies Panel | 11 |
| 4.7. News Filters Panel | 12 |
| 4.8. News Data Panel | 12 |
| 4.9. File Menu | 14 |
| 5. Ravenpack: Attributes and Filters | 15 |
| 6. Thomson Reuters: Attributes and Filters | 17 |
| 7. News Impact Measure | 19 |
| 8. Strategies for Asset Allocation | 22 |
| 8.1. Momentum Strategy | 22 |
| 8.1.1. Momentum Rate Of Change (ROC) | 23 |
| 8.1.2. Relative Strength Index (RSI) | 23 |
| 8.2. Momentum Strategy in the NAT Tool | 24 |
| 8.3. Contrarian Strategy in the NAT Tool | 26 |
| 9. Display Graphics and Back Testing | 28 |
| 9.1. Plot Sentiment | 29 |
| 10. Empirical Studies Examples | 30 |
| 10.1. Ravenpack Sentiment Index, DAX30, 2007 to 2010 | 30 |
| 10.2. Momentum Strategy, Dow Jones Industrial Average, 2007 to 2010 | 33 |
| 11. References | 35 |

1. Scope and Purpose

Scope:

News Analysis Tool Kit has been developed as a software tool which provides support to empirical studies of the application of News sentiment metadata in the financial models for asset allocation as well as risk control. The scope of the manual is that of explaining how these features are relevant for such studies. There is an additional section not yet integrated into the manual which supplies the user specification of an API which can be used by an end user to embed the leading features of the tool in an application (say investment strategy or risk control application) of their own.

Purpose: The purpose of the manual is to set out the explanations of the designed features of the tool and also that of the API so that the manual can be used as reference document both for using the software system as an investigation tool and subsequently creating embedded applications using software components.

2. Overview and outline

The News Analytics Toolkit (NAT) is designed to provide the end user a tool which extracts company-specific news data and market data. Essentially, this tool collates and consolidates news sentiment data and market data for equities.

The GUI and the filtering features make NAT a very valuable tool with which to study news sentiment (analysis) of equities and to enhance existing financial models. Typical financial models which can be enhanced by News Analytics include:

- High-Frequency trading models
- Low-Frequency fund rebalancing models
- Risk analysis (market risk and liquidity risk) models

The NAT connects to a database of pre-processed quantified news data supplied by Ravenpack or Thomson Reuters. A raw news database comprises a large number of news items covering around 28,000 companies. An end user can use NAT effectively to search, filter and export news data for a portfolio of selected companies. The archived News data is available from 2005 until the present date. Using this tool the user can search over the entire database and apply logical filters for different

NAT Tool User Manual

attributes (parameters). The choice of filter includes: news sentiment score, news relevance score for a given company, novelty of the news, news category, news type and date.

News data are company-specific, and the end user is able to select the companies he/she wishes to get news data for. Companies can be filtered, among others, by country, ISIN code and stock market indices. In the NAT tool some of the major world indices are covered: FTSE100, CAC40, DAX30, Eurostoxx50, Hang Seng, Topix, Sensex and Dow Jones 30. NAT also processes market data for the companies in these assets (universes). This makes it easy for an analyst to design studies which uses both filtered news data and market data for a specified period, followed by back testing.

3. Information flow and data sources

A short review of news analytics focusing on its applications in finance is given in this section; it is an abridged version of the review chapter in the Hand Book compiled by one of the authors (Mitra & Mitra, 2011). It is widely recognized that news plays a key role in financial markets; the sources and volumes of news continue to grow. New technologies that enable automatic or semi-automatic news collection, extraction, aggregation and categorization are emerging. Further, machine-learning techniques are used to process the textual input of news stories to determine quantitative sentiment scores.

There is a strong yet complex relationship between market sentiment and news. The arrival of news continually updates an investor's understanding and knowledge of the market and influences investor sentiment. Traders and other market participants digest news rapidly, revising and rebalancing their asset positions accordingly. As markets react rapidly to news, effective models which incorporate news data are highly sought after. This is not only for trading and fund management, but also for risk control. Major news events can have a significant impact on the market environment and investor sentiment, resulting in rapid changes to the risk structure and risk characteristics of traded assets. Though the relevance of news is widely acknowledged, how to incorporate this effectively, in quantitative models and more generally within the investment decision-making process, remains an unanswered question.

News Analytics (NA) technology can also be used to aid traditional non-quantitative fund managers in monitoring the market sentiment for particular stocks, companies, brands and sectors. This technology aids fund managers from the minutiae of repetitive analysis, such that they are able to better target their reading and research. They reduce the burden of routine monitoring for fundamental managers. The basic idea behind these NA technologies is to automate human thinking and reasoning. Traders, speculators and private investors anticipate the direction of asset returns as well as the size and the level of uncertainty (volatility) before making an investment decision. They carefully read recent economic and financial news to gain a picture of the current situation. Using their knowledge of how markets behaved in the past under different situations, people will implicitly match the current situation with those situations in the past most similar to the current one. News analytics seek to introduce technology to automate or semi-automate this approach. By automating the

NAT Tool User Manual

judgement process, the human decision maker can act on a larger, hence more diversified, collection of assets. These decisions are also taken more promptly (reducing latency). In (Leinweber, 2009), this process is referred as intelligence amplification (IA).

As shown in Figure 1: An outline of information flow and modelling architecture, news data are an additional source of information that can be harnessed to enhance (traditional) investment analysis. Yet it is important to recognize that NA in finance is a multi-disciplinary field which draws on financial economics, financial engineering, behavioural finance and artificial intelligence (in particular, natural language processing).

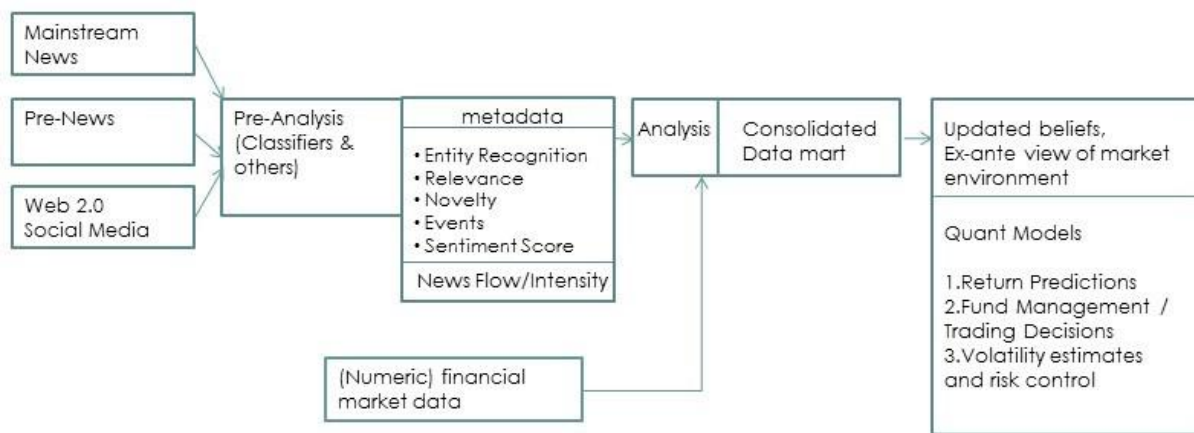


Figure 1: An outline of information flow and modelling architecture

4. User Interface: Control Panel

In summary, the NAT Tool allows the end user to export news data for selected companies based on several different filters. Figure 2 presents a general overview of the NAT Tool.

The left panels allow the user to select the period and companies from which the data will be retrieved. Also, the left panel contains an option to choose the source of news: Thomson Reuters or Ravenpack. The central panels present the companies that resulted from the appropriate query. The right panels contain the news data filters and query results.

NAT Tool User Manual

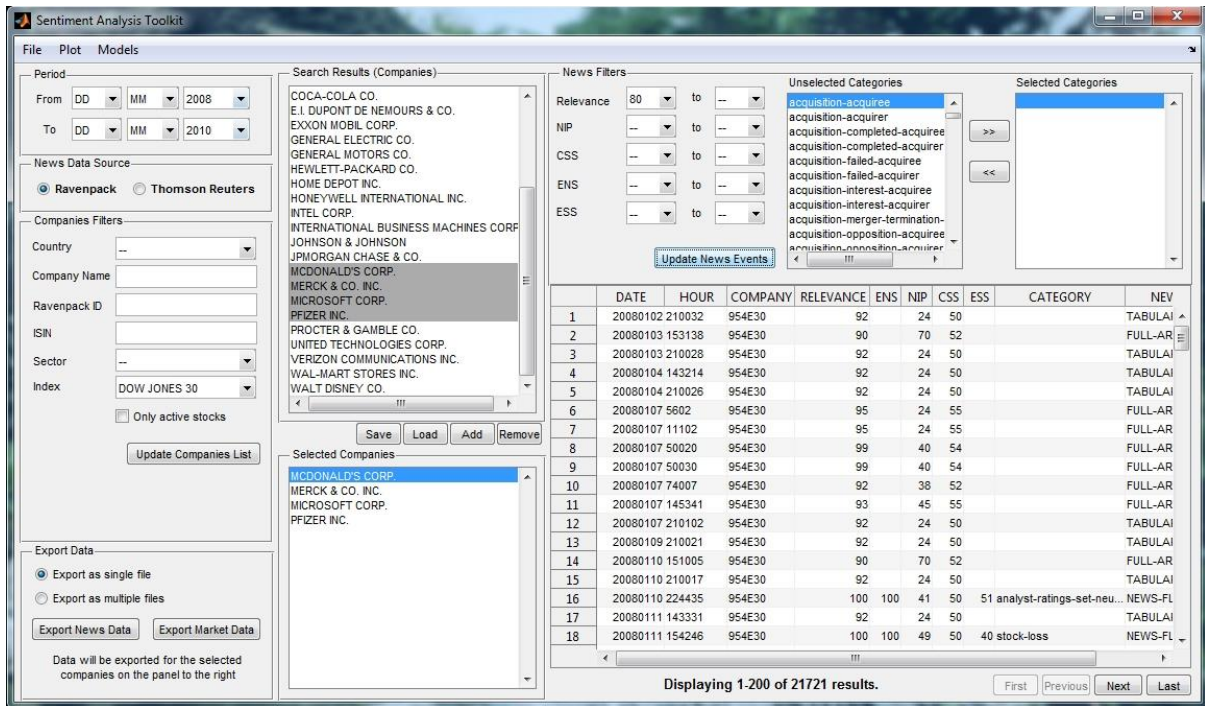


Figure 2: General View of the NAT Tool

The right panels are responsible for the filters that can be used on a news data query and the resulting news data is also displayed here.

To exit the tool, either click on the **File** menu followed by **Quit** or click the **X** in the top right of the screen.

In the next sections, each of the panels will be explained in detail.

4.1. Period Panel



Figure 3: Period Panel

The upper left panel on the screen is the **Period Panel**. This panel is shown in Figure 3. The time period chosen is used to:

- Limit the news data that will be queried from the database
- If a certain index is selected and the **Only active stocks** checkbox is ticked, the query will exclude companies that were not part of the index on the selected period.

NAT Tool User Manual

There are different ways in which the user can select a period:

- If only the **From** fields are filled, the period chosen will comprise of all data from the chosen date.
- If only the **To** fields are filled, the period chosen will comprise of all data older and up to the chosen date.
- If the user does not specify a **day** in the **From** field but specifies a **month** and a **year**, the query will select all data that is newer or equal to the **first** day of the chosen month/year.
- If the user does not specify a **day** in the **To** field but specifies a **month** and a **year**, the query will select all data that is older or equal to the **last** day of the chosen month/year.
- If the user does not specify a **day** nor a **month** in the **From** field but specifies a **year**, the query will select all data that is newer or equal to the **1st of January** of the chosen year.
- If the user does not specify a **day** nor a **month** in the **To** field but specifies a **year**, the query will select all data that is older or equal to the **31st of December** of the chosen year.

Example 2.1.1: if the user wishes to retrieve news data for the years 2008 and 2009, it must only select 2008 as the year on the **From** field and 2009 as the year on the **To** field.

Note: an error message will pop up if the user tries to make a query where a day is specified but no month or year, and when a month is specified but no year, in both **From** and **To** fields.

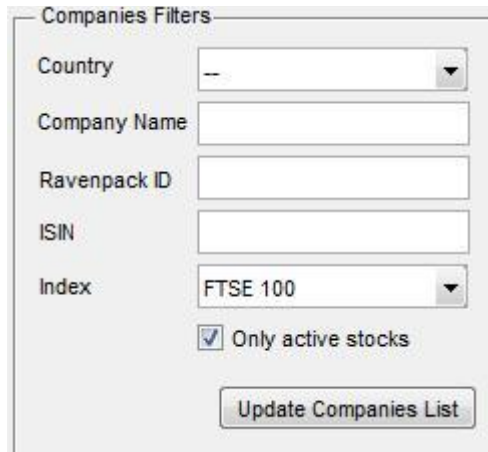
4.2. News Data Source



Figure 4: News Data Source

The NAT tool currently supports news data from two different sources: Ravenpack and Thomson Reuters, and there is an option to alternative between both sources. To do so, just click on the desired news data source in the **News Data Source** panel, as shown in Figure 4.

4.3. Companies Filters Panel



The screenshot shows a panel titled "Companies Filters". It contains the following elements from top to bottom: a "Country" dropdown menu with "--" selected; a "Company Name" text input field; a "Ravenpack ID" text input field; an "ISIN" text input field; an "Index" dropdown menu with "FTSE 100" selected; a checkbox labeled "Only active stocks" which is checked; and a button labeled "Update Companies List".

Figure 5: Companies Filters Panel

The **Companies Filters Panel**, located on the centre-bottom left corner of the screen, allows the user to select filters and make queries over the companies list. This panel is shown in Figure 5. Before making a query, a series of filters can be chosen:

- **Country:** Only companies from the selected country will be retrieved from the database. The list of countries depends on the contents of the database.
- **Company Name:** The user can type the full name or part of a name of a company, only companies whose name matches the typed text will be retrieved from the database. This filter is NOT case sensitive.

Note: It is possible to type 2 or more different names or extracts of a name. The query will treat the different extracts as a logical **OR**.

Example 2.2.1: If the user wishes to retrieve data from both the banks Barclays and HSBC, the user can type in the **Company Name** field the following: "Barclays hsb", without the quotes. Both companies will be retrieved from the database when the **Update Companies List** button is pressed.

- **Ravenpack ID:** The Ravenpack data assigns an ID to each company. If the ID is known, it can be used as a filter to retrieve specific companies from the database. The user can type the full ID or part of the ID of a company, only companies whose Ravenpack ID matches the typed text will be retrieved from the database. This filter is NOT case sensitive.

Note: It is possible to type 2 or more different IDs or extracts of an ID. The query will treat the different extracts as a logical **OR**. Example 2.2.1 applies to this filter when the company names are replaced by the appropriate Ravenpack IDs.

- **ISIN Code:** The ISIN Code of a company can be used as a filter to retrieve specific companies from the database. The user can type the full ISIN code or part of the ISIN code of a company, only companies whose ISIN code matches the typed text will be retrieved from the database. This filter is NOT case sensitive.

NAT Tool User Manual

Note: It is possible to type 2 or more different ISINs or extracts of an ISIN. The query will treat the different extracts as a logical **OR**. Example 2.2.1 applies to this filter when the company names are replaced by the appropriate ISIN Code.

- **Index:** This field can be used to filter companies of a specific stock market index, like the FTSE100 or CAC40. The indices shown will depend on the database used.
- **Only Active Stocks:** If this checkbox is ticked, only stocks that were part of the chosen **Index** for at least one day of the **Period** selected will be retrieved from the database. If no index is chosen, then this checkbox will have no effect.

After choosing the appropriate filters, the query is executed when the **Update Companies List** button is pressed. The results of this query will be shown in the **Search Results (Companies)** panel.

4.4. Export Data Panel

The **Export Data Panel**, located on the bottom left corner of the screen, allows the user to export both market and news data. The panel is shown in Figure 6: Export Data Panel. Data is exported in **csv** format and can be produced in one single file or in multiple files, depending on which option is chosen.

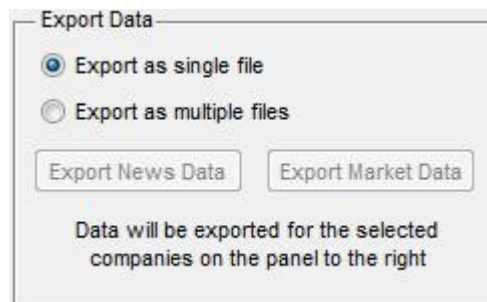


Figure 6: Export Data Panel

- **Export News Data** button: When the user clicks this button, news data will be exported in **csv** format. The news data to be exported is the one shown in the **Export News Data** panel.
- **Export Market Data** button: When the user clicks this button, daily market data will be exported in **csv** format for the companies shown in the **Selected Companies** panel. Moreover, data will be exported for the period selected in the **Period** panel.

The exported **csv** file(s) contain(s) the following data:

- List of selected companies, containing name, country code, Ravenpack ID and ISIN Code. If data is exported as multiple files, instead of the list of selected companies, each file will contain data for a single company, the company being identified in the file name.

NAT Tool User Manual

- Chosen Filters and respective values (an empty entry means no value was specified for that filter)
- Market Data / News Data

4.5. Search Results (Companies) Panel

The **Search Results (Companies)** panel can be seen in Figure 7. The companies shown in this panel are the result of the companies' query that is executed when the **Update Companies List** button is pressed on the **Companies Filters** panel.

The companies in the list are not necessarily the companies that will be used to retrieve news data from the database. Among the search results, the user must select one or more companies for which news data will be retrieved.

To select one company among the search results, just click on the company name. **Shift** and **Ctrl** keys can be used to select multiple companies. Notice, however, that the companies selected in this panel will not be used to export market data or as filters to search for news data. For that, it is necessary to create a portfolio of companies in the **Selected Companies** panel.

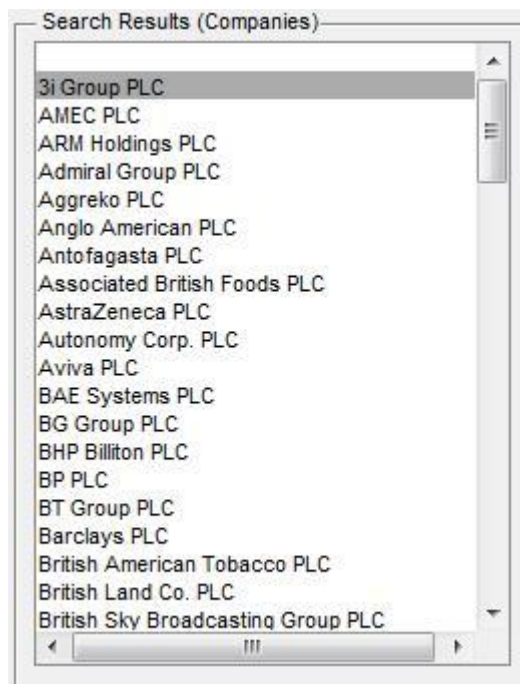


Figure 7: Search Results (Companies) Panel

4.6. Selected Companies Panel

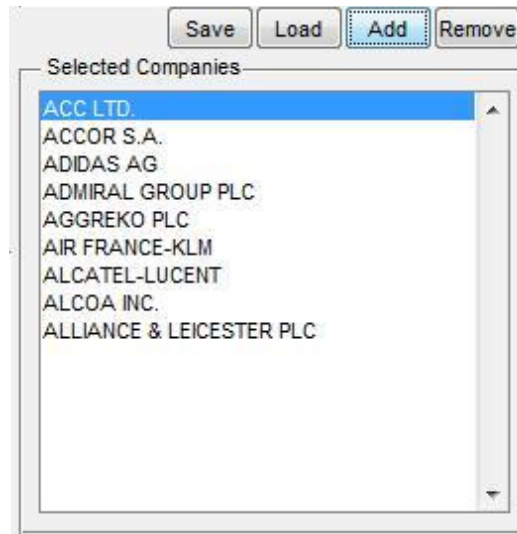


Figure 8: Selected Companies Panel

The **Selected Companies** panel can be seen in Figure 8. The companies that appear on this panel are the actual companies for which news data will be retrieved when the **Update News Events** button is pressed in the **News Filters** panel. This panel can be seen as a portfolio of companies for which we wish to produce data.

To **add** companies to the **Selected Companies** panel, first select some companies from the **Search Results (Companies)** panel. Then, click the **Add** button, which can be found between the **Search Results (Companies)** panel and the **Selected Companies** panel.

To **remove** companies from the **Selected Companies** panel, first select the companies to be removed in the **Selected Companies** panel. This can be done by clicking on the name of the company. **Shift** and **Ctrl** keys can be used for multiple selections. Then, click the **Remove** button.

To **save** the portfolio for further use, click on the **Save** button and choose a filename and path for a **csv** file to be saved.

To **load** a previously saved portfolio, click on the **Load** button and choose a previously saved file.

4.7. News Filters Panel

The News Filter Panel varies according to what news data source is selected. Please refer to Ravenpack: and **Error! Reference source not found.** for specific descriptions of Ravenpack and Thomson Reuters News Filters.

4.8. News Data Panel

NAT Tool User Manual

| | DATE | HOUR | COMPANY | RELEVANCE | ENS | NIP | CSS | ESS | CATEGORY | NEV |
|----|----------|--------|---------|-----------|-----|-----|-----|-----|----------------------------|---------|
| 1 | 20080102 | 210032 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 2 | 20080103 | 153138 | 954E30 | 90 | | 70 | 52 | | | FULL-AR |
| 3 | 20080103 | 210028 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 4 | 20080104 | 143214 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 5 | 20080104 | 210026 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 6 | 20080107 | 5602 | 954E30 | 95 | | 24 | 55 | | | FULL-AR |
| 7 | 20080107 | 11102 | 954E30 | 95 | | 24 | 55 | | | FULL-AR |
| 8 | 20080107 | 50020 | 954E30 | 99 | | 40 | 54 | | | FULL-AR |
| 9 | 20080107 | 50030 | 954E30 | 99 | | 40 | 54 | | | FULL-AR |
| 10 | 20080107 | 74007 | 954E30 | 92 | | 38 | 52 | | | FULL-AR |
| 11 | 20080107 | 145341 | 954E30 | 93 | | 45 | 55 | | | FULL-AR |
| 12 | 20080107 | 210102 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 13 | 20080109 | 210021 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 14 | 20080110 | 151005 | 954E30 | 90 | | 70 | 52 | | | FULL-AR |
| 15 | 20080110 | 210017 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 16 | 20080110 | 224435 | 954E30 | 100 | 100 | 41 | 50 | 51 | analyst-ratings-set-neu... | NEWS-FL |
| 17 | 20080111 | 143331 | 954E30 | 92 | | 24 | 50 | | | TABULAI |
| 18 | 20080111 | 154246 | 954E30 | 100 | 100 | 49 | 50 | 40 | stock-loss | NEWS-FL |

Displaying 1-200 of 21721 results. [First] [Previous] [Next] [Last]

Figure 9: News Data panel

When the **Update News Events** button is pressed in the **News Filters** panel, the query results are displayed in the **News Data** panel, shown in Figure 9. News data from Ravenpack and Thomson Reuters have different fields, so the user will see a different table depending on the data source. The fields for Ravenpack are:

- Date (YYYYMMDD)
- Hour (HHMMSS): If the number of characters is less than 6, than by completing with left zeros the correct format can be obtained. For example: Item 13 in Figure 9 is 73320, this means 07:33:20 (7:33am, 20 seconds)
- Company (actually the company's Ravenpack ID)
- Relevance
- ENS
- NIP
- Category
- ESS

For Thomson Reuters, the fields are:

- Headline
- Date (YYYYMMDD)
- Hour (HHMMSS): If the number of characters is less than 6, than by completing with left zeros the correct format can be obtained.

NAT Tool User Manual

- Company (actually the company's Ravenpack ID).
 - **Note:** Even though this is a Thomson Reuters filter, the Ravenpack ID is the only unique identifier we have available for each company.
- Relevance: Score from 0 to 1
- Sentiment: -1 for probably negative events, 0 for probably neutral and 1 for probably positive.
- Sent. Score: Score from 0 to 100 calculated as shown in Section **Error! Reference source not found.**
- Nov. 12h: Number of related news events in the 12 hours prior to the news event timestamp
- Nov. 1d: Number of related news events in the 24 hours prior to the news event timestamp
- Nov. 3d: Number of related news events in the 3 days prior to the news event timestamp
- Nov. 5d: Number of related news events in the 5 days prior to the news event timestamp
- Nov. 7d: Number of related news events in the 7 days prior to the news event timestamp
- Prob. Pos: Probability of news event being positive
- Prob. Neu: Probability of news event being neutral
- Prob. Neg: Probability of news event being negative

When the news data query yields more than 200 news events, the results are split into pages of 200 events each. To navigate between pages, the **First**, **Previous**, **Next** and **Last** buttons are provided.

4.9. File Menu




Figure 10: File Menu

The **File Menu** (shown in Figure 10) provides alternative ways for some of the NAT Tool functionalities. The menu items are:

- **Load Companies:** Same functionality as the **Load** button between the **Search Results (Companies)** panel and the **Selected Companies** panel.
- **Save Companies:** Same functionality as the **Save** button between the **Search Results (Companies)** panel and the **Selected Companies** panel.
- **News Data Source:** Alternative way to choose between Ravenpack and Thomson Reuters as the source of news data.
- **Export News Data:** Alternative way to export news data, this menu item allows choosing to export the data as single or multiple files.

NAT Tool User Manual

- **Export Market Data:** Alternative way to export market data, this menu item allows choosing to export the data as single or multiple files.
- **Quit:** Terminates the application.

5. Ravenpack: Attributes and Filters

The **News Filters** panel is where the query for news data is prepared. The contents of the panel depend on the source of news data (Thomson Reuters or Ravenpack) that is selected as each source has different filters.

When Ravenpack is selected in the **News Data Source** panel, the **News Filters** panel will look as in Figure 11. A series of filters can be selected for the news data query:

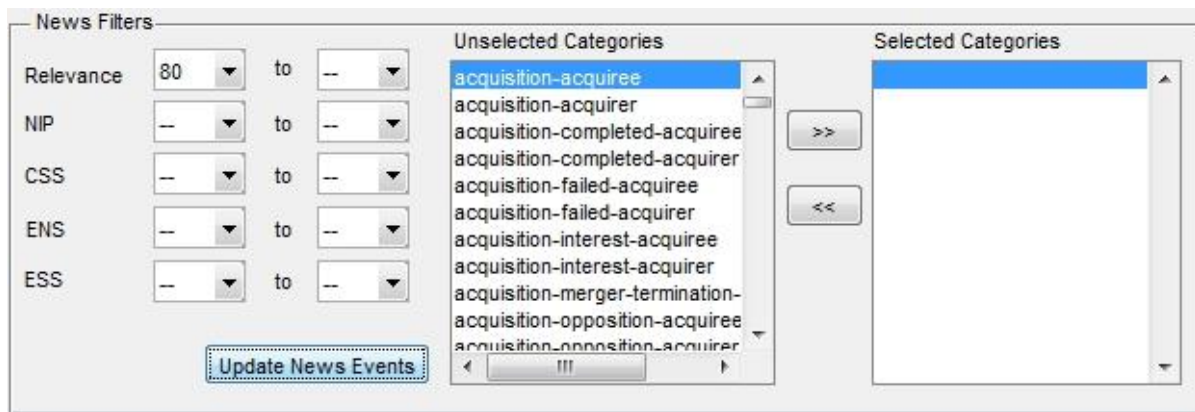


Figure 11: Ravenpack News Filters Panel

- **Relevance:** The relevance of a news event for a particular company, a score calculated by Ravenpack. Relevance scores range from 1 to 100. The user can select up to two fields, a lower bound and an upper bound.
 - If the user chooses values for both the lower and upper bounds, the tool will retrieve news data with a Relevance score greater and equal to the lower bound and lower and equal to the upper bound.
 - If the user only chooses a value for the lower bound, the tool will retrieve news data with a Relevance score greater and equal to the lower bound.
 - If the user only chooses a value for the upper bound, the tool will retrieve news data with a Relevance score lesser and equal to the upper bound.
- **NIP:** NIP stands for News Impact Projection, a score calculated by Ravenpack. The NIP scores range from 1 to 100. The user can select up to two fields, a lower bound and an upper bound.
 - If the user chooses values for both the lower and upper bounds, the tool will retrieve news data with a NIP score greater and equal to the lower bound and lower and equal to the upper bound.

NAT Tool User Manual

- If the user only chooses a value for the lower bound, the tool will retrieve news data with an NIP score greater and equal to the lower bound.
- If the user only chooses a value for the upper bound, the tool will retrieve news data with an NIP score lesser and equal to the upper bound.
- **CSS:** The CSS stands for Composite Sentiment Score. It is a score calculated by Ravenpack. CSS scores range from 1 to 100. The user can select up to two fields, a lower and upper bound.
 - If the user chooses values for both the lower and upper bounds, the tool will retrieve news data with a CSS score greater and equal to the lower bound and lower and equal to the upper bound.
 - If the user only chooses a value for the lower bound, the tool will retrieve news data with a CSS score greater and equal to the lower bound.
 - If the user only chooses a value for the upper bound, the tool will retrieve news data with a CSS score lesser and equal to the upper bound.
- **ENS:** In Ravenpack news data, ENS stands for Event Novelty Score. It is a score from 0 to 100, where a score of 100 represents a completely novel news event and a score of 0 represents a news event that has been widely repeated.
 - If the user chooses values for both the lower and upper bounds, the tool will retrieve news data with an ENS score greater and equal to the lower bound and lower and equal to the upper bound.
 - If the user only chooses a value for the lower bound, the tool will retrieve news data with an ENS score greater and equal to the lower bound.
 - If the user only chooses a value for the upper bound, the tool will retrieve news data with an ENS score lesser and equal to the upper bound.
- **ESS:** In Ravenpack news data, ESS stands for Event Sentiment Score. It is a score from 0 to 100, where a score of 100 represents very positive news and a score of 0 represents very negative news. ESS is apparently similar to CSS but they differ in how they are calculated.
 - If the user chooses values for both the lower and upper bounds, the tool will retrieve news data with an ESS score greater and equal to the lower bound and lower and equal to the upper bound.
 - If the user only chooses a value for the lower bound, the tool will retrieve news data with an ESS score greater and equal to the lower bound.
 - If the user only chooses a value for the upper bound, the tool will retrieve news data with an ESS score lesser and equal to the upper bound.
- **Categories:** Ravenpack divides news data according to categories. The user can select one or more categories to filter the news data. The **Categories** field is divided in two small panels:
 - **Unselected Categories Panel:** The left panel contains all categories that will NOT be used to filter the news data query.
 - **Selected Categories Panel:** If there are any categories in the right panel, only news data that are of any of these categories will be retrieved from the database when the **Update News Events** button is pressed.

NAT Tool User Manual

To add categories from the **Unselected Categories** panel to the **Selected Categories** panel, first select one or more (using the **Shift** or **Ctrl** keys) categories on the **Unselected Categories** panel and press the **>>** button.

To remove categories from the **Selected Categories** panel, first select one or more (using the **Shift** or **Ctrl** keys) categories on the **Selected Categories** panel and press the **<<** button.

Note 1: The categories present in the **Selected Categories** panel will not appear in the **Unselected Categories** panel.

Note 2: If no categories are selected, companies from all the categories will be retrieved from the database when the **Update News Events** button is pressed.

Note 3: There are several news items that are not categorized. These will be retrieved from the database only when there are no categories in the **Selected Categories** panel.

When the appropriate filters are chosen, the news data query is executed by pressing the **Update News Event** button. The results will be displayed in the bottom-right panel.

6. Thomson Reuters: Attributes and Filters

When Thomson Reuters is selected in the **News Data Source** panel, the **News Filters** panel will look as in Figure 12: Thomson Reuters News Filters Panel. A series of filters can be selected for the news data query:

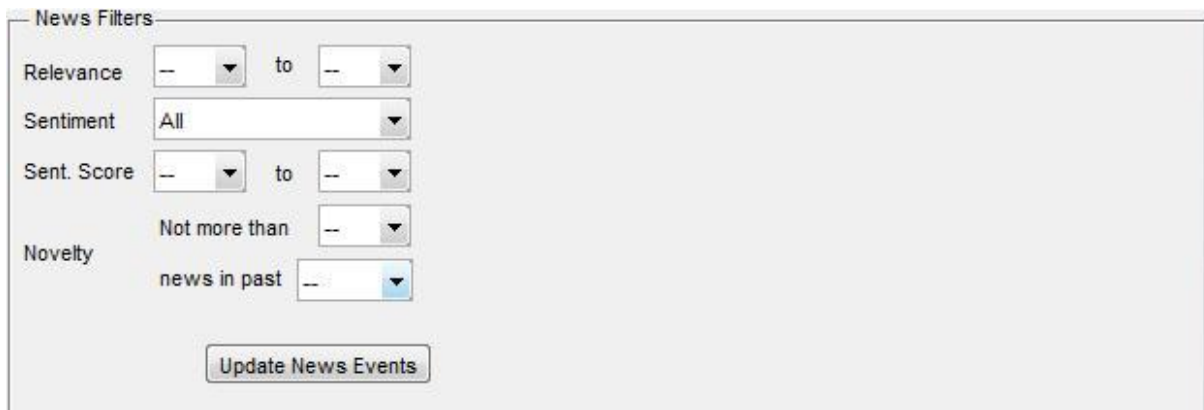


Figure 12: Thomson Reuters News Filters Panel

- **Relevance:** The relevance of a news event for a certain company is a score calculated by Thomson Reuters and is already present in the news data database. Relevance scores range from 0 to 1, but it has been rounded and presented as a score from 0 to 100. The user can select up to two fields, a lower and upper bound.

NAT Tool User Manual

- If the user chooses values for both the lower and upper bounds, the tool will retrieve news data with a Relevance score greater and equal to the lower bound and lower and equal to the upper bound.
- If the user only chooses a value for the lower bound, the tool will retrieve news data with a Relevance score greater and equal to the lower bound.
- If the user only chooses a value for the upper bound, the tool will retrieve news data with a Relevance score lesser and equal to the upper bound.
- **Sentiment:** Thomson Reuters classifies news events as most probably positive, negative or neutral. The options on this filter are:
 - **Strictly (Positive, Neutral, Negative):** If any of these options are chosen, the news data query will return only news events that are positive, neutral and/or negative.
 - **Excluding (Positive, Neutral, Negative):** If any of these options are chosen, the news data query will return all news events with the exception of events which are positive, neutral and/or negative.
- **Sent. Score:** Thomson Reuters provides positive, neutral and negative probabilities for each news event. The probabilities sum to one. By applying a simple formula to the probabilities, it is possible to obtain a sentiment score from 0 to 100, similar to the **CSS** field in the Ravenpack news data. The formula is given by the expression:
 - **Score = $100 \times (Prob_{pos} + \frac{Prob_{neu}}{2})$** (more details can be found at the end of this section)
 - The user can select up to two fields, a lower and an upper bound.
 - If the user chooses values for both the lower and upper bounds, the tool will retrieve news data with a Sent. Score greater and equal to the lower bound and lower and equal to the upper bound.
 - If the user only chooses a value for the lower bound, the tool will retrieve news data with a Sent. Score greater and equal to the lower bound.
 - If the user only chooses a value for the upper bound, the tool will retrieve news data with a Sent. Score lesser and equal to the upper bound.
- **Novelty:** While Ravenpack assigns a score from 0 to 100 on how new a news event is, Thomson Reuters deals with novelty differently. The filter allows the user to choose news events which at most have a certain number of related news in a fixed past period of time. For example, if the user chooses **Not more than 5 news in past 3 days**, only news events that have at most 5 other events on the same topic, in the 3 days prior to the news event timestamp, will be selected.

Thomson Reuters' news sentiment engine analyses and processes each news story that arrives as a machine readable text. Through text analysis and other classification schemes the engine then computes for each news event: (i) relevance, (ii) entity recognition and (iii) sentiment probabilities as well as a few other attributes. News event sentiment can be positive, neutral and negative and the classifier assigns probabilities such that

$$Prob(positive) + Prob(neutral) + Prob(negative) = 1.0$$

We turn these three probabilities into a single sentiment score in the range 0-100 using the following equation:

$$Sent = 100 * (Prob(positive) + \frac{1}{2}Prob(neutral))$$

where *Sent* denotes a single transformed sentiment score. We find that such a derived single score provides a relatively better interpretation of the mood of the news item. Thus the news sentiment score is a relative number which describes the degree of positivity and negativity in a piece of news.

7. News Impact Measure

In order to compute the impact of news events over time we first find an expression which describes the attenuation of the news sentiment score. The impact of a news item does not solely have an effect on the markets at the time of release but also over finite periods of time that follow. To account for this prolonged impact, we have applied an attenuation technique to reflect the instantaneous impact of news releases and the decay of this impact over a subsequent period of time. The technique combines exponential decay and accumulation of the sentiment score over the time buckets under observation. We take into consideration the attenuation of positive sentiment to the neutral value and the rise of negative sentiment also to the neutral value and accumulate (sum) these sentiment scores separately. By separating the positive and negative sentiment scores we avoid an exact cancellation which may be interpreted as no news.

We first shift the sentiment score by subtracting 50 (the value corresponding to the neutral sentiment score = 0) to compute

$$SENT = Score - 50$$

Thus the score *SENT* lies between -50 and +50. During the trading day, as news arrives it is given a sentiment value.

Although news arrives asynchronously we work out the aggregated impact of all news in the following way. Let

POS denote the set of news with positive sentiment value $SENT > 0$;

NEG denote the set of news with negative sentiment value $SENT < 0$;

$PNews(k, t_k)$ denote the sentiment value of the k^{th} positive news at time of arrival t_k , and $k \in POS$;
 $PNews(k, t_k) > 0$.

$NNews(k, t_k)$ denote the sentiment value of the k^{th} negative news at time of arrival t_k , and $k \in NEG$;
 $NNews(k, t_k) < 0$.

Let λ denote the exponent which determines the decay rate. More details into this rate are given in the end of this section. The cumulated positive and negative sentiment scores at time t are given by:

$$PSent(t) = \sum_{\substack{k \in POS \\ t_k \leq t}} PNews(k, t_k) e^{-\lambda(t-t_k)},$$

NAT Tool User Manual

$$NSent(t) = \sum_{\substack{k \in NEG \\ t_k \leq t}} NNews(k, t_k) e^{-\lambda(t-1)},$$

The arrival of more news items lead to higher values of accumulation; this therefore takes into account the news intensity, that is, the news flow.

The impact is illustrated in **Figure 13**: which shows the impact score (with attenuation and accumulation) for the asset JP Morgan during the month of August 2008, with positive and negative sentiment represented separately.




Figure 13: JP Morgan August 2008: News impact score (accumulated and aggregated) for positive (blue line) and negative (red line) sentiment respectively

In order to compute the aggregated impact of all the news items at time t , we sum the positive and negative:

$$CumSent_t = \sum_{i=1}^t Psent(t) + NSent(t)$$




Figure 14: Decay rate

Figure 14: Decay rate shows the decay rate of news sentiment. The parameter λ_{90} specifies that the news sentiment decays exponentially to half of its initial value over a 90 minute period.




Figure 15: AIG Sentiment vs. Bid Price

In Figure 15: AIG Sentiment vs. Bid Price we plot the overall sentiment for the asset AIG over one month of data in which the vertical axis on the left is the sentiment score, the vertical axis on the right is the bid price and the horizontal axis is time (calculations are made on a minute basis).

The work of (Leinweber D. ..., 2011) and also our earlier work (Mitra L. M., 2009) as well as the present study indicates that accumulation of the sentiment leads to a fairly good fit with asset price in general and with asset price volatility in particular.




Figure 16: AIG Cumulated Sentiment and Asset Prices

The cumulated sentiment and asset prices are plotted and displayed in Figure 16: AIG Cumulated Sentiment and Asset Prices.

8. Strategies for Asset Allocation

The NAT Tool provides three news enhanced strategies for asset allocation, namely Momentum, Contrarian and Value Strategy.

8.1. Momentum Strategy

Momentum is defined as the tendency of either high or poorly performing stocks to continue exhibiting high or poor performance for a long period, i.e. of the order of a year or longer (Jegadeesh & Titman, 1993). Momentum is regarded as an anomaly because this behaviour is not explained by the capital asset pricing model (Sharpe, 1964).

Conventionally, a momentum portfolio is selected by a simple ranking of historic total stock returns over a given time period, then buying the top winners and going short on the losers while the assets from the ends are formed into equally weighted portfolios. The possible causes of momentum have been explored in some depth. Fama and French (Fama & French., 1996) discussed the embarrassment of the three-factor model in failing to explain momentum.

In particular, explaining momentum pricing has become one of the principal battlefields in finance between the behaviourists and the rationalists.

The behaviourists focus almost exclusively on the mechanism by which new information or news is embedded in prices, if investors are prone to exhibiting various psychological biases. Following a decision to buy, investors are more likely to later buy more of the stock if they receive further good news than they are likely to sell if they receive bad news. Investors are believed to be generally conservative; they are slow in updating their prior beliefs in the event of good (bad) news. Thus, prices do not adjust completely when new information arrives, but will adjust later if confirming news arrives.

Rationalists, on the other hand, focus on how "minimally rational" investors reacting to unpredictable changes in market conditions could induce the momentum anomalies. Some empirical observations suggest that as market conditions improve, news slowly diffuses into the prices of similar stocks. If investors could rely on information already processed with positive sentiment, they could improve the performances.

Momentum can be calculated in different ways and the NAT Tool currently provides two different strategies, Rate of Change and Relative Strength Index.

8.1.1. Momentum Rate Of Change (ROC)

Momentum is usually defined as the absolute difference in stock:

$$\text{momentum} = \text{close}_{\text{today}} - \text{close}_{n \text{ days ago}}$$

Rate of change scales by the oldest value closing price, so as to represent the increase as a fraction:

$$\text{rate of change} = \frac{\text{close}_{\text{today}} - \text{close}_{n \text{ days ago}}}{\text{close}_{n \text{ days ago}}} \times 100$$

Momentum, in general, refers to prices which continue their trend. Momentum and ROC indicators show trend by remaining positive while an uptrend is sustained or negative while a downtrend is sustained. A crossing up through zero may be used as a signal to buy, or a crossing down through zero as a signal to sell. How high (or how low) the indicators get shows how strong the trend is.

8.1.2. Relative Strength Index (RSI)

The Relative Strength Index (RSI) is a technical indicator used in the technical analysis of financial markets. It was developed by J. Welles Wilder (Wilder, 1978). It should not be confused with Relative Strength.

Relative Strength Index is classified as a momentum oscillator, measuring the velocity and magnitude of directional price movements. Momentum is the rate of the rise or fall in price. It is computed as the ratio of higher closing prices to lower closing prices; stocks which had stronger positive changes have a higher RSI than stocks which had stronger negative changes. The RSI is most typically used on a 14 day timeframe, measured on a scale from 0 to 100, with high and low levels

NAT Tool User Manual

marked at 70 and 30, respectively. Shorter or longer timeframes are used for alternately shorter or longer outlooks. More extreme high and low levels-80 and 20, or 90 and 10-occur less frequently but indicate stronger momentum.

For each trading period an upward change U or downward change D is calculated. Up periods are characterized by the close being higher than the previous close:

$$U = \text{close}_{\text{now}} - \text{close}_{\text{previous}} \\ D = 0$$

Conversely, a down period is characterized by the close being lower than the previous period (note that in this situation D is a positive number):

$$D = \text{close}_{\text{previous}} - \text{close}_{\text{now}} \\ U = 0$$

If both closing prices are equal, U and D are zero. The average U and D are calculated using an n -period exponential moving average (EMA). The ratio of these averages is the Relative Strength:

$$RS = \frac{EMA(U, n)}{EMA(D, n)}$$

If the average of D values is zero, then the RSI value is defined as 100. The Relative Strength is then converted to a Relative Strength Index between 0 and 100:

$$RSI = 100 - \frac{100}{1 + RS}$$

The exponential moving averages should be appropriately initialized with a simple average using the first n values in the price series. The level of the RSI is a measure of the stock's recent trading strength.

8.2. Momentum Strategy in the NAT Tool

The Momentum Strategy for portfolio selection is one of the most popular strategies in the financial market. In the NAT Tool, it is possible to generate portfolios based on news enhanced momentum strategies that can be compared to the traditional momentum strategies.

To access the Momentum Strategy plotting tool, the user must select the **Models** menu followed by the **Momentum** menu item. An interface like the one in Figure 17 should appear.

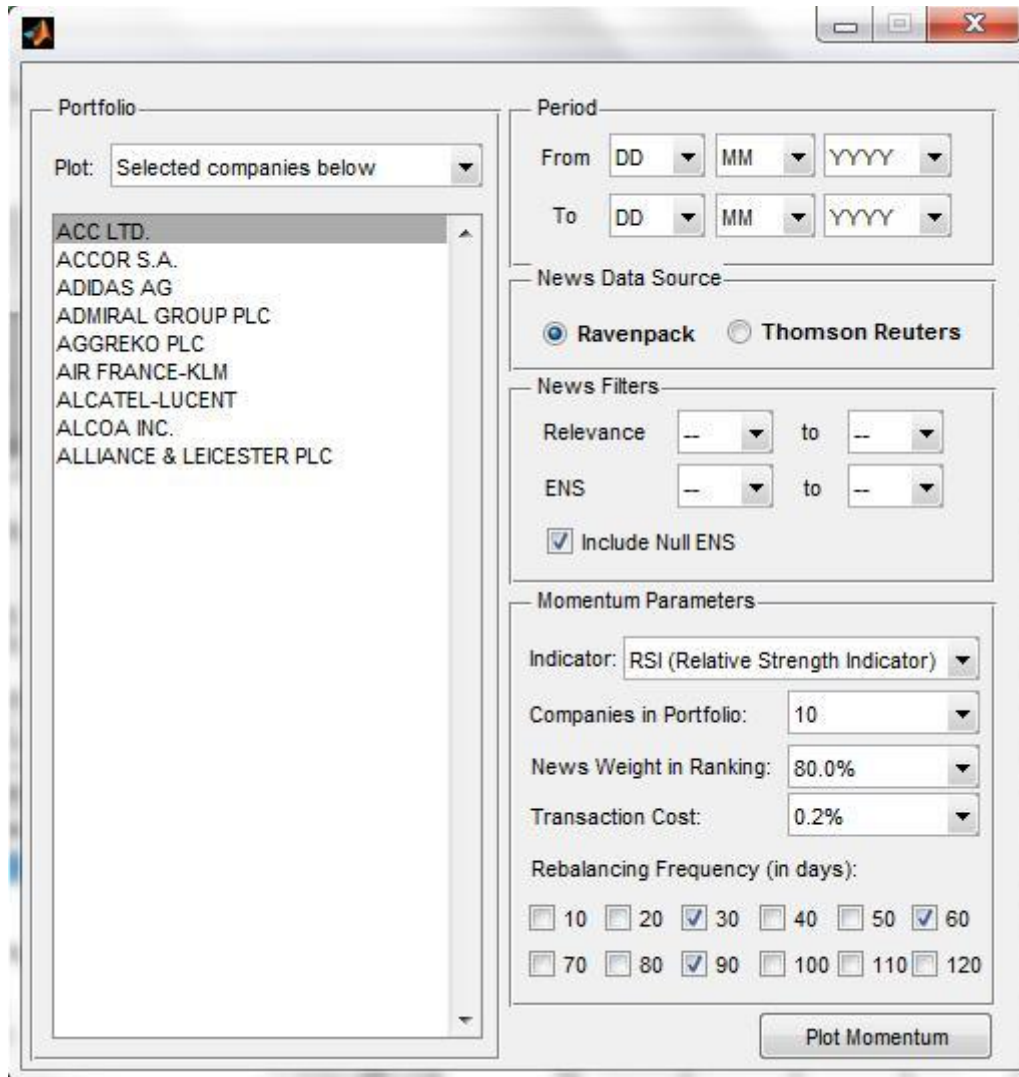


Figure 17: Momentum Strategy

The Momentum Model will plot two (or three) lines in a graphic, containing returns for traditional momentum portfolios and returns for news enhanced momentum portfolios. It is also possible to include, in some cases, a third line representing a market index for comparison.

The Momentum Interface is split in several panels. They are the following:

- **Portfolio** panel: This panel can be found on the left side of the screen. It contains exactly the same companies that are shown in the **Selected Companies** panel in the main NAT Tool window. **Data from the companies present in this list will be used to generate the charts.**
 - Instead of the companies shown in the list, the user can choose from a series of well known market indices. On the top of the panel, the user can select the **Plot selected companies below** or choose the CAC40, DAX30, DJI30, FTSE100 or SENSEX60 indices. If any of the indices is selected, the companies shown on the list will be **ignored** and the momentum strategy will be calculated for the companies that are part of the

NAT Tool User Manual

selected index. Moreover, when (and only when) an index is selected, a third line containing the index returns will be added to the resulting graphic.

- **Note:** Even though it is possible to highlight companies in the **Portfolio** panel, it does not have any effect on the tool or charts. The charts will always contain information from **ALL** companies shown in the list.
- **Period** panel: Dates selected in this panel will be used to filter the sentiment plots' underlying data. This panel works precisely as the **Period** panel in the NAT Tool's main window.
- **News Data Source** panel: Defines the news data source that will be used to generate the sentiment plots.
- **News Filters** panel: Filters that apply to the sentiment charts. News data can be filtered by Relevance and Novelty for both Ravenpack and Thomson Reuters.
- **Momentum Parameters** Panel: This panel contains the parameters that can be chosen to personalize the Momentum Strategy. These parameters are:
 - **Indicator:** There are many ways to calculate momentum. The NAT Tool provides two of them: RSI (Relative Strength Indicator) and ROC (Rate of Change). The user can select which of these two strategies will be used in calculating the momentum portfolios.
 - **Companies in Portfolio:** Specifies the fixed number of companies to be held in the momentum portfolios.
 - **News Weight in Ranking:** Momentum strategies consist of ranking companies according to some criteria and choosing the best ranked companies. The news enhanced momentum consists of creating two rankings: a traditional one and a news ranking based on positive sentiment score. Both rankings are merged to create the news enhanced momentum portfolio. This parameter allows the user to choose which weight will be applied to the news ranking. A value of 80% means that the final merged ranking will be split into 80% importance for the news ranking and 20% importance for the traditional ranking.
 - **Transaction Cost:** The percentage of trading cost that will be applied every time the portfolios are rebalanced.
 - **Rebalancing Frequency (in days):** The fixed time intervals at which the portfolios will be rebalanced. It is possible to choose multiple values. For example, if the values of 60 and 90 are chosen, it is as if the investor split his or her money to be invested into two different portfolios, one rebalanced every 60 days, the other rebalanced every 90 days. The plotted line is the average return between these two fictional portfolios.

When the user clicks the **Plot Momentum** button, the generated graph will appear on the screen. It will contain the returns for the traditional momentum portfolio, the returns for the news enhanced portfolio and possibly the returns for the chosen index.

8.3. Contrarian Strategy in the NAT Tool

The Contrarian Strategy is almost the same as the Momentum Strategy. The difference lies in the fact that the contrarian strategy will choose the worst performing companies in terms of market data,

while momentum usually will choose the best performing companies. The news enhanced contrarian strategy will choose companies based on bad momentum and good news.

The parameters and panels are precisely the same as the Momentum interface.

8.4. Value Strategy (see (Uryasev))

The Value Strategy is a strategy that attempt to maximise returns by buying securities that seem underpriced and / or selling securities that seem overpriced based on some sort of fundamental analysis. Here we present a news-based value strategy that intends to identify underpricing / overpricing through the use of news impact analysis. This strategy is traded daily, the position is opened and closed on the same day. It is composed of three main stages:

- **Calibration:** In this stage we use past data to compute historical news impact, which will be later used to calibrate the model. Calibration should be done on a rolling window basis (moving along with time).
- **Definition of Range:** In this stage historical impact calculations are sorted and outliers excluded to define a range of impact measures from negative to positive.
- **Buy / Sell Decision:** Current impact is calculated and compared to historical range in order to decide the daily position (long / short).

Implementation details of the Value Strategy are supplied to purchasers of the NAT Tool / API.

In the NAT Tool we apply value strategy to stock market indices. Let C be the set of companies in a certain index, the index impact at time t is calculated as:

$$CumSent_{Index,t} = \sum_{c=1}^C CumSent_{c,t}$$

where $CumSent_{c,t}$ is the cumulative impact of company $c \in C$ at time t .

The screenshot shows the 'Value Strategy' tool interface. It consists of several panels:

- Portfolio:** A dropdown menu showing 'CAC40'.
- News Data Source:** Two radio buttons, 'Ravenpack' (selected) and 'Thomson Reuters'.
- News Filters:** Two pairs of dropdown menus for 'Relevance' and 'ENS', each with a 'to' label and another dropdown. A checkbox 'Include Null ENS' is checked.
- Period:** Two date pickers labeled 'From' and 'To', each with 'DD', 'MM', and 'YYYY' dropdowns.
- Parameters:** Three dropdown menus: 'News Hours' (10:05), 'In Sample Days' (100), and 'Transaction Cost' (0.05%).
- Plot Value Strategy:** A button at the bottom right.

Figure 18: Value Strategy

To access the Value Strategy tool, the user must select the **Models** menu followed by the **Value Strategy** menu item. An interface like the one in Figure 18: Value Strategy should appear.

The Value Strategy Interface is split in several panels. They are the following:

- **Portfolio** panel: Allows the user to choose to which index the strategy will be applied.
- **Period** panel: Dates selected in this panel will be used to filter the sentiment plots' underlying data. This panel works precisely as the **Period** panel in the NAT Tool's main window.
- **News Data Source** panel: Defines the news data source that will be used to generate the sentiment plots.
- **News Filters** panel: Filters that apply to the sentiment charts. News data can be filtered by Relevance and Novelty for both Ravenpack and Thomson Reuters.
- **Parameters** Panel: This panel contains the parameters that can be chosen to personalize the Momentum Strategy. These parameters are:
 - **News Hours:** The time at which impact will be calculated. Notice that trading uses closing prices so a position is opened at the end of a day and closed at the end of the next day.
 - **In Sample Days:** Specified the calibration period.
 - **Transaction Cost:** The percentage of trading cost that will be applied every time the portfolios are rebalanced.

When the user clicks the **Plot Value Strategy** button, the generated graph will appear on the screen. It will contain the returns for the index and the returns for the news enhanced value strategy portfolio.

9. Display Graphics and Back Testing

The NAT tool provides efficient means to display graphs generated with historical news and market data. This gives the user the possibility of back testing in order to get a better understanding of the influence of news in market prices. This tool is described below.

9.1. Plot Sentiment

It is possible to generate Sentiment Charts by merging news and market data. To access the Plot Sentiment interface, it is necessary to click on the **Plot** menu and then on the **Sentiment Chart** menu item. A screen like the one shown in Figure 19 should be visible, where the following panels can be found:

- **Portfolio** panel: This panel can be found on the left side of the screen. It contains exactly the same companies that are shown in the **Selected Companies** panel in the main NAT Tool window. **Data from the companies present in this list will be used to generate the charts.**
 - **Note:** Even though it is possible to highlight companies in the **Portfolio** panel, it does not have any effect on the tool or charts. The charts will always contain information from **ALL** companies shown in the list.
- **Period** panel: Dates selected in this panel will be used to filter the sentiment plots' underlying data. This panel works precisely as the **Period** panel in the NAT Tool's main window.
- **News Data Source** panel: Defines the news data source that will be used to generate the sentiment plots.
- **News Filters** panel: Filters that apply to the sentiment charts. News data can be filtered by Relevance and Novelty for both Ravenpack and Thomson Reuters.
- **Include** panel: This panel defines which type of data will be plotted in the resulting graphs. A maximum of **two** lines can be selected. The options are:
 - **Market Data:** The value of an equally-weighted portfolio containing all companies in the **Portfolio** panel.

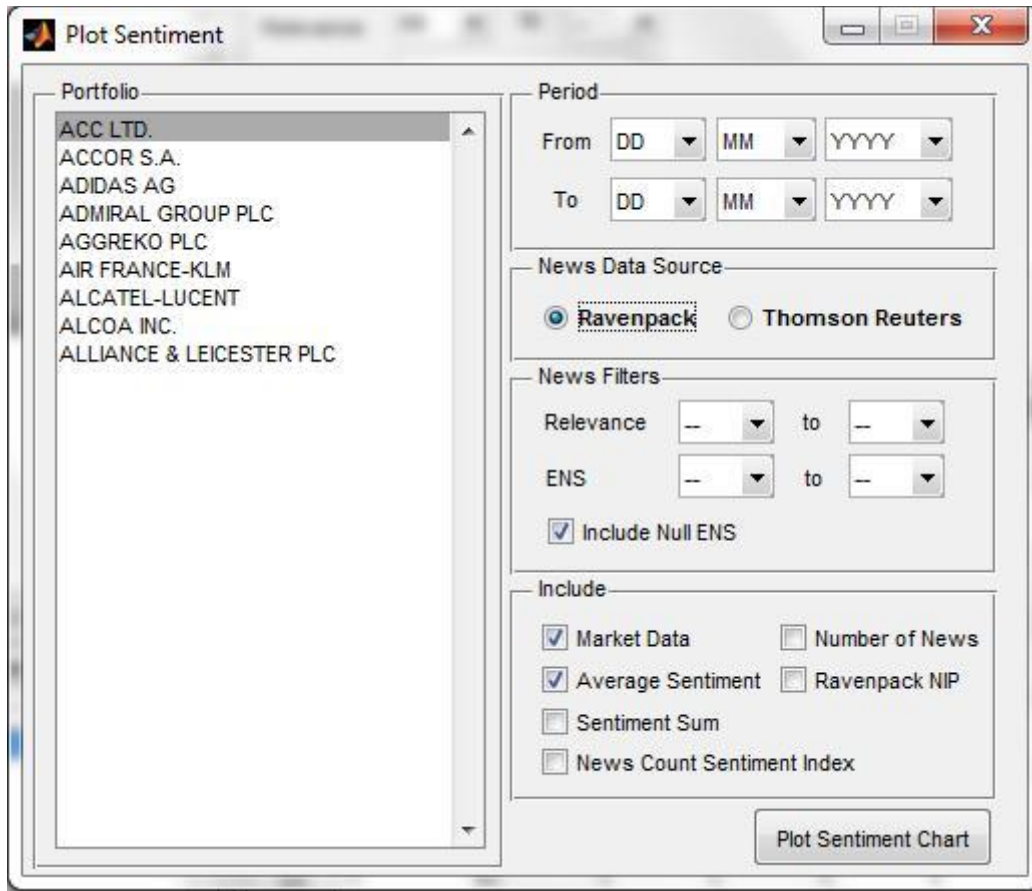


Figure 19: Plot Sentiment

- **Average Sentiment:** Moving Average of the News Sentiment. The field used is CSS for Ravenpack and Sent. Score for Thomson Reuters
- **Sentiment Sum:** Moving Absolute Sum of News Sentiment. The field used is CSS for Ravenpack and Sent. Score for Thomson Reuters
- **News Count Sentiment Index:** Sentiment Index as calculated in *Introducing the RavenPack Sentiment Index*, by Peter Hafez (2011).
- **Number of News:** Moving average number of news.
- **Ravenpack NIP:** Moving Average of Ravenpack NIP score. This is not available if Thomson Reuters is selected in the **News Data Source** panel.

When the user clicks the **Plot Sentiment Chart** button, the generated graph will appear on the screen.

10. Empirical Studies Examples

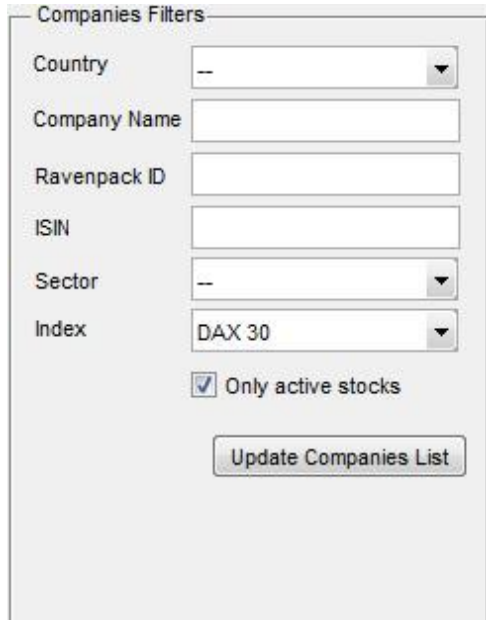
In this section, we will show a few empirical examples as a guideline to using the NAT Tool.

10.1. Ravenpack Sentiment Index, DAX30, 2007 to 2010

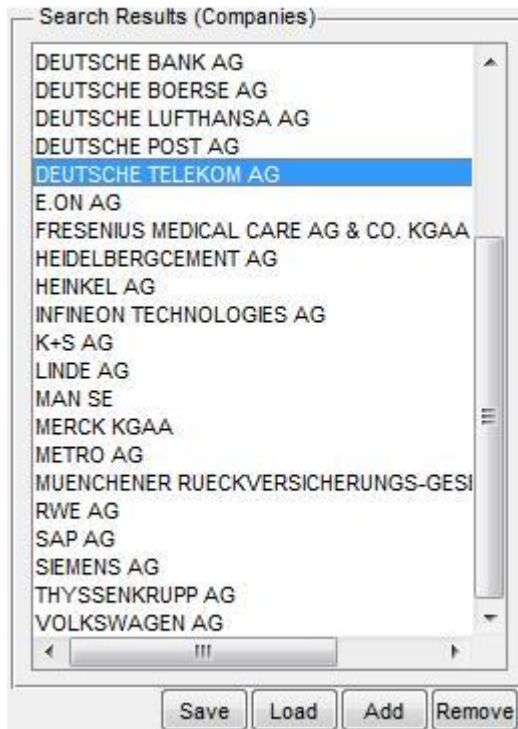
NAT Tool User Manual

In this example, we want to compare the Ravenpack news sentiment index with the stock prices for the DAX30 companies during the years 2007-2010. We are assuming that the Ravenpack news database is available and it includes news for the aforementioned period.

After opening the tool, the first step is to choose the DAX30 companies.



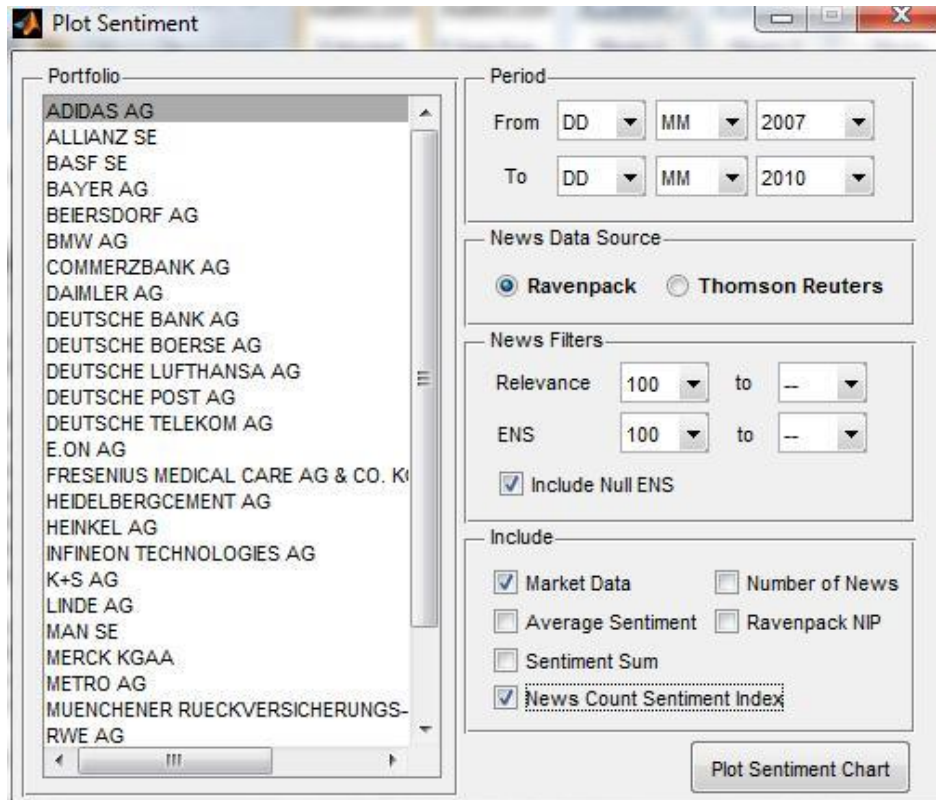
Select the DAX 30 index in the **Companies Filters** panel. Let us include only the active stocks, the surviving DAX 30 companies. Next, click the button **Update Companies List**.



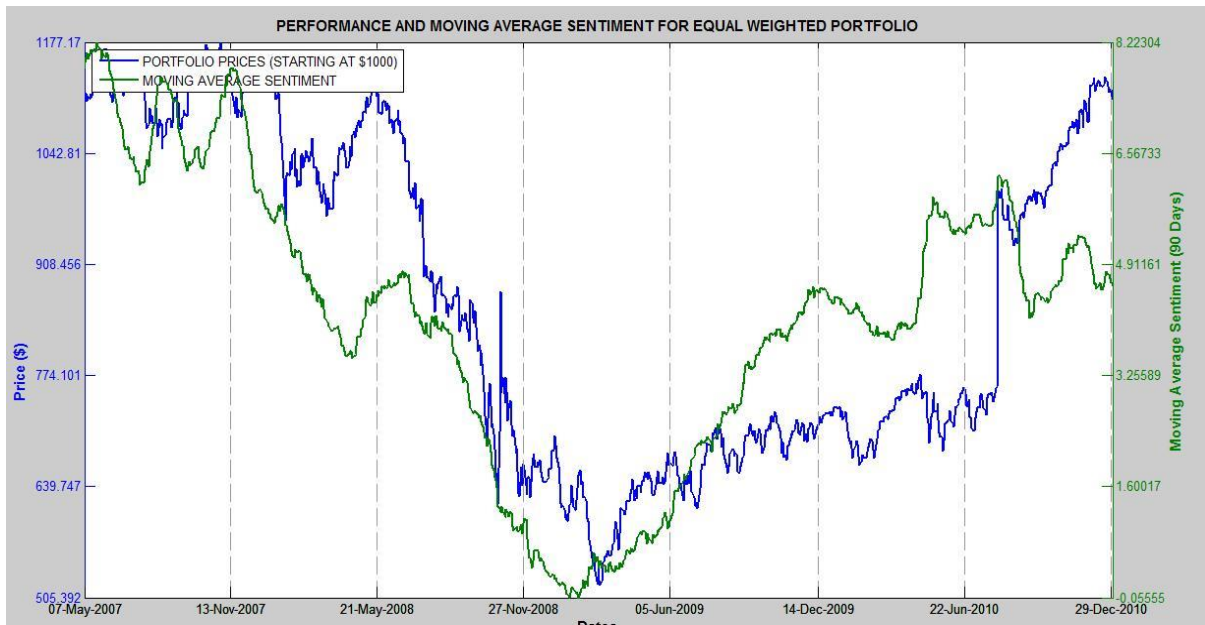
All companies will appear in the **Search Result (Companies)** panel. The companies are not selected yet to the portfolio. Click on any company and then press **Ctrl+A** to select all companies. Next, click the **Add** button and all companies should appear in the **Selected Companies** panel.

After this step, click on the **Plot** menu and then the **Sentiment Chart** menu item.

NAT Tool User Manual



The **Plot Sentiment** window will appear. Choose the year 2007 in the **From** field and 2010 in the **To** field in the **Period** panel. Make sure Ravenpack is selected as **News Data Source**. Let us select only news that received a 100 **Relevance** score, and very novel news according to Ravenpack **ENS** field. Set the values as in the **News Filters** panel shown in the figure above. The **Include Null ENS** checkbox is ineffective in this example, as every news event that has a Relevance value of 100 will have an ENS score. Next, choose the Market Data and News Count Sentiment Index. Finally, click the **Plot Sentiment Chart** button. The following graphic will appear:

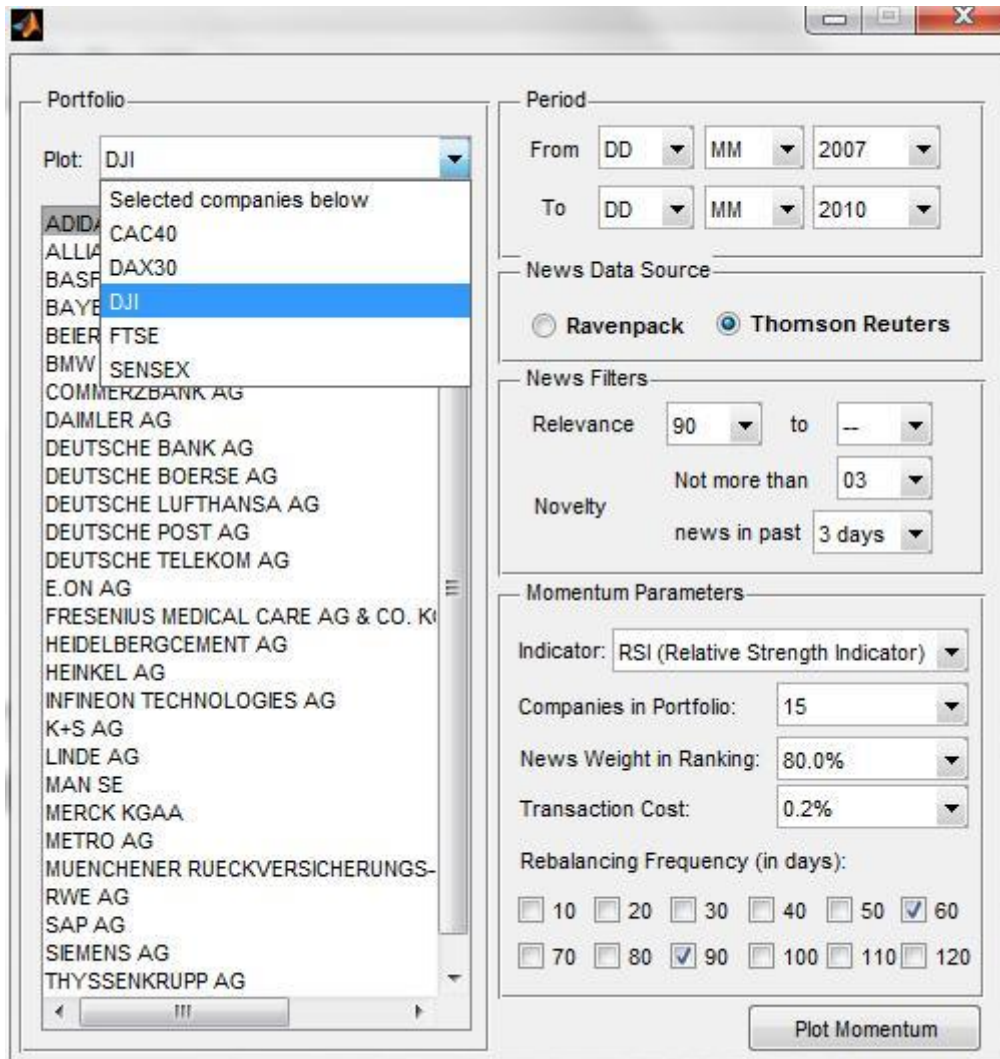


The left y axis and blue line represent the prices of an equally-weighted portfolio containing the stocks from DAX30, while the green line and y axis are the 90-day moving average sentiment index as defined in (P. Hafez; J. Xie, 2011).

10.2. Momentum Strategy, Dow Jones Industrial Average, 2007 to 2010

In this example, we are looking to compare the pure Momentum Strategy with the News Enhanced Momentum Strategy. We will use data from the companies present in the Dow Jones Industrial Average Index, from 2007 to 2010.

It is possible to select any set of companies (portfolio) in the main interface as in the previous section. However, we will do slightly different here. The first step is to open the **Models** menu and click on the **Momentum** menu item.



Let us assume we are starting from where we left in the previous section, the DAX companies were selected. When we open the menu item, we will see all companies in the left side of the panel. Above that there is a list where the first option is **Plot: Selected companies below**. The Momentum interface allows the user to plot the companies of built-in indices instead of the selected companies. This allows us to also plot the index performance as a means of comparing with the Momentum strategies.

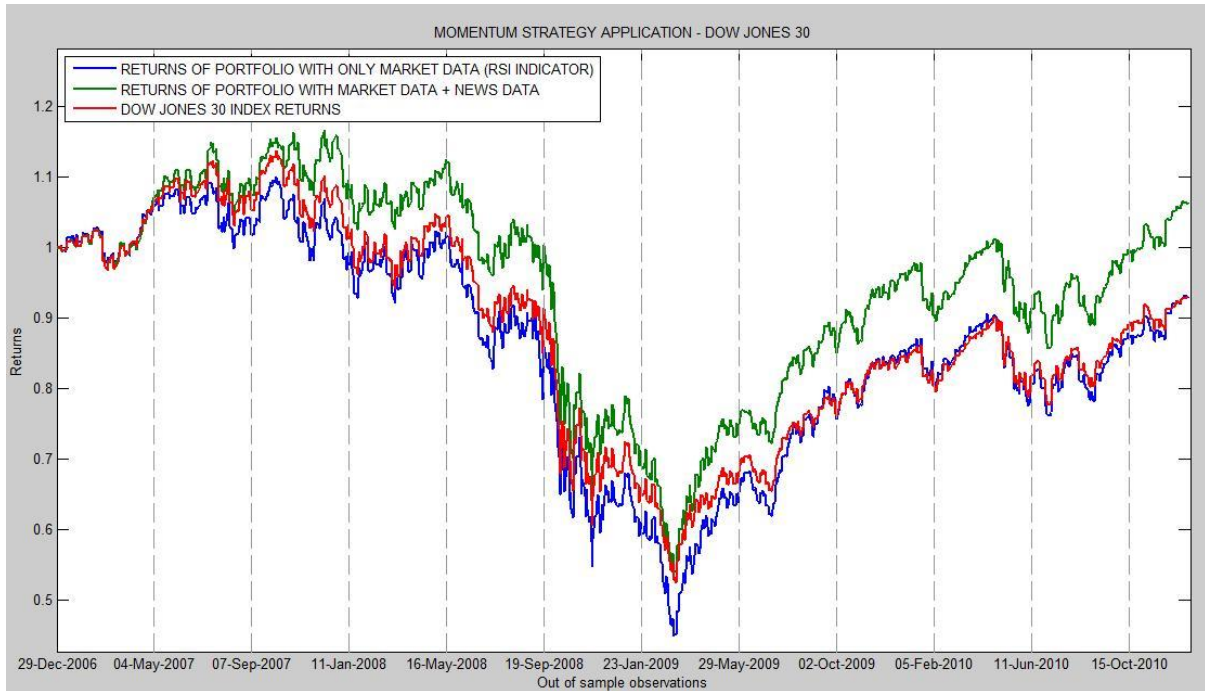
On the **Plot:** list in the left side of the interface, select DJI (Dow Jones Industrial Average). In the **Period** Panel, select from 2007 to 2010. Make sure that **Thomson Reuters** is selected in the **News Data Source** panel. In the **News Filters** panel, let us select now all companies which have a relevance value of at least 90. Novelty in Thomson Reuters is dealt with differently from Ravenpack, so we will choose to include in our calculations only news event that have no more than 3 related news events in the 3 days (72 hours) prior to that one in particular.

The next step is to choose the **Momentum Parameters**. For this example, we will choose the RSI indicator; we will fix the number of companies in our portfolios to 15. In the news enhanced momentum, we create two rankings: the traditional momentum ranking and a news based one. Thus, we define the weight the news ranking will have when combining both; we set this value to 80%. We

NAT Tool User Manual

include transaction costs at 0.2% and we choose the rebalancing frequency to happen every 60-90 days. This means that we are splitting our initial investment in two equal-valued portfolios: The first will be rebalanced every 60 days and the second every 90 days.

Finally, we click on the **Plot Momentum** button. The following graph will appear:



All portfolios are scaled. The red line represents the Dow Jones 30 Index performance. The blue line is the traditional momentum strategy, while the green line plots the performance of the news enhanced momentum strategy.

11. References

Fama, E., & French, K. (1996). Multifactor explanations of asset pricing anomalies. *Journal of Financial Economics*, 51 (1), 55-84.

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for market efficiency. *Journal of Finance*, 48, 65-91.

Leinweber, D. ... (2011). Relating News Analytics to Stock Returns. *Chapter 6 in Mitra, G. and Mitra, L "The Handbook of News Analytics in Finance"*.

Leinweber, D. (2009). *Nerds on Wall Street*. John Wiley & Sons.

Mitra, G., & Mitra, L. (2011). *The Handbook of News Analytics in Finance*. London: John Wiley & Sons.

NAT Tool User Manual

Mitra, L. M. (2009). Equity portfolio risk (volatility) estimation using market information and sentiment. *Quantitative Finance*, 9 (8), 887-895.

P. Hafez; J. Xie. (2011). *Introducing the Ravenpack Sentiment Index*.

Sharpe, W. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19 (3), 425-442.

Uryasev, S. (n.d.). Private Communication.

Wilder, W. (1978). *New Concepts in Technical Trading Systems*. Trend Research.