

OUTPUTS

The Jacobs Levy Markowitz Simulator

Bruce I. Jacobs, Kenneth N. Levy, and Harry M. Markowitz

Output Files

When JLMSim runs in DA mode, it produces seven output files, which the user can inspect and reformat directly. Each output file has a five-part name—one part describing the particular type of output file, one the version of JLMSim that produced the run, two parts describing the particular case, and one part describing the file's format. For example

```
Daily Reports for Case M25c testcase -1.csv
```

consists of the following five parts.

1. The report type—Daily Reports for Case—is one of the seven output files described in more detail below.
2. The JLMSim model (or version) number, M25c, is supplied by the program. (Cases run with different versions of the JLMSim have different numbers.)
3. The case description, "testcase," is specified by the user in the JLMSimInput.txt file. The user is free to use any mnemonic system for identifying cases.
4. The -1 indicates that, for this run, the user selected as the starting random number seed the default setting. The case description specified by the third and fourth parts of each output file name will be the same as the case name the user entered either via the Input GUI or entered directly into the JLMSimInput.txt file.
5. The .csv at the end of the file name indicates that this is a comma-separated variable file, which the user can open with Microsoft Excel or a similar spreadsheet program.

The seven output files consist of the following:

1. The Daily Reports file contains one line for each day of the simulation giving closing prices, daily volumes, and other information characterizing the activity and end-of-day status of the system for that day.
2. The Estimates file shows the expected return and standard deviation estimates of each statistician for each security at the end of each simulated month.
3. The JLMSimInput file is an exact copy of the JLMSimInput.txt file presented to the run, the only difference being that the case description is appended to the name of the output file, as in

```
JLMSimInput for Case M25c testcase -1.txt
```
4. The Messages file reports any problems encountered during the run. One of its final outputs is the first unused random seed, which may be used as an input to the next run if the user wants to rerun a simulation with current inputs but a different starting random seed.
5. The Order Impact file reports various information concerning the 25 largest buy and 25 largest sell attempts for each security by investors of each investor template.
6. The Trace file presents detailed information for a few selected investors each time they reoptimize.
7. TheWhoDidWhat file shows, for each security and simulated day, various information about the investors of the investor templates that bought and sold that security on that day.

The following sections provide more details about each of these files.

The Daily Reports Output File

The Daily Reports file of a run contains one line for each day of the simulation. Each line includes closing prices, daily volumes, and other information characterizing the activity and end-of-day status of the system for that day. JLMSim writes the Daily Reports file as a comma-separated variable file with a .csv extension, thus:

```
Daily Reports for Case M25c testcase -1.csv
```

Successive sets of columns of the Daily Reports file present the

- day number,
- prices of individual securities,
- equal-weighted and capitalization-weighted market indices,
- shares traded by individual securities and total market volume,
- information about the buy and sell books for each security,

- macro variables such as total stock value, total borrowing, and total lending, and
- status and profitability of investors of various investor templates.

Contrary to the counting convention used inside JLMSim, the day numbers in output files are labeled from a base of one; i.e., day number n in the file is the n th simulated day. Let NrS be the number of securities. The next NrS columns of the Daily Reports file present the number of shares of each security traded. The following column presents the sum of these volumes—namely, the total number of shares traded in the market. These numbers tend to be large, so it is convenient to reformat them to include commas.

The next $3 \times NrS$ columns of the Daily Reports file present information concerning the buy books—i.e., the list of buy orders—for each of the NrS securities. The first three columns present information concerning security 0. Subsequent triplets of columns present similar information concerning the remaining securities. The first column of each triplet tells how many orders are on the buy book for that security at the end of that day. The next two columns of each triplet report the highest and lowest bid prices among the buy orders on the book. If there are no orders on the book for a particular security on a particular day, meaningless symbols ("-1.#J" and "1.#J" respectively) appear as the low and high bids.

Following the $3 \times NrS$ columns presenting buy book information are another $3 \times NrS$ columns presenting similar sell book information.

The next ten columns of the Daily Reports file display macro information as of the end of each simulated day. The information in these columns consists of:

- the total value of all stock at the day's closing prices,
- the total value of the cash positions (or, equivalently, lending to brokers) as of the end of the day,
- the total investor borrowing as of day's end,
- total bad debt (the unpaid debt of all investors that became defunct on or before the end of the current day),
- the current value of withdrawals requested but not yet paid out,
- the annual lending rate, with the interest received at the end of the day being the rate shown divided by 240 (the number of trading days per year) multiplied by the amount lent as of the end of the day,
- the annualized rate investors pay the broker on unpaid balances as of the end of the day (computed in the same manner as the lending rate, above),
- accumulated interest received on lending since the start of simulation,
- accumulated interest paid on borrowing since the start of simulation, and
- accumulated value of rebates received on short positions since the start of simulation.

Let $NrIT$ be the number of investor templates in the simulation. The last $9 \times NrIT$ columns of the Daily Reports file summarize the performance of investors in each template through the end of each day. The first six of each of these nine columns contain the following information:

- the total number of investors of that template,
- the number of these that are OK, that is, neither defunct nor currently subject to a margin call,
- the number of active investors, those that are OK and not poor,
- the number of poor investors, those whose total current wealth is less than the small purchase slack input parameter and for whom reoptimizations and trading activity are not performed,
- the number of investors currently subject to a margin call, and
- the number of defunct investors, which are permanently out of the simulation because all their assets have been liquidated following a margin call and they are still in debt.

So-called poor investors may migrate from poor to active if they hold securities that increase in value or if they receive a random cash deposit. The last three of each of the nine columns provides a performance summary. Each of these columns reports a summary statistic concerning the current wealth of the investor divided by the investor's initial wealth. If the investor is defunct or has negative wealth but assets that are not yet completely liquidated, wealth is counted as zero. The three numbers reported are

- the average (arithmetic mean) of current wealth divided by initial wealth among all investors of the given template,
- the standard deviation of the same, and
- the largest value of current over initial wealth among investors of the template.

The Estimates Output File

The Estimates output file provides one line for each simulated month showing each statistician's expected (i.e., mean) return estimate and standard deviation estimate for each security. The first line is a copy of the name of the file.

The first column displays the simulated month. The next NrS columns show the expected return estimates of statistician 0 for every security. Following the expected return estimates for statistician 0 are the standard deviation estimates by the same statistician. Following the estimates for statistician 0 are the expected return and then the standard deviation estimates for each other statistician in turn.

The JLMSimInput Output File

The input parameters for a JLMSim run are read from the `JLMSimInput.txt` stored in the same folder as the JLMSim program. The program produces an output file whose contents are identical to `JLMSimInput.txt` but whose name is changed to include the case description, such as the following:

```
JLMSimInput for Case M25c testcase -1.txt
```

This allows the user to check the exact inputs to a given run at some later date.

The Messages Output File

While running a simulation, JLMSim writes messages to two destinations—the user's screen and the Messages file. The Messages file is a small file containing information regarding the run, such as the case name, information regarding optimizations and error or warning messages from the internal mean-variance optimizer, the last random number seed used, and total values and numbers of deposits and withdrawals.

Messages to the screen include one line per simulated day reporting the month and day within the month, the values of the equal-weighted and cap-weighted market indices at the close of the day, and the number of shares traded during the day.

Order Impact Analysis Output File

The Order Impact Analysis file has a name such as

```
Order Impact Analysis for Case M25c testcase -1 t GE 100.csv
```

It starts with the words "Order Impact Analysis for," followed by the case description. The file name then gives the time at which the analysis is to start; t GE 100 signifies a time greater than or equal to day 100. The start time is an input supplied by the user. Its function is to avoid peculiarities of the simulated system during its initial burn-in period. Entering a value of 0 will start the impact analysis from the beginning of the simulation; entering a value of 1 will start the impact analysis after the first day of simulation, etc. Finally, .csv denotes that this is a comma-separated variable file.

For a given investor template and security number, and for buy and sell orders separately, the Order Impact Analysis file reports information about the 25 largest buy and sell orders for the security from investors of the specified template. The file has 16 columns. The first column indicates the investor template. The second column indicates the sequence numbers of the specific investors of this template that had the largest purchase orders. The file refers to a specific investor using a hyphenated indicator of the form xxx-yyy, where xxx is the template number and yyy is the sequence number. The sixth column indicates the size of the order. The seventh column indicates how many shares were actually purchased. The eighth column shows the time at which the largest order (by investors of this template for this security) occurred. The remaining three columns of the table show, as of the time the order was placed, the highest bid for this security, the lowest offer, and the average of these two numbers. The last five columns of the Order Impact Analysis file show the average completed price, the time of completion, the highest bid at that time, the lowest offer at that time, and the average of the bid and offer.

The Trace Output File

The Trace output file produced by a JLMSim run has a name of the following form:

```
Trace file for Case M25c testcase -1.txt
```

where, again, the second part is a description of the simulation run. This file can be read using any standard text editor. The contents of the Trace output file may be divided into two parts: start-up data, or information read in or computed before the simulation starts, and ongoing data, or information that develops after simulation starts.

The kind of start-up data to be found at the beginning of a Trace file include data concerning the user-specified factor model used to generate returns for the statisticians' databases at the start of the simulation. (This information may also be found in the `JLMSimInput.txt` file.) Other start-up data include each statistician's initial estimate of expected return for each security and the securities' covariance matrix, the expected return estimates used by each portfolio analyst, including those for cash, borrowing and short positions, and information about the efficient frontiers computed by each of the portfolio analysts.

Two kinds of reports follow the start-up information in the Trace file—reports on the status of selected individual investors after each of their reoptimizations and an end-of-day report on the status of all securities.

The body of the post-reoptimization report is a table that, for each security including cash (security number NrS) and debt (security number NrS+1), shows:

- sec, the security's number,
- XUnits, the number of shares held by the investor,
- XValue, the dollar value of these shares,
- XFraction, the ratio of XValue to the investor's current wealth;
- XIdeal, the ideal fraction (the fraction in the portfolio in the efficient set that maximizes $E-KV$ for the investor's K ,
- DFraction, the amount by which the investor will try to increase or decrease Xfraction at this time (an amount less than or equal to the difference between actual and ideal),
- Price, the current price of the security, and
- DUnits, the number of shares the investor will seek to buy (+) or sell (-).

The body of the end-of-day report shows, for each security (not including cash and borrowing):

- Sec Nr., the security's sequence number,
- from the Buy List,
 - size, the number of buy limit orders on the book,
 - low, the lowest bid price of buy orders,
 - high, the highest bid price of buy orders,
- from the Sell List,
 - size, the number of sell limit orders on the book,
 - low, the lowest offer of the sell orders,
 - high, the highest offer of the sell orders,
- Trade Price, the price at which the last trade in this security took place,
- Curr Price, the current price of the security (adjusted at the end of the day if it is outside the bid/asked range),
- Trade Time, the time at which the last trade took place, and
- Tick, whether the last time "tick" changed was a trade at an uptick or downtick.

The WhoDidWhat Output File

For each day and for each security, the WhoDidWhat output file shows which investors bought and sold the security. Within a day, the information is grouped into subtables, each with columns for ten securities

Each subtable has one row for each investor template, showing:

- Inv Temp, the template number,
- Bought, the total number of shares bought by investors in this template,
- Nr Buyers, the number of individual investors of this template who bought this security,
- Seq Largest, the sequence number of the investor of this template who bought the largest number of shares,
- Amt Largest, the number of shares purchased by the investor who bought the most, and
- Acct Status, the account status of the largest buyer (e.g., 0 = OK).

On the sell side, the table shows the amount sold, the number of sellers, and the sequence number, amount sold, and account status of the largest seller.

Output Graphical User Interface

The JLMSim Output Graphical User Interface (or Output GUI) is an interface program designed to assist users in examining some of the most important output of the JLMSim program. In particular, it provides the following functions:

- a quick and easy way to view general results of the simulation, such as total deposits, total number of deposits, total withdrawals, total number of withdrawals, and the next available random seed;
- a simple mechanism for viewing charts of prices, volumes, indices, and equilibrium return estimates;
- the ability to print price charts generated by the JLMSim program; and

- the ability to export Daily Reports files to Excel spreadsheets.

The Output GUI uses the output files generated by the JLMSim program. It reads the program-generated copy of the `JLMSimInput.txt` file to determine parameters such as simulation run length and numbers of entities. It reads the Messages file to determine the total deposits, the total number of deposits, the total withdrawals, the total number of withdrawals, and the next available random seed. In the DA mode, it reads the Daily Reports file to obtain price and volume information for generated price and index series, or in the CME mode, it reads the Trace file to obtain equilibrium return estimates.

Running the Output GUI

After having run JLMSim itself, run the Output GUI from the Start menu by clicking Start/Programs/JLMSim/JLMSimOutput. When the program first runs, it will present you with the screen shown in Figure 1. This screen represents an empty tabbed workbook. Each worksheet in the workbook provides a different level of analysis of a JLMSim case (i.e., a JLMSim run). The following sections describe the process of loading a case and exporting the results of a run to Excel.

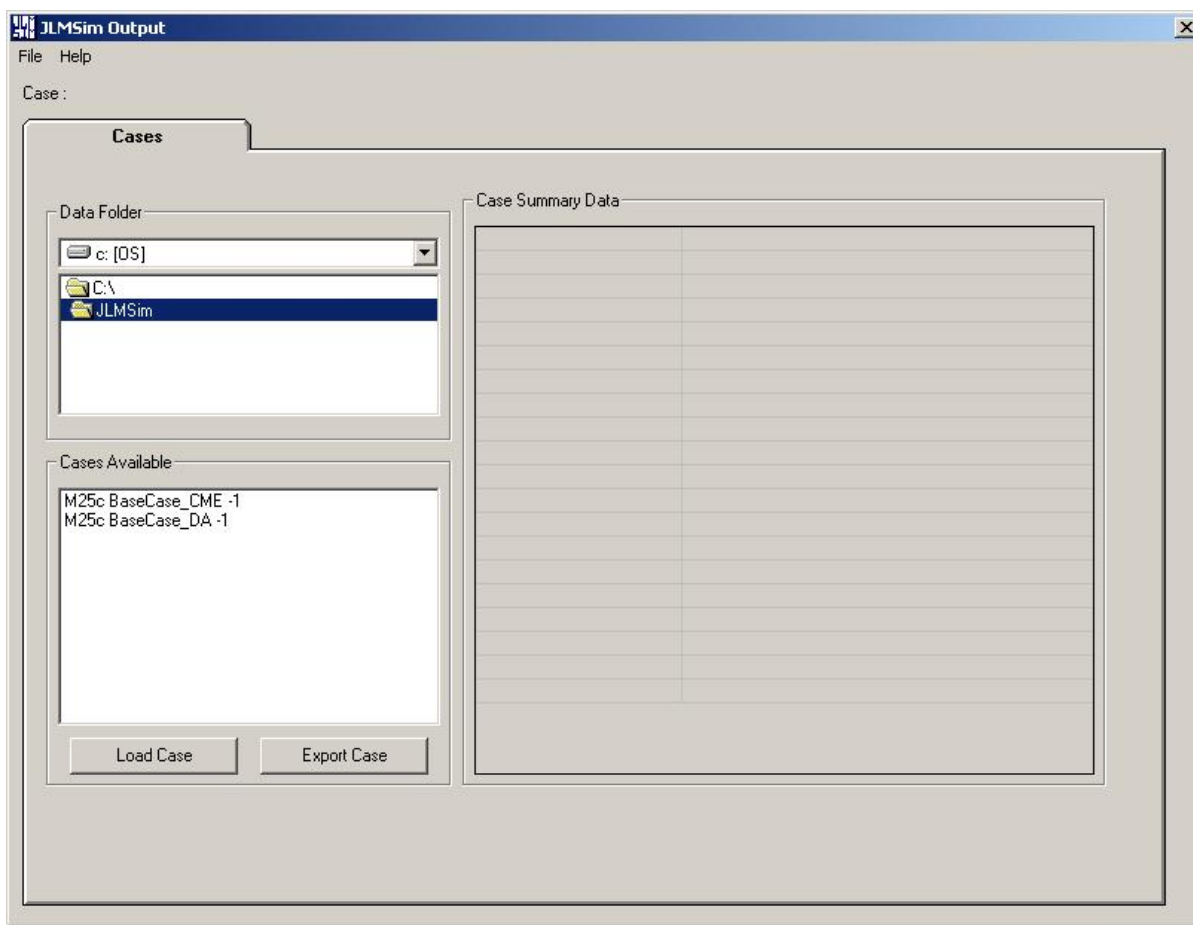


Figure 1: The Empty Workbook.

Loading a Case

The startup screen shown in Figure 1 shows all run results available in the current working directory. By default, the current working directory is set to `C:\JLMSim`. However, if this directory is not available, the Output GUI chooses the working directory to be the one from which it was started. The current directory is displayed in the directory tree in the Data Folder pane of the Cases worksheet. You can use the directory tree to navigate to any desired folder. The Output GUI will automatically display, in the Cases Available pane, the names of any JLMSim output runs available in the chosen directory.

To load a case whose name is displayed in the Cases Available pane, either double-click on the name of the case or highlight the name of the case and click the Load Case button. The Output GUI will then load data pertaining to the case and present a summary similar to that shown in Figure 2.

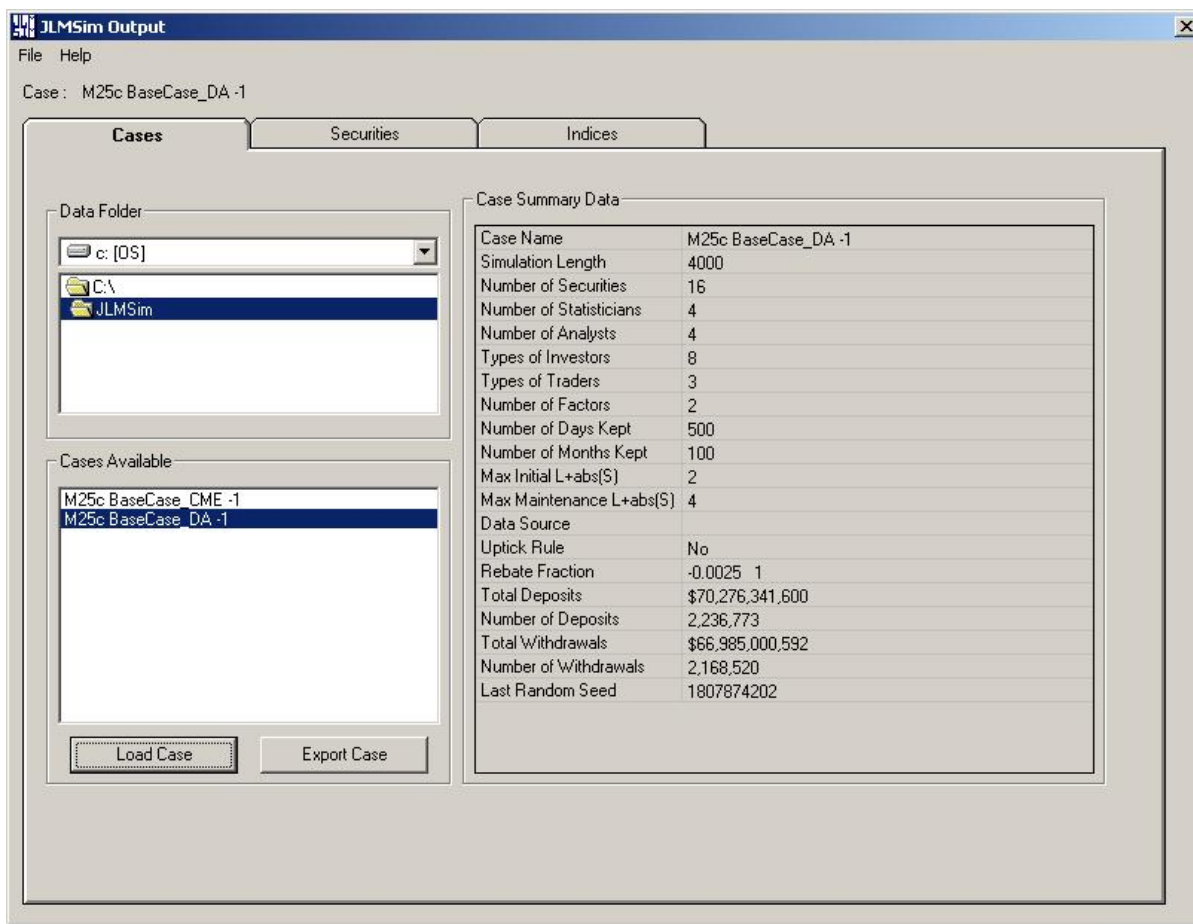


Figure 2: The Basic Information Worksheet in DA Mode.

The Case Summary Data provides an overview of the loaded case. The summary includes basic parameter and entity information gleaned from the carbon copy of the `JLMsimInput.txt` file used to generate the case and basic output data obtained from the Parameters file.

Exporting Daily Reports to Excel

Once you have loaded a case, you can convert the Daily Reports file to Excel by clicking the "Export Case" button or by pulling down the File menu and selecting the Export Case item. The Output GUI will then spend a few moments creating an Excel file, presenting you with a standard file export dialog box that you can use to choose a suitable file name and to select an appropriate directory for receiving the Excel file.

The saved Excel file will be formatted in a convenient fashion. For example, the worksheet will format all numbers to show two decimal places and to have commas separating groups of three digits. The column headers will be written in a bold font, and the worksheet will have panes frozen at cell B5.

Note that you must have Microsoft Excel installed on your computer in order for the export function to work.

Output GUI in DA mode

The Output GUI will automatically detect if the loaded case results from a run of JLMsim in the DA mode, and if so, will present the ability to view security price and volume charts, or index charts.

Securities Price Charts

Once you have a case loaded into the Output GUI, you can view charts of security prices and volumes from that case by clicking on the Securities tab.

The first time that you switch to the Securities tab, you will be presented with a blank worksheet. To plot the price and volume of a security, select one of the securities from the drop-down box at the top left of the worksheet and then click the Plot button. The Output GUI will display price and volume charts as shown in Figure 3.

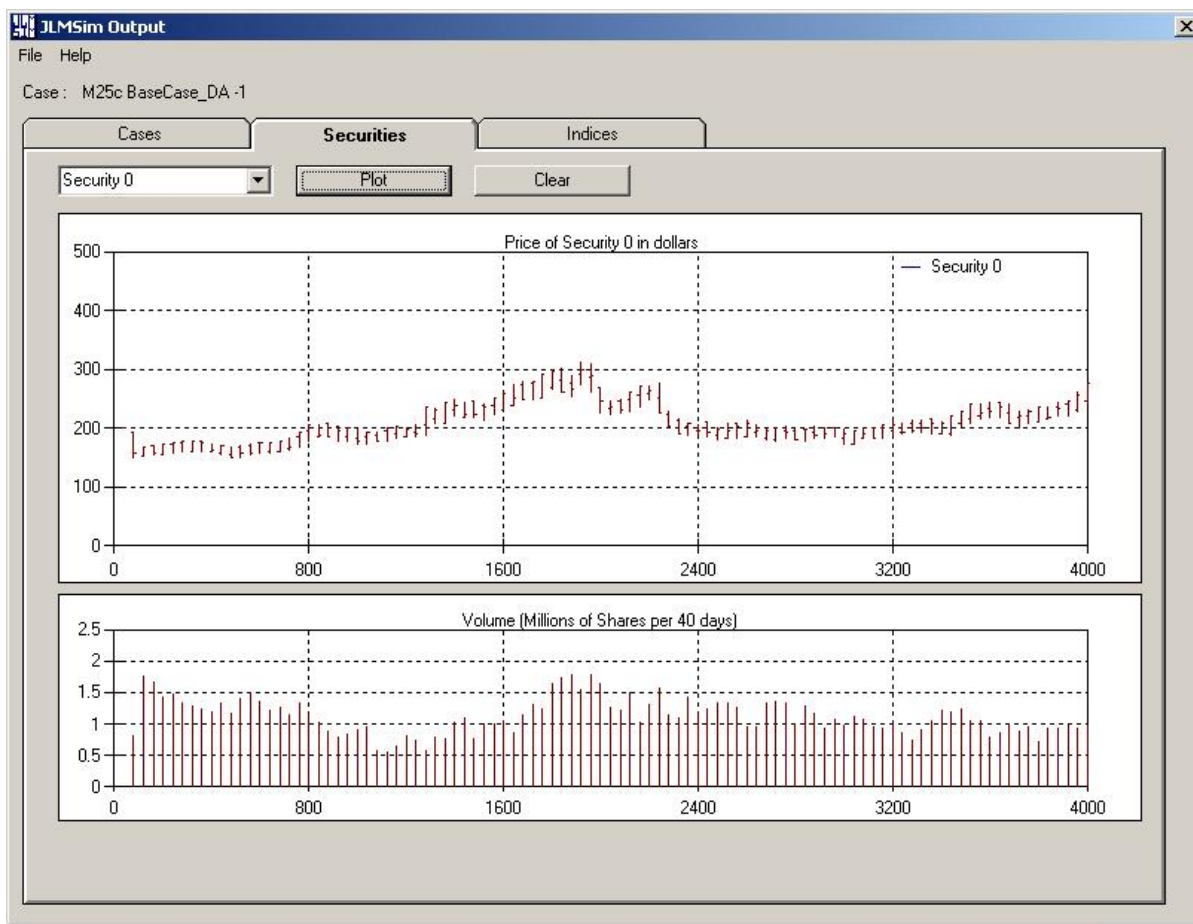


Figure 3: Price and Volume Charts for a Single Security.

The chart has 100 divisions along its time axis, and the axis is scaled so that the results of the entire simulation run are displayed. Thus, for example, for the security shown in Figure 3, the simulation run spanned 4000 days, so each division corresponds to a 40-day period. The price chart shows the opening, high, low and closing prices for each 40-day period, and the volume chart shows the total volume in each forty-day period.

To plot an additional series, select the desired security from the drop-down box and click Plot again. You can repeat this process of adding more plots to view as many securities as you like. When data from more than one security are plotted, the graphs change from open, high, low, close, and volume bars to continuous lines, as shown in Figure 4.

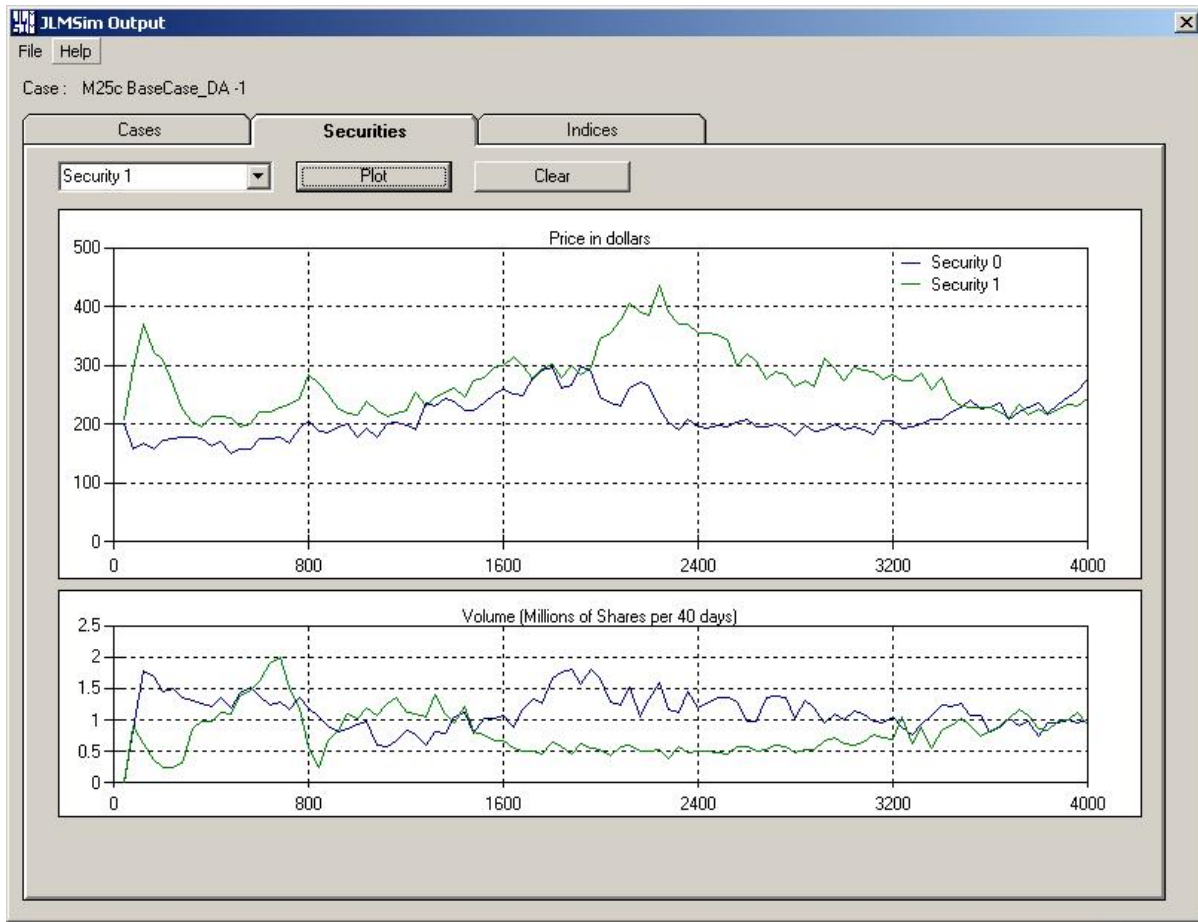


Figure 4: Price and Volume Charts for Multiple Securities.

Clicking the Clear button removes all of the plots from the current worksheet (while keeping the current case in memory), so that charts of other securities from the same case can be plotted.

Printing Hardcopies of Security Charts

To print a hard copy of the chart currently displayed, pull down the File menu, and select Print Chart. The Output GUI will present you with a standard print dialog box, from which you can select an appropriate printer.

Viewing Index Price and Volume Charts

To view the price and volume information for the equal-weighted and capitalization-weighted indices for the currently loaded case, click on the Indices tab. The Output GUI will display a chart similar to that shown in Figure 5.

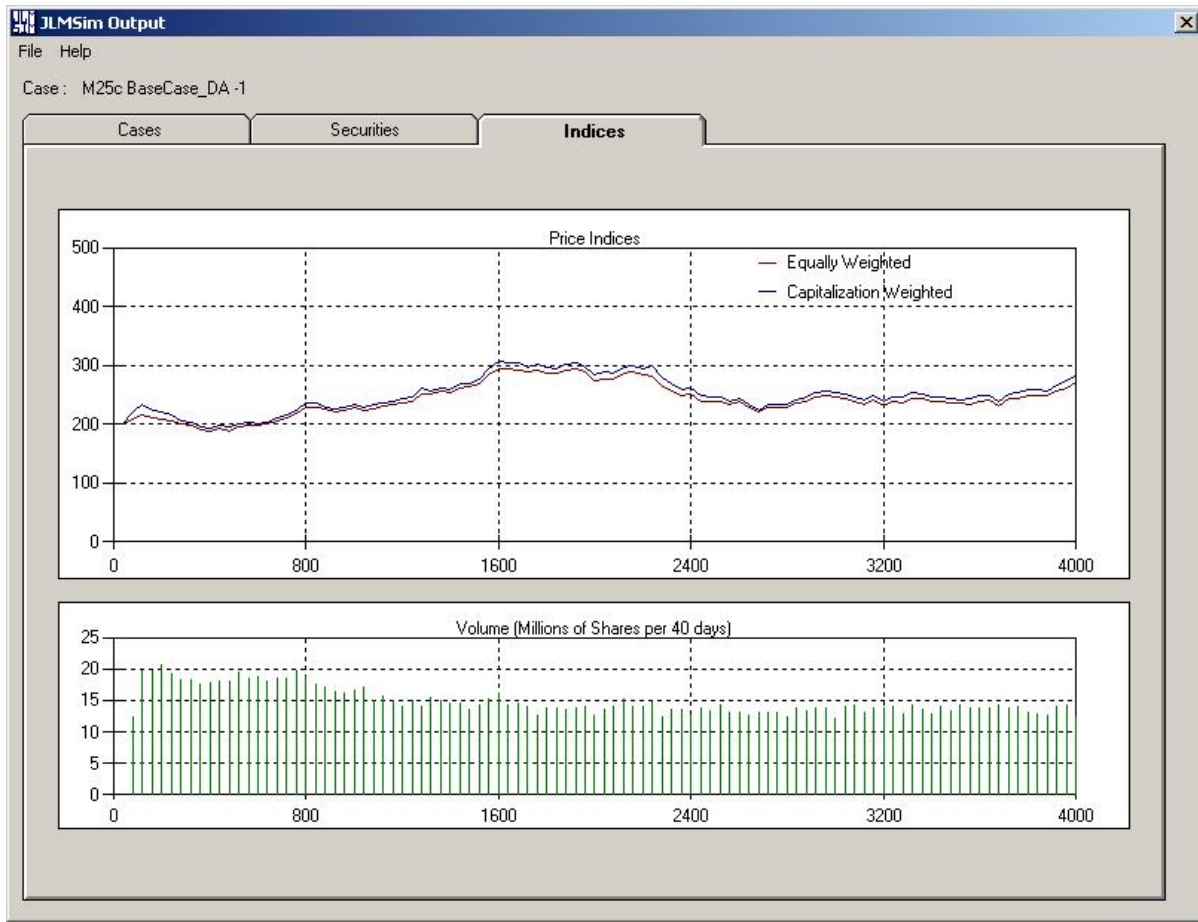


Figure 5: Price and Volume Charts.

Output GUI in CME mode

The Output GUI will automatically detect if the loaded case results from a run of JLMSim in the CME mode, and if so, will present the ability to view security price estimate convergence charts and equilibrium return estimate convergence charts. Figure 6 shows the appearance of the Basic Information worksheet after a typical CME case is loaded.

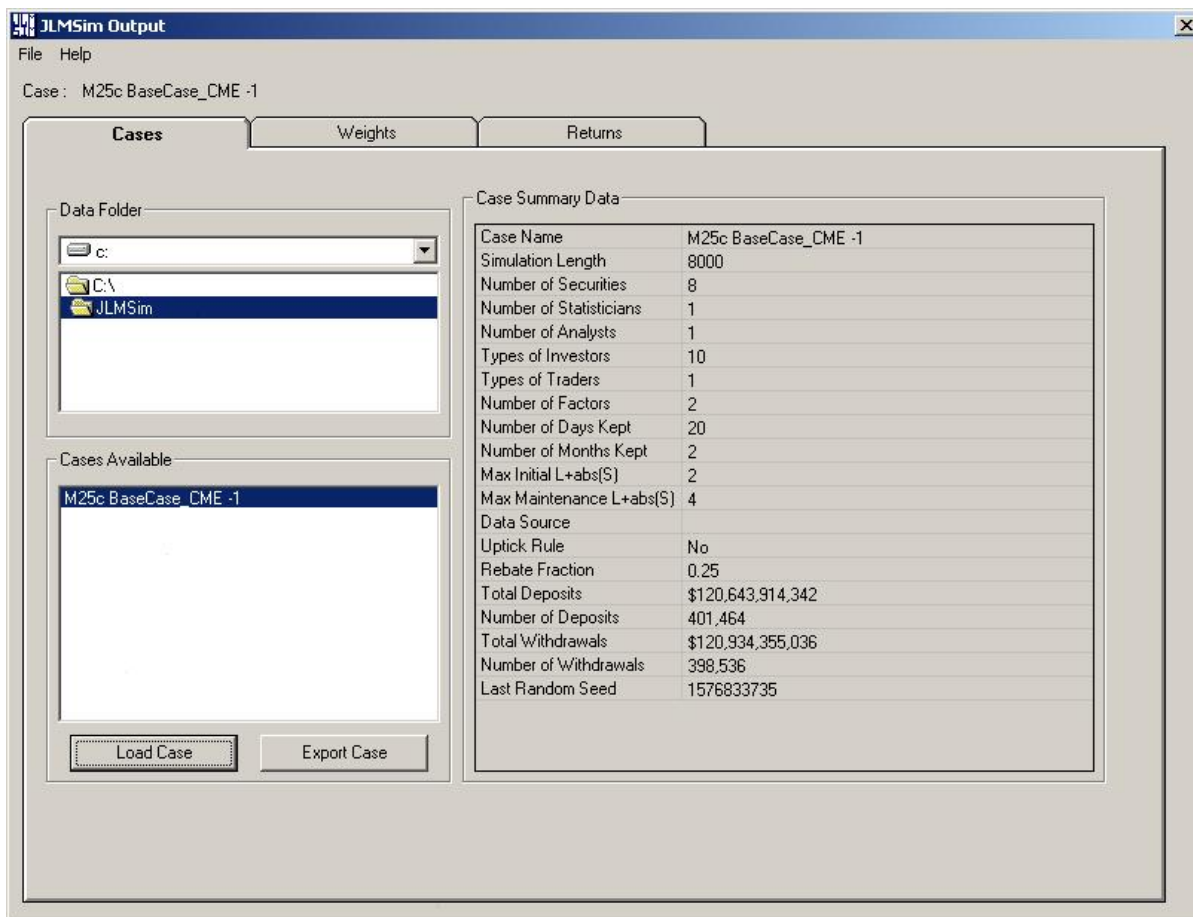


Figure 6: Basic Information Worksheet in CME Mode.

When a CME case is loaded, new tabs titled Weights and Returns appear. The Weights tab contains plots of the convergence of the simulated market portfolio toward its target weights. This is illustrated in Figure 7. The user has control over both the total number of iterations to plot and the sampling interval.

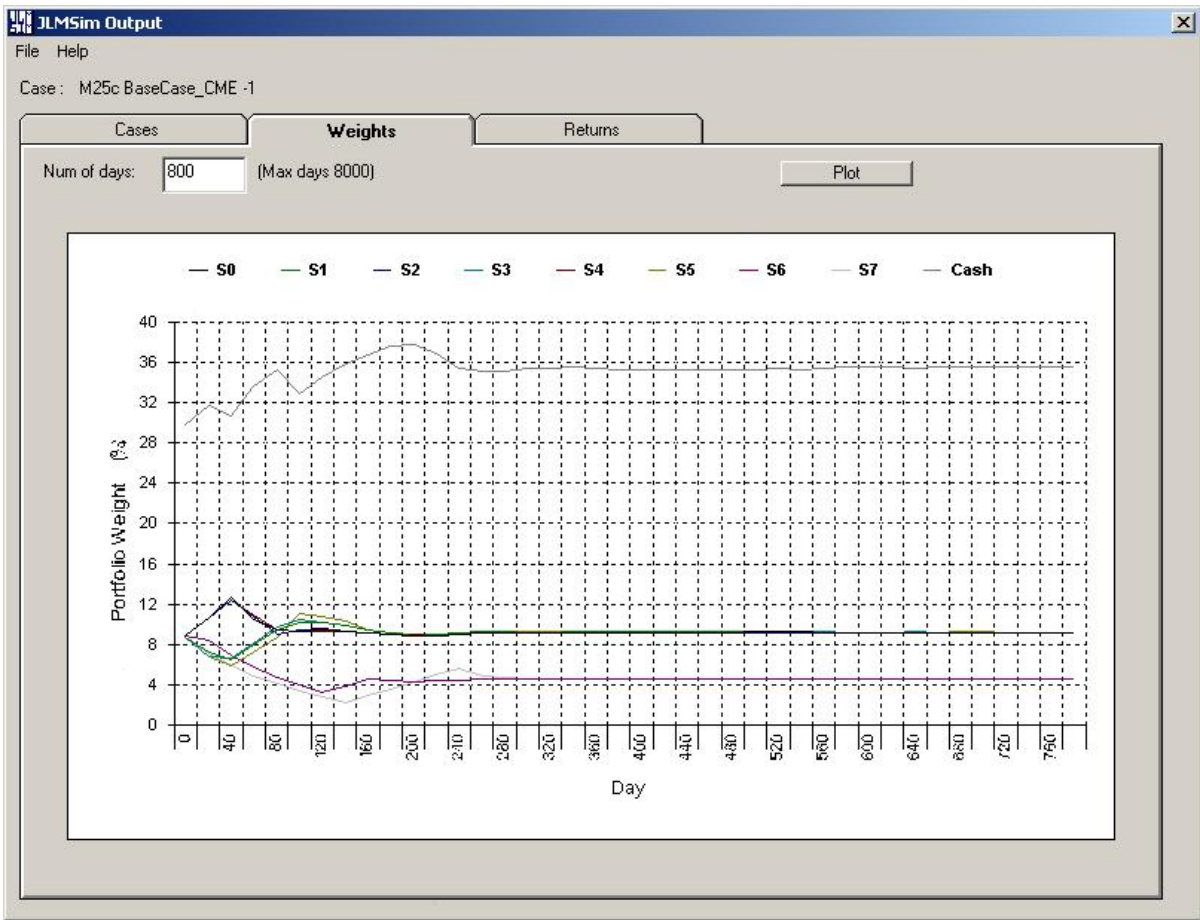


Figure 7: Iterative Convergence of Portfolio Weights toward Target Values in CME Mode.

The Returns tab contains plots of the convergence of the estimated capital market equilibrium returns. This is illustrated in Figure 8. The user has control over the total number of months (i.e., groups of 20 iterations to plot). This is illustrated in Figure 8.




Figure 8: Iterative Convergence of Estimated Capital Market Equilibrium Returns in CME Mode.