nQuery Advanced

MANUAL

Statsols

nQuery Advisor® Sample Size Calculator

v. 8.5.1.0 User Manual

Statsols

Statistical Solutions Ltd.,
One International Place, 100 Oliver Street, Suite 1400,
Boston, MA 02110

Statistical Solutions Ltd.
Building 4500, Airport Business Park,
Cork, Ireland,

Registered in Ireland Reg No. 233638

US: +1 (617) 830 1836 EU/ROW: +353 21 48 39 100
Web: http://www.statsols.com Email: sales@statsols.com

Acknowledgements

We would like to sincerely thank all those who made the production of nQuery Advisor from Statsols possible. The Statsols R&D Team: David Lemoine, Steven Keady, Cathal O'Gorman, Christine Carey, Ronan Fitzpatrick, Brian Fox and Paul Murphy and the Statsols Team: Brendan Nyhan, David Cleary, Sorcha Conaghan, Liam Burke, Laura O'Donovan, Diana Scriven, Caroline Costello, Helen Murphy, Kevin Sievewright and Mary Byrne. Special Thanks to: Dan Andrei Chirac, Andrei Banc, Claudiu Balanean and Cristina Pop from Lasting Software, Eimear Keyes, Adrian O'Leary, Aodhan O'Leary, Luke Rothwell, Kieran Logan, Tricia Baalfe, Dr. Janet Elashoff, Andrew Grannell, Brian Sullivan, Eoghan Murphy, Aisling Leen, Paul Kenneally, Louise Byrne, Layla Tobin, Garry McCarthy and Dr. Niall Fitzgerald.

IMPORTANT: READ BEFORE INSTALLING SOFTWARE

nQuery Advisor License Agreement

This document sets forth the terms and conditions of the license and the limited warranty for nQuery Advisor, hereinafter referred to as "nQuery". Proceeding with this installation constitutes your acceptance of this license agreement with respect to all accompanying nQuery software received by you. If you do not accept this agreement, you may cease the installation within 10 days of receipt of installation key and your money will be refunded.

THIS LICENSE AGREEMENT is made and entered into by and between Statistical Solutions Limited, an Irish registered company, hereinafter referred to as "LICENSOR", and you, hereinafter referred to as "LICENSEE".

WHEREAS, LICENSOR has ownership rights of nQuery ("SOFTWARE") and LICENSEE desires to obtain from LICENSOR, and LICENSOR desires to grant to LICENSEE, a license to use the aforementioned nQuery.

NOW, THEREFORE, in consideration of the mutual covenants, conditions and terms hereinafter set forth, and for other good and valuable consideration, the parties hereto hereby agree as follows:

I. GRANT

A. Subject to the provisions contained herein, LICENSOR grants to LICENSEE a non-exclusive, non-transferable license for the use of SOFTWARE ("LICENSE") which shall not and will not be copied by the LICENSEE and will not be passed by the LICENSEE to any other person or party, or will the whole or any part of it be disclosed by the LICENSEE to any other or party.

B. Unless otherwise specifically agreed in writing by the LICENSOR, SOFTWARE may be used solely by LICENSEE.

C. No title or ownership rights to SOFTWARE are transferred to LICENSEE by this Agreement.

II. MAINTENANCE

A. LICENSOR agrees to provide maintenance for SOFTWARE for the duration of this Agreement.

B. MAINTENANCE includes the following services and materials.

1. Updates and improvements to SOFTWARE, as these become available.
2. Reasonable written or telephone consultations.

C. Maintenance will be provided only for the release of SOFTWARE which is current at the time maintenance is required.

III. REGISTERED USE

A. LICENSEE agrees that it will maintain SOFTWARE (including any future forms of SOFTWARE provided to LICENSEE) in confidence.

B. LICENSEE agrees not to knowingly provide, distribute or otherwise make available, in whole or in part to those outside of LICENSEE'S site, SOFTWARE, or any derivative thereof, including without limitation program listings and code.

IV. ACCEPTANCE AND PAYMENTS

A. All terms and conditions of this Agreement shall be binding upon purchase of LICENSE by LICENSEE.

B. LICENSEE, for and in consideration and as a condition of purchasing this LICENSE, hereby agrees to pay over to LICENSOR an annual fee. The yearly fee in effect will become due unless LICENSEE has properly exercised its right to terminate this Agreement pursuant to paragraph V.B. below.

V. TERM

A. This Agreement shall become effective upon purchase of LICENSE by LICENSEE and shall remain in effect for an initial period of one (1) year, and thereafter shall be automatically renewed for successive periods of one (1) year, unless terminated as provided in this Agreement.

B. LICENSEE or LICENSOR shall have the right to terminate this Agreement, at the end of any licensing period upon written notice, which shall be given to the other party at least sixty (60) days prior to the expiration of the yearly licensing period. In the event that notice of termination is not given by either party at least 60 days prior to the expiration of the license period the Agreement shall be deemed effective for another one year.

C. Either party shall have the right to terminate this Agreement at any time upon violation by the other party of any of the provisions of this Agreement.

D. Upon termination, of this Agreement, LICENSEE shall destroy all installations of SOFTWARE, LICENSEE shall promptly forward written notice to LICENSOR that all programs and materials containing SOFTWARE have been destroyed and are no longer in use or usable by LICENSEE.

VI. WARRANTIES

A. LICENSOR warrants that it has the right to grant this License.

B. LIMITED WARRANTY AND REMEDY. Both parties hereto agreeing that this is not a consumer contract, LICENSOR warrants that for sixty (60) days after purchase of LICENSE by the LICENSEE, SOFTWARE shall reasonably perform in accordance with the accompanying documentation (manuals and computer aided instructions). LICENSOR specifically DOES NOT warrant that SOFTWARE will operate uninterrupted or error free.

LICENSEE'S SOLE AND EXCLUSIVE REMEDY for a failure of SOFTWARE to perform in accordance with the accompanying documentation is for the LICENSEE to return SOFTWARE to LICENSOR in writing of the non-performance within sixty (60) days of purchase of the media. LICENSOR'S sole obligation within a reasonable time after receiving notification of non-performance shall be, at LICENSOR'S option to:

1. correct demonstrable and significant program errors.
2. provide the LICENSEE with functionally equivalent software, or
3. provide or authorise a refund.

THE LICENSOR'S LIMITED WARRANTY SET FORTH ABOVE IS IN LIEU OF ANY OTHER WARRANTIES. LICENSOR DISCLAIMS ANY IMPLIED WARRANTY INCLUDING, BUT NOT LIMITED TO ANY IMPLIED WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE OR MERCHANTABILITY. LICENSOR DISCLAIMS ALL OTHER WARRANTIES INCLUDING, BUT NOT LIMITED TO, THOSE MADE BY DISTRIBUTORS OF LICENSOR'S SOFTWARE, BUT DOES ATTEMPT TO DISCLAIM OR EXCLUDE LIABILITY FOR PERSONAL INJURIES WHERE THE LICENSOR HAS AN OBLIGATION FOR SAME UNDER LAW.

A. In no event shall LICENSOR become liable to LICENSEE, or any other party, for any loss or damages, consequential or otherwise, including but not limited to time, money or goodwill arising from the use, operation or modification of SOFTWARE by LICENSEE. LICENSEE further agrees to indemnify, defend and hold harmless, its agents or employees from and LICENSOR against all loss or expense resulting from liability of any nature arising or resulting from the acts or omissions of LICENSEE, its employees or agents, in its use, possession or operation of SOFTWARE.

B. If any of the provisions, or portion thereof, of this Agreement is invalid under any applicable statute of law, they are to that extent to be deemed omitted. This Agreement shall constitute the entire agreement between the parties hereto with respect to the subject matter hereof, and shall supersede all oral, promotional or advertising presentations. This Agreement may not be modified except in writing signed by an authorised representative of LICENSOR. The prevailing party in the legal action brought

to construe the terms of or enforce this Agreement shall be entitled to reasonable legal fees and costs. Under no circumstances shall this Agreement or any of the rights granted to LICENSEE hereunder be sold, assigned or sub-licensed, voluntarily or by operation of law, to any other person or entity, and any such purported sale, assigned or sublicenses shall be void.

VII. This agreement shall be governed by the laws of Ireland and the parties agree to the exclusive jurisdiction of the Irish Court.

VIII. DATA PROTECTION

Each of the parties shall observe all of their respective obligations under all legislation and regulations relating to the protection of personal data (including Regulation (EU) 2016/679) that arise in connection with this Agreement. The parties acknowledge that LICENSEE will not be disclosing or transferring personal data to LICENSOR and that LICENSOR will not be processing personal data on behalf of LICENSEE.

Contents

1 nQuery Basics

This chapter is a quick guide for users to familiarize themselves with how to access nQuery Advisor (referred to as nQuery from hereon) and conduct a basic sample size analysis. Many individual elements of using nQuery are described in more detail in subsequent chapters.

Note that installation is not covered in this tutorial as installation guides will have been provided to you with the application and are also available online in your nQuery account.

1.1 System Requirements

Minimum system requirements for nQuery 8.0.0.0

- Processor: 2 core 1 GHz processor
- Memory: 3 GB of available hard disk space
- RAM: 2GB
- OS: Windows 7 Service Pack 1 (Home Premium) or higher
- Other: .NET Framework 4.6.2

1.2 Help and Support

The easiest way to get statistical or technical support is to contact Statsols using their nQuery Community User account (see subsection 1.4.1) or to use the Help Centre button found in the bottom left of the nQuery application (see subsection 1.4.5). The support centre can also be found at the following URL:

- info.statsols.com/help-center

1.3 Starting nQuery

There are two main ways to open nQuery. Firstly, nQuery can be opened by double-clicking on the desktop icon (if selected during installation) and then nQuery will be

automatically launched. Secondly, it can be found in the Windows Start menu by clicking on the Windows icon in the bottom-left. In each OS, the easiest way to find the application is to enter "nQuery" into the search dialog and select the "nQuery" application when it appears in the results. Alternatively, to find the application manually, use the following guide for your respective OS:

Windows 10: Click the Start menu in the bottom-left. In the "All Apps" menu on the left-hand side of the Start menu, go to the "N" section in the alphabetised list, select the "nQuery" folder and select the "nQuery" application.

Windows 8/8.1: Click the Start menu in the bottom-left. Select the "All Apps" menu using the arrow at the bottom of the screen. Scroll to the nQuery application or enter "nQuery" in the search window in the top-right and select the "nQuery" application.

Windows 7: Click the Start menu in the bottom-left. Then select "All Programs" and you can locate nQuery in the alphabetised list. Click on this folder and select the "nQuery" application.

1.4 nQuery Start-up and Layout

Once the user has launched nQuery, the application will appear as illustrated in Figure 1.1. The five major elements in the application on start-up are the Home tab (1), the toolbar (2), the menu bar (3), the tab menu (4) and the information bar (5). These are summarised in the following section.

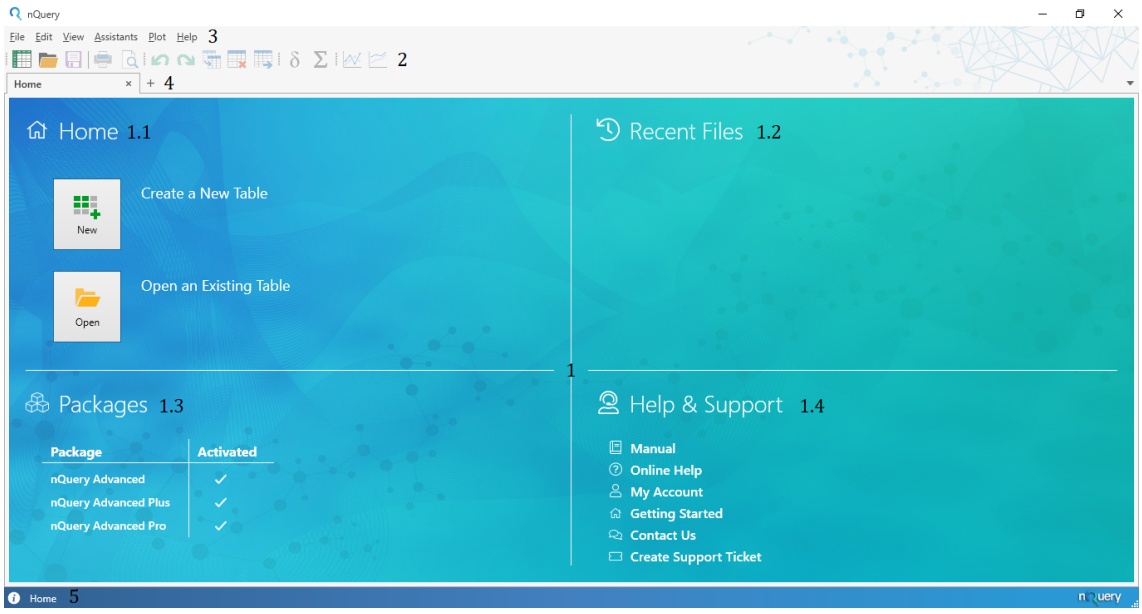

Figure 1.1: nQuery on Start-up

1.4.1 Home Tab

When you open nQuery, the software will open the "Home" tab. This window contains options to open new tables or pre-existing nQuery files, open recent nQuery files, see which nQuery modules are currently active and support options if you need further help in using the software. The four quadrants (1.1 to 1.4 from top-left to bottom-right) in the Home tab are described in more detail below.

1.4.1.1 Home

The Home quadrant (top-left) contains two buttons: "Create a New Table" and "Open an Existing Table".

Selecting the "Create a New Table" button will open the "Select Test" window where you can open a new table. This window is described in section 1.5

Selecting the "Open an Existing Table" button will open the "Open" window where you can open any pre-existing nQuery (.nqt) file. This is standard Windows "Open File" dialog.

1.4.1.2 Recent Files

The Recent File quadrant (top-right) will contain a list of all save files recently opened in nQuery. To open a file, select the required file from this list and the save file will open automatically.

1.4.1.3 Packages

The Packages quadrant (bottom-left) will contain a list of the nQuery packages which are currently active under your current nQuery license. Packages which are active will have a tick to the right of the package name. For details on nQuery packages and activating them, see section 7.2. To purchase additional packages, see `www.statsols.com`.

1.4.1.4 Help & Support

The Help & Support quadrant (bottom-right) contains a range of help, support and troubleshooting options for nQuery. These options are as follows:

1. Manual: This will open the user manual (this document)
2. Online Help: This will open the nQuery Knowledge Base website. This website provides a comprehensive set of instructions for installation, using the software and common troubleshooting issues and solutions.

3. My Account: This will open the nQuery Account Administrator Portal. Your nQuery account website will contain access to the relevant installers, activation keys and installer guides.
4. Getting Started: This will open the "How to Use nQuery" website. This website contains videos and instructions on all the basic elements of using nQuery for sample size determination.
5. Contact Us: This will open the "nQuery FAQ" website. This website will contain answers to frequently asked questions (FAQs) about nQuery and provides an "Ask a Question" facility where you can send a custom query to us online by filling out the relevant form.
6. Create a Support Ticket: This will bring you the "Log a Support Ticket" website. Select the "Log a Ticket" option to start this process. Note an nQuery account will be required to use this facility.

1.4.2 Toolbar

The toolbar provides easy access to the most commonly used functions in nQuery. The toolbar options can be split into four main categories: file options, print options, table options, side-table options and plotting options. The toolbar is shown in detail in Figure 1.2.

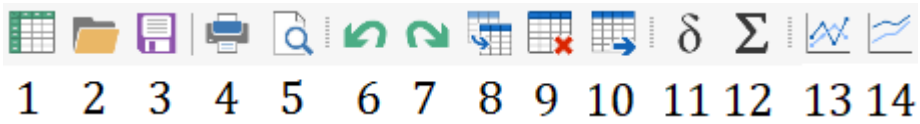

Figure 1.2: Toolbar

There are fourteen options in the menu bar. These are as follows:

1. New Table: Open the table select screen
2. Open Table: Open nQuery table save files
3. Save Table: Save the currently open nQuery table
4. Print: Print the currently open nQuery table
5. Print Preview: Open Print Preview of currently open nQuery table
6. Undo Action: Undo last action within currently open nQuery table
7. Redo Action: Redo last undone action with currently open nQuery table
8. Copy Table: Copy all contents of currently open nQuery table
9. Clear Table: Clear all entries from currently open nQuery table

10. Fill Right: Fill all entries to the right with the same selected cell value
11. Compute Effect Size: Open the Effect Size side-table (if available)
12. Compute Covariance Matrix: Open the Covariance Matrix side-table (if available)
13. Plot Power vs Sample Size: Plot Power vs Sample plot for selected column(s)
14. Plot User Selected Rows: Open custom plot dialog for selected column(s)

These options will be explored in further detail later in this manual.

1.4.3 Menu Bar

There are six options on the menu bar: File, Edit, View, Assistants, Plot and Help. These are highlighted in Figure 1.3.

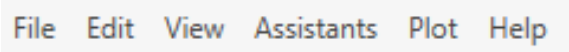

Figure 1.3: Menu Bar

The **File** menu allows the user to open a new or previously saved design table, as well as enabling the user to save a design and allowing the user to exit nQuery whenever they wish. Design tables can be saved as .nqt format, which is the Statsols file format for nQuery. It also includes the options for printing, the option menu and to close nQuery.

The **Edit** menu enables the user to undo and redo actions within the current nQuery table (shortcuts are provided in text). The user has options to copy and clear a table or fill a table using the Fill Right option. The fill right option is where the user, when defining multiple columns, enters certain information into a column and can copy this information across the remaining empty columns.

The **View** menu options are initially unavailable (greyed out) until the user opens a design table except for the "Home" option. The "Home" option will open or close the initial "Home" tab in nQuery. Once a table has been opened, options appear enabling the user to open or close the table elements of the output statement, the specify multiple factors tool, the help panel and the notes tool. These table elements are described in section 1.8.

The **Assistants** menu gives access to several useful tools that can help with a user's calculation. These include access to the "Calculate Effect Size" and "Specify Covariance Matrix" side-tables, access to utilities for deriving the standard deviation, access to cumulative distribution functions for multiple statistical distributions, a data entry tool, a survival parameter converter tool, a Bayesian Posterior Error calculator, the Report feature and a shortcut to the windows calculator.

The **Plot** menu is initially unavailable until the user opens a design table. Once a table has been opened and filled appropriately, the user can use this menu to create Power vs. Sample Size and Plot User Selected rows plots. There are also table specific options for Boundary and Inverse Boundary plots (group sequential design tables e.g. Two Means GST - MTT12) and the Survival vs Time plot (simulation survival tables - STT3/STT3u).

The **Help** menu gives access to the nQuery user manual (this document), the "About" page containing the nQuery version information and license agreement, the Installation (IQ) and Operational Qualification (OQ) tools and utilities for activating/renewing a license, enabling add-on modules and for checking for the latest nQuery updates.

Below is a complete list of menu options:

File: New, Open, Save, Save As, Print, Print Preview, Recent (horizontal menu), Options, Close

Edit: Undo, Redo, Fill Right, Copy Table, Clear Table

View: Home, Specify Multiple Factors, Output, Help, Notes

Assistants: Compute Effect Size, Compute Covariance Matrix, Standard Deviation, Data Entry, Distribution Functions, Survival Parameter Converter, Posterior Error Rate Calculator, Report, Windows Calculator

Plot: Power vs. Sample Size, User-Selected Rows, Survival vs Time plot (STT3/STT3u only), Multiple Boundary (Group Sequential Design only)

Help: nQuery Manual, Activate/Renew License, Enable Modules, Check for Updates, Installation Qualification, Operational Qualification, About

1.4.4 Tab Menu

The tab menu provides an easy way to navigate between different tables within nQuery and quickly create fresh copies of open tables. It is shown in Figure 1.4 (note that we have opened some additional tables in nQuery for illustration purposes)

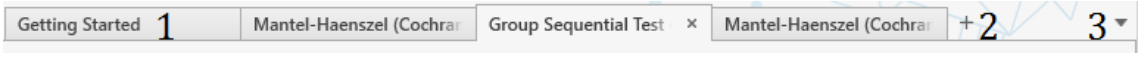

Figure 1.4: Tab Menu

There are three main elements to the tab menu: the tabs, the fresh table shortcut and the table navigation drop-down.

1. Tabs: Individual tabs for each table opened in the session and the "Getting Started" tab. The tab also has an "x" button on its right-hand side to close an individual tab

2. Fresh Tab: If the Getting Started tab is selected or no tabs are available, will open the "Select Test Design and Goal" menu. If a table is selected, it will open a new clean version of that table.
3. Tab Navigation Menu: If selected, will open a drop-down menu of all the open tabs. When a tab is selected, this table is displayed in nQuery.

1.4.5 Information Bar

The information bar provides information on the currently opened table, the currently selected cell and access to the nQuery help centre. It is shown in Figure 1.5 (note that we have selected a cell in a table for illustration purposes)

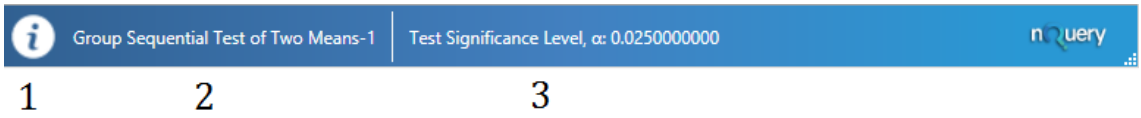

Figure 1.5: Information Bar

There are three main elements to the information bar: the help centre, the table name and cell information

1. Help Centre: Selecting the information symbol will open the online help centre in a browser. It provides links to guides, troubleshooting and create support ticket
2. Table Name: Provides the name of the currently open table
3. Cell Information: Provides the name and the exact value (to 10 decimal places) in the currently selected table cell

1.5 Opening a Design Table

The next aspect of the interface that will be reviewed is the opening of a new design table. There are three ways in which the user can open a new design table in nQuery:

1. Selecting New from the File menu
2. Selecting the New option from the menu bar 
3. Selecting the "+" symbol in the tab menu

For option 3, if the Getting Started tab or no tabs are available this will open the "Select Test Design and Goal" window and if a table tab is selected this will open a new fresh copy of that table.

When these options are selected, the "Select Test Design and Goal" window will open. This is shown in Figure 1.6.

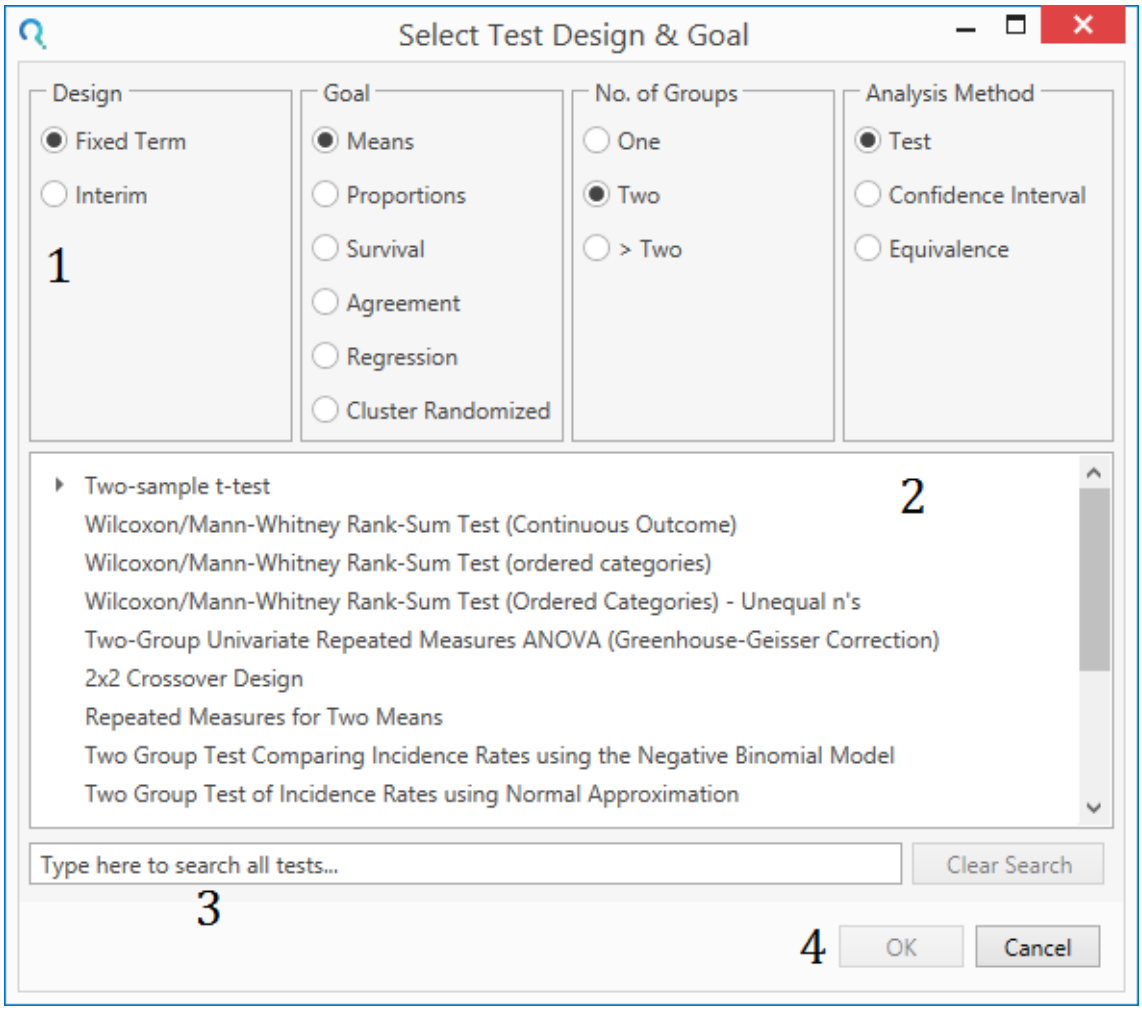

Figure 1.6: Select Test Design and Goal Window

The Select Test Design and Goal window consists of four main elements: the table selection menus (1), the table selection window (2), the table search bar (3) and the OK/Cancel buttons (4).

In nQuery, there are two main ways to find a specific design table: using the table selection menus or using the table search bar.

Table Selection Menus: The table selection menu consists of four radio button columns. These are for the Design, the Goal, the No. of Groups and the Analysis Method. A brief description of these follows:

1. Design: Specify whether the study analysis will use a Fixed Term (all data analysed at study end) or Interim (analysis of data mid-study) design.

2. Goal: Specify the type of data which will be analysed in the study. Choose between Means, Proportions, Survival, Agreement, Regression or Cluster Randomized.
3. No. of Groups: Specify the number of groups which will be compared in the study. Choose between One (against standard), Two or Greater than Two.
4. Analysis Type: Specify the type of analysis and hypothesis type which will be used in the study. Choose between Test, Confidence Interval, Equivalence

In the Goal menu, Means will select all tables related to continuous, ordinal, incidence rate or variance data analysis or testing. Proportions will select tables related to categorical data analysis or testing. Survival will select tables related to time-to-event data analysis or testing. Agreement will select tables related to agreement, correlation or diagnostic screening data analysis or testing. Regression will select tables related to regression analysis or testing for continuous, binary, survival and incidence rate data. Cluster Randomized will select tables related to the analysis or testing of cluster randomized data.

In the Analysis Type menu, Test will select tables where the power of an inequality/superiority hypothesis test will be used. Confidence Interval will select tables where a specific width of the confidence interval is targeted. Equivalence will select tables where the power of an equivalence or non-inferiority hypothesis test will be used.

Table Selection Window: For a given set of table selection menu options or a specific table search query, the table selection window will display all of the tables which are consistent with the menu options or search query. For example, in Figure 1.6 the options selected from the Table Selection Menu are Fixed Term > Means > Two Groups > Test which corresponds to tables which use a fixed term analysis for the power of a test for a two sample superiority hypothesis. Relevant options for this combination of menu options include common tests such as Two Sample t-tests, Wilcoxon/Mann-Whitney Rank-Sum (U) tests and Two Group Repeated Measure ANOVA tests.

To open a table, select the desired table from this Window and select OK in the bottom-right of the window.

If the number of available tables for a given selection criteria is high, you can use the scroll bar on the right-hand side to find tables further down the list for the given query.

Note that some types of statistical test have a large number of variants. These variants are then contained within a sub-menu to reduce the number of options shown by default for easier searchability. These sub-menus are indicated by a "▶" symbol. To open the sub-menu select the ▶ symbol and the options will appear automatically below the sub-menu title. For example, in Figure 1.6 the option of "Two-Sample t-test" is a sub-menu option. If we select the ▶ drop-down menu

then the options will appear below the "Two-Sample t-test" option. This is shown in Figure 1.7.

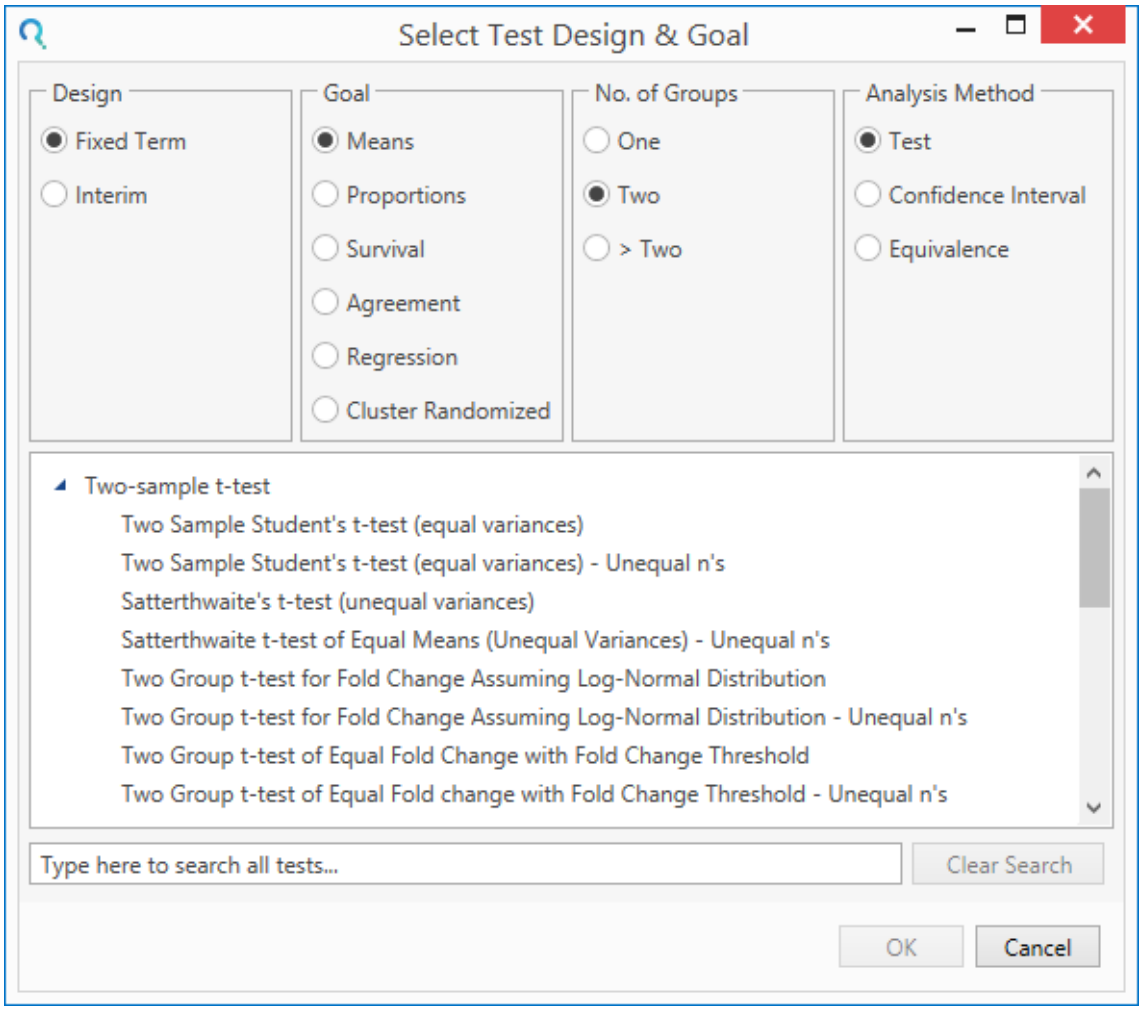

Figure 1.7: Select Test Design and Goal Drop-down

Table Search Bar: The table search bar allows the user to search for tables using a given text search query. For example if a t-test analysis was being considered, entering "t-test" into this search bar field would show all of the potential t-test tables in nQuery. The tables which match the query will be shown in the Table Selection Window. To open the table, select the table in Table Selection Window and select "OK". An example of this is shown in Figure 1.8.

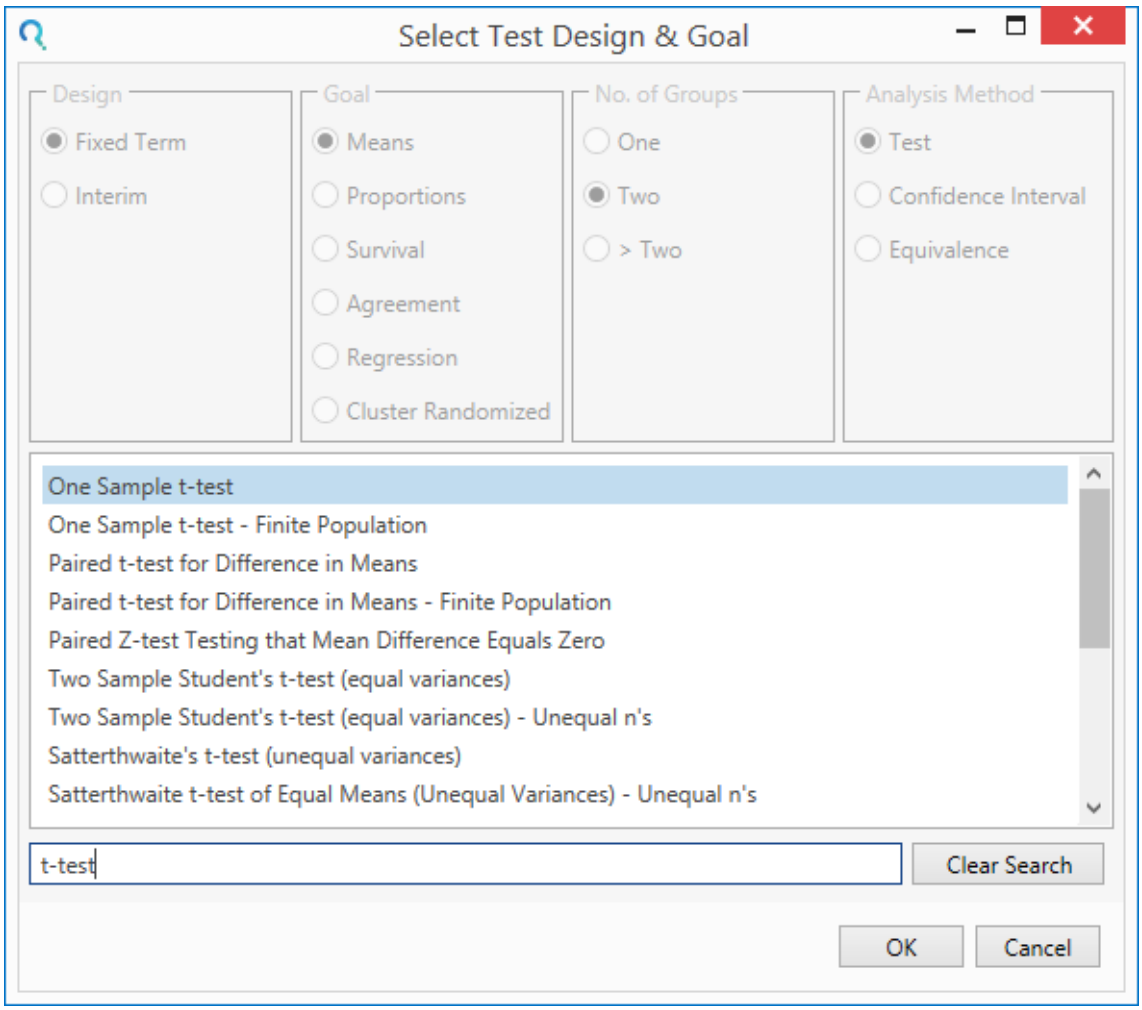

Figure 1.8: Table Search Query Example

Note that while a search query is active the Table Select Menu will be disabled (greyed out). To re-activate the Table Select Menu, click the "Clear Search" button.

OK/Cancel Buttons: These buttons complete the table selection process. When the desired table from the Table Selection Window is selected, click "OK" and that table will be opened in nQuery. To close the Select Test Design and Goal window without opening a table, select the "Cancel" button. Note that selecting the close window option in the top-right is equivalent to selecting "Cancel".

1.6 Design Table Layout

nQuery design tables have a large number of shared design elements. This section will outline the default layout of an nQuery design table and a brief description of these elements and their usage. Many of these elements will be covered in more

detail in future chapters. The default layout of the One Way Analysis of Variance (ANOVA) table is shown in Figure 1.9.

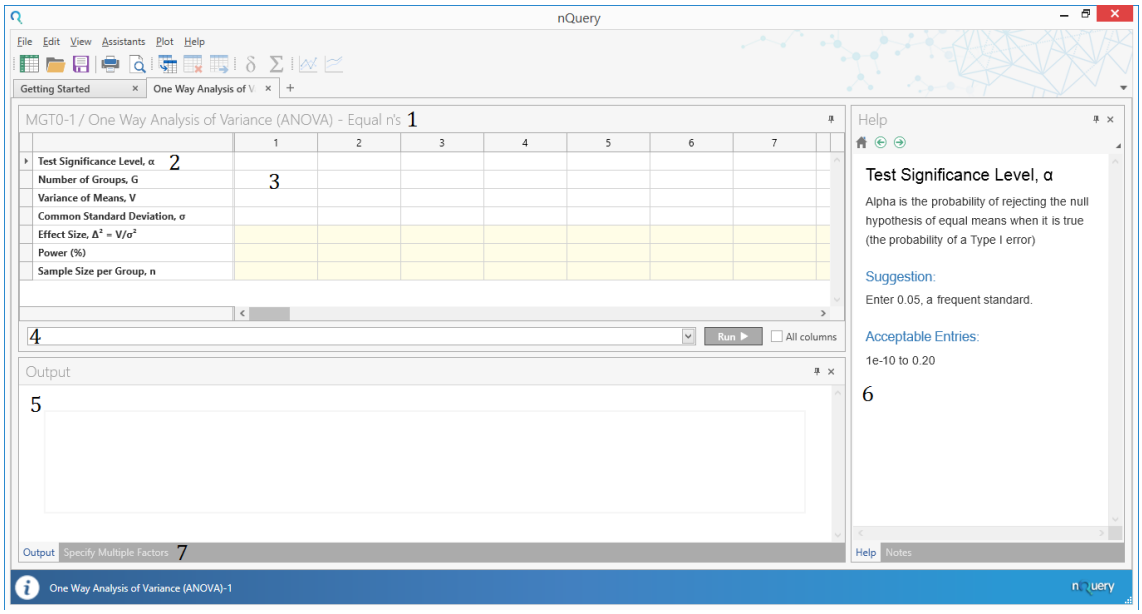

Figure 1.9: One Way Analysis of Variance Design Table

When a design table is opened, there are seven major elements of interest: the table title, the design parameter row-names, the main design table, the solver selection menu, the Output window, the Help window and the window selection tabs.

Table Title: This is the full name of the design table for the currently open table

Design Parameter Row-Names: The names of the design parameters required for the calculations in this design table.

Main Design Table: The table where the values of the design parameters will be entered for a specific calculation.

Solver Drop-down Menu: The menu containing the names of the solvers in this design table and the "Run" button to manually activate solvers.

Output Window: The window containing the output statement summary for a solver calculation.

Help Window: The window containing the help cards for each parameter row.

Window Selection Tabs: Tabs to navigate between a window's options. In the output window, we can select the Specify Multiple Factors Tool and in the Help window, the Notes tool.

The usage of these elements will be explored in the next section using a tutorial example for nQuery.

1.7 Using nQuery Design Tables

1.7.1 Introduction

In this section, a basic example of using an nQuery design table is provided. In addition, the usage of the additional tools and options to assist in conducting a sample size or power calculation is illustrated. Some of these elements will be explored in detail later in this manual.

For this example, the One-Way ANOVA table used in section 1.5 and section 1.6 will continue to be used. This table is shown in Figure 1.10.

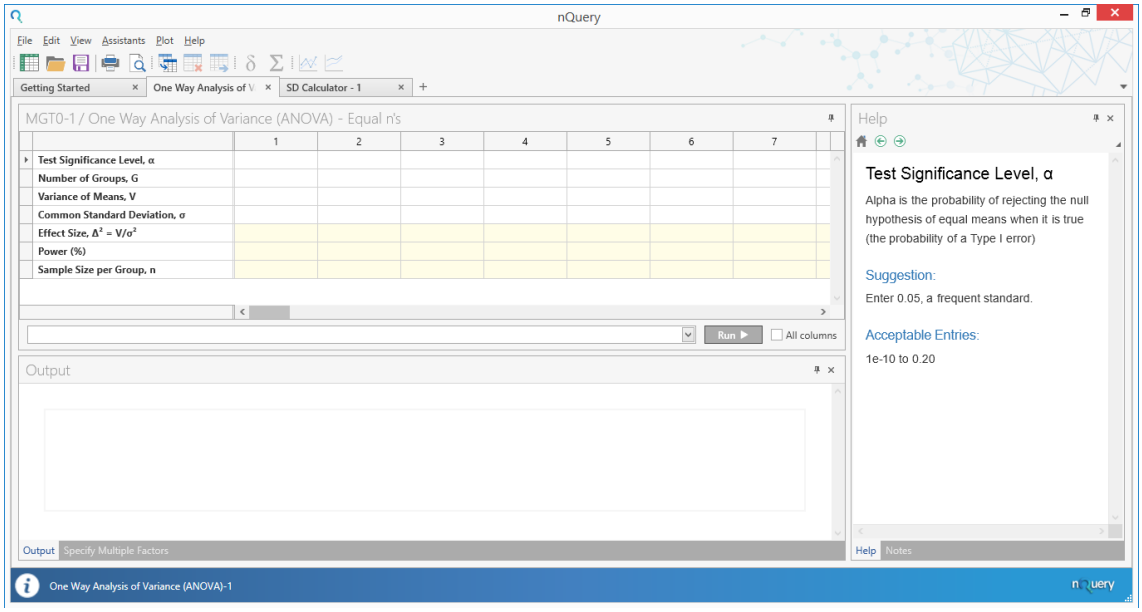

Figure 1.10: Clean One-Way ANOVA Table Example

In the main design table, each row corresponds to a design parameter which is set by the user and each column corresponds to an individual design specification. nQuery allows users to easily calculate the parameters of interest using multiple design specifications using its intuitive spreadsheet format.

For each row in the left-most column, the definition of that row parameter is given. In this example table, the first row corresponds to the test significance level of the proposed study design. For example, this could be set to 0.05, an industry standard.

To conduct a calculation the full set of mandatory design parameters must be filled in a column except for the desired value of interest (e.g. sample size or power). In this table, the design parameters are the test significance level, the number of groups, the variance of means, the common standard deviation, the effect size, the power and the sample size per group.

1.7.2 Table Help

To find additional information or guidance for a design parameter, use the help cards provided for that parameter. If a specific table row is selected, the help window on the right-hand side will dynamically update to relevant help for that design parameter. These help cards have four main potential sections: Main Text, Suggestion, Acceptable Entries and Aid. The **Main Text** section provides a definition of the parameter for the design. The **Suggestion** section provides advice on common values for a parameter or advice on how to derive the value from other known parameters. The **Acceptable Entries** section provides the range of values which are enter-able in a given cell with values outside this range not being usable in the table solvers. The **Aid** section provides information on table assistants which can be used to derive this cell value from other known information (e.g. Effect Size side table, Data Entry tool)

In Figure 1.10, the help card for the test significance level is shown on the right-hand side. This shows the definition of the significance level as the Type I error, the suggested value of 0.05 and the acceptable entries of 1E-8 to 0.20. In Figure 1.11, the help card for the Variance of Means row is shown. In this help card, there is an example of an Aid section which refers to the usage of an Effect Size side-table to derive the value for this parameter.

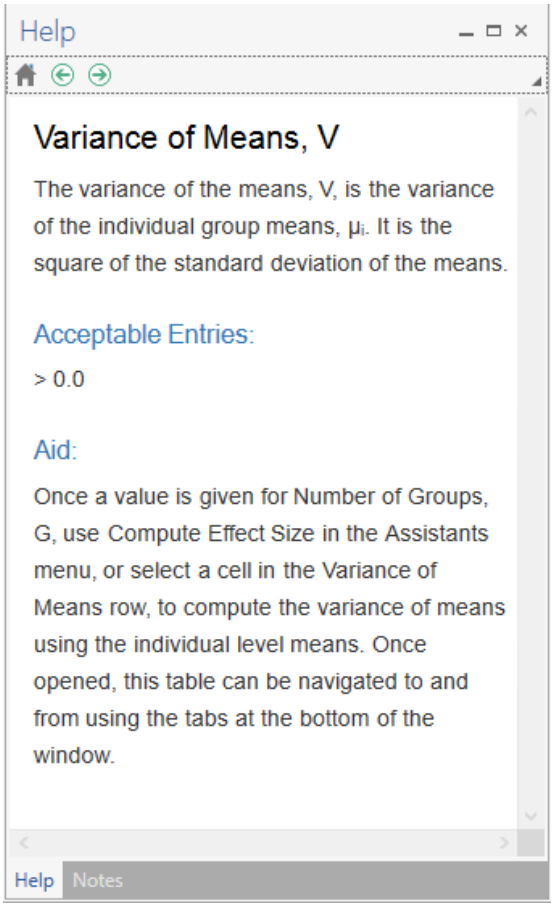

Figure 1.11: Variance of Means Help Card Example

If you want an overall summary for the table, select the Home icon at the top of the Help window. This will provide a brief summary on how to use the table and provide the reference materials used for the design table solver methods. The Home Card for this table is shown in Figure 1.12.

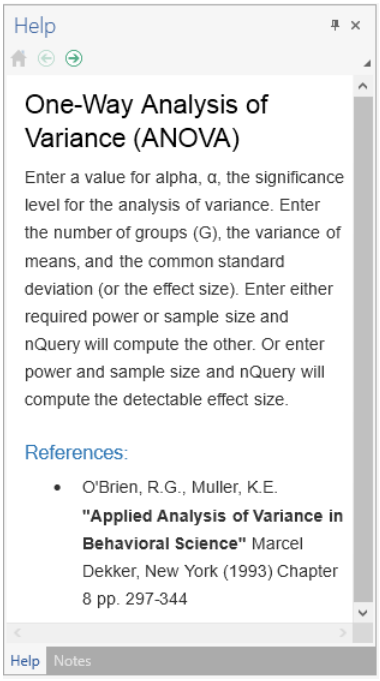

Figure 1.12: Home Card Example

1.7.3 Activating a Table Solver

To conduct a calculation of interest, we require the usage of a table solver. Rows which are highlighted in orange correspond to a table solver i.e. rows for which there is a solver available. In this example, there are solvers available for the effect size, the power and the sample size.

The easiest way to activate a specific solver is to fill in all of the other mandatory design parameters in a table column except for the parameter of interest. The other methods for selecting and activating solvers will be discussed in later chapters.

After the values for the design parameters which will be used in this study design is decided, enter these into the relevant column. In this example, assume a sample size calculation is requested and that values given in Table 1.1 were used for the mandatory design parameters. This design corresponds to a study which will use a one-way analysis of variance to test the null hypothesis that 3 independent group means are equal at the 0.05 significance level with 80% power, assuming a "medium" effect size of 0.5 [Cohen, 1988]. Note that in this table, the variance of means and common standard deviation are optional parameters and are thus not required for a solver to activate. These rows are provided for user convenience to derive the effect size.

Table 1.1: Example Parameter Values

Parameter	Value
Test Significance Level	0.05
Number of Groups	3
Effect Size	0.5
Power	80

When these values are entered into column one of the design table, the sample size is automatically calculated due to being a solver row and all other mandatory parameters being specified. This is illustrated in Figure 1.13

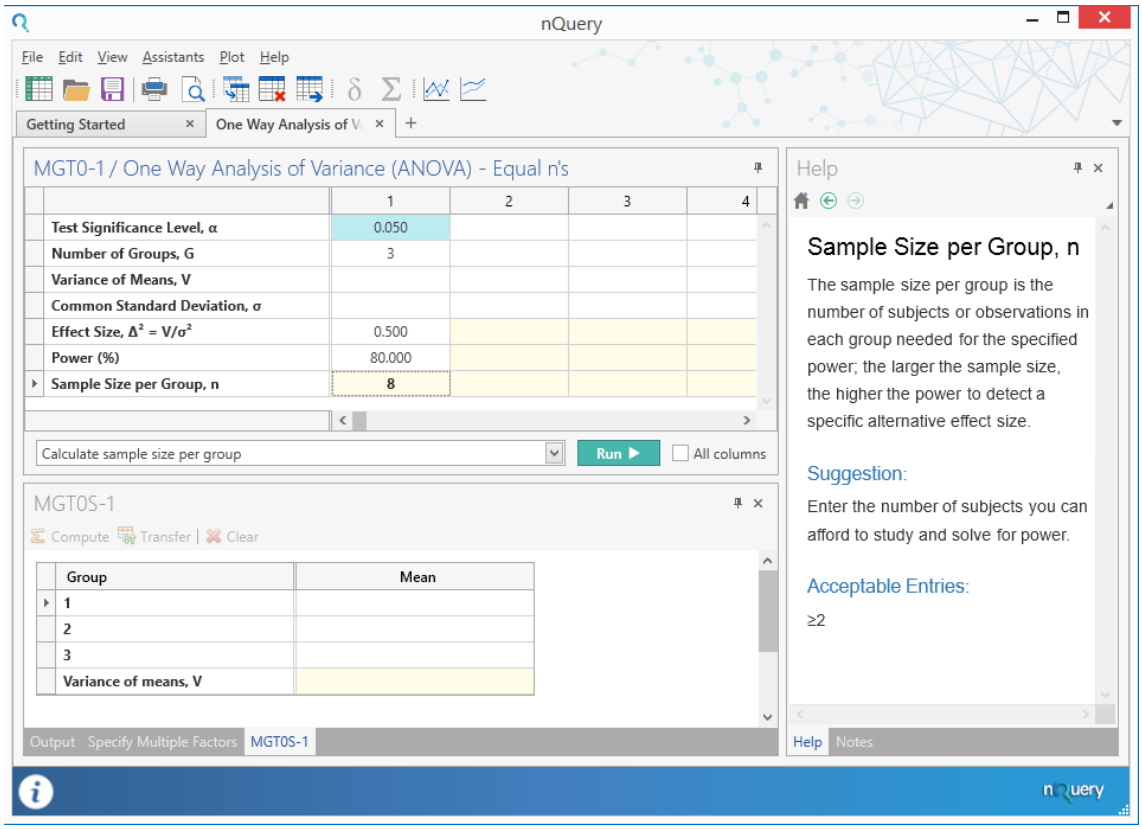

Figure 1.13: Example Solver Answer

In this example, the sample size per group required is 8. This corresponds to a total sample size of 24 (8 x 3). Note that when a solver is activated the solver output (in this case the sample size per group in column 1) will be the only cell still highlighted in orange and that the text of the solver answer will be highlighted in bold.

The other solvers for power and effect size can be activated in a similar fashion. Note that the "Calculate sample size per group" solver is now activated in this column

unless this column is subsequently fully cleared using the "Clear Table" option or "Clear" from the right-click context menu while the column cells are highlighted. For more detail on how to manually select and change solvers, see section 2.1.

This very basic example illustrates the ability of nQuery to easily and quickly find the appropriate sample size for your study.

1.7.4 Effect Size Side-table and Auto-Calculations

Let's expand upon this example by using the optional elements and effect size side table to derive the sample size for a study design where the user can take advantage of the additional information they have available to find better estimates for the required design parameters.

For this example, assume that the common (within-group) standard deviation is assumed to be equal to 1.1. Assume that the expected group means are also known and are equal to the values in Table 1.2.

Table 1.2: Tutorial Group Means Example

Group #	Mean Value
1	1.1
2	2.3
3	3.0

To use these group mean values to derive the Variance in Means, use the Effect Size side-table. To open this side-table, the number of groups must be specified. Once the number of groups is specified in a column, the "Calculate Effect Size" options will become active in the Assistants menu and menu bar when that column is selected. However, the easiest way to open the side-table in a column is to select the derived row in that column and the side-table will automatically open below the main design table (i.e. in the same window as "Output").

Once the side-table is open, enter the group means into the relevant rows. This will activate the "Calculate" button in the top-left of the side-table window. Select "Calculate" and this will show the calculated value of approximately 0.616 in the Variance of Means row in the side-table. This activates the Transfer button. Selecting "Transfer" will transfer the value for the Variance of Means to the relevant row for the current side-table's column of the main table. This is shown in column two in Figure 1.14.

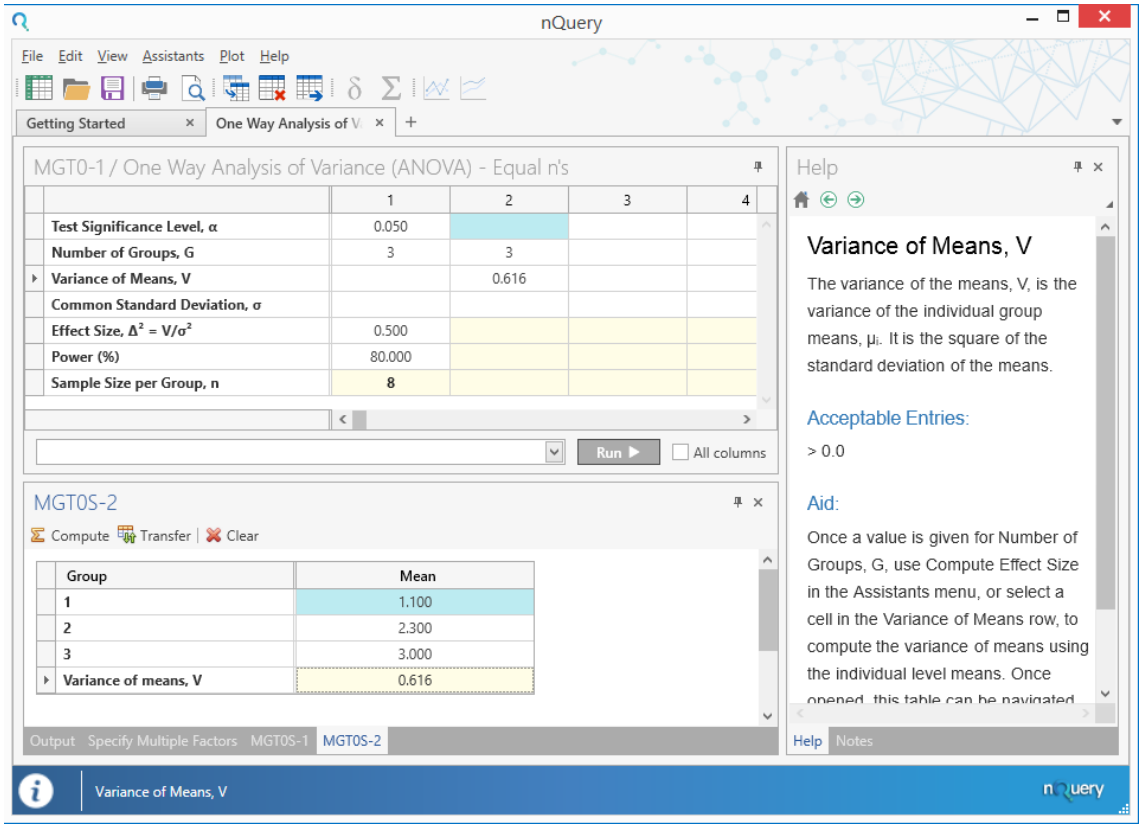

Figure 1.14: Side Table Example

Now enter the common standard deviation value of 1.1. In this table, the effect size can be derived from the variance in means and the common standard deviation and thus when these are both specified the effect size will automatically update to derived value for the values set for those two parameters. In this example, this equals approximately 0.509. This is shown in Figure 1.15.

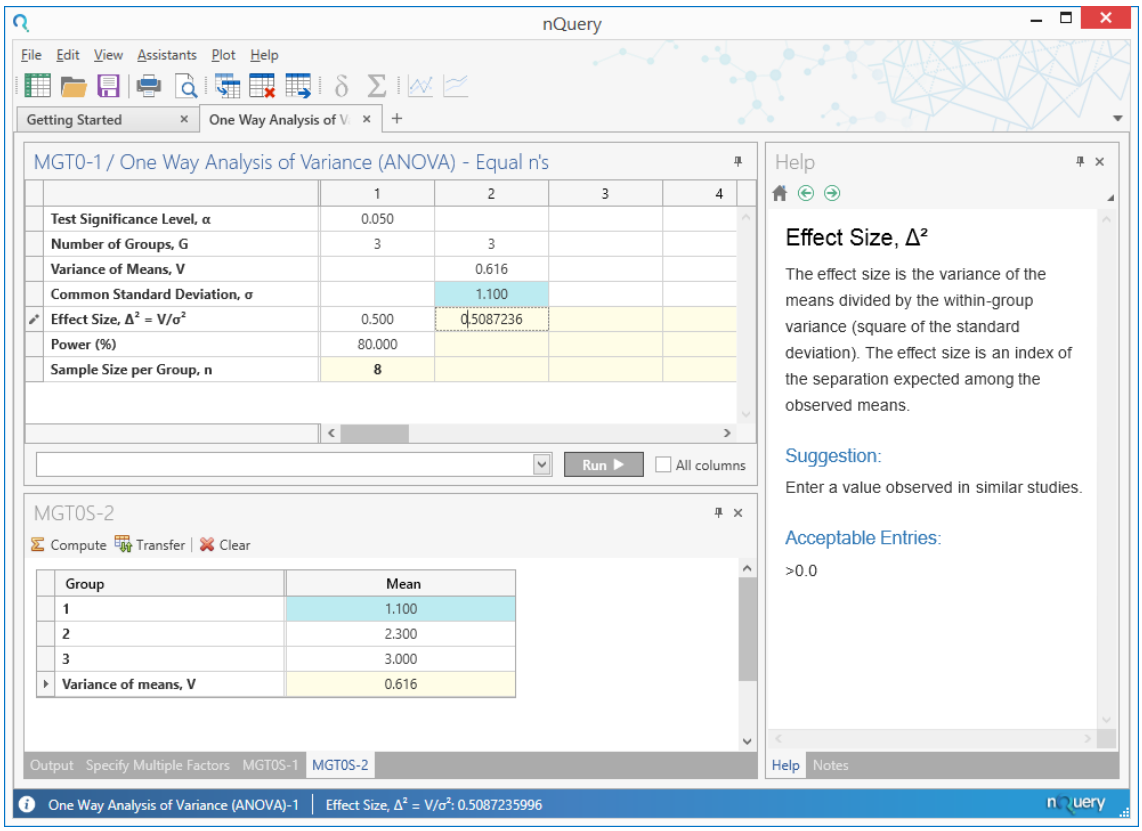

Figure 1.15: Effect Size Auto-Calculation Example

Note that entering two out of three of these parameters (variance, standard deviation, effect size) will always return the third value and that changing one after all three are specified will cause one of the other parameters to change to ensure consistency (e.g. updating the variance of means would cause the common standard deviation to update).

To complete the example, enter the same significance level (0.05) and power (80) as for the previous example. This will give the same sample size per group of 8.

1.7.5 Saving and Opening nQuery Files

nQuery comes with a save file format named .nqt. This file format allows users to save their work in a design table and use it again in future sessions. The save file will retain the design table, side-table, output statements, plots and notes from the time of saving.

To save an nQuery file, the user can use the Save or Save as... options from the File menu or by using the Save option  icon in the tool bar. If the currently open design table has been saved before then selecting the save option will automatically

overwrite that file with the current state of the design table. If the design table has not been saved before or the user selects the Save as... option then the Save menu will appear. Select the desired folder and edit the default save name in the File Name field if desired. Select Save to save the file.

To open an nQuery file, the user can select the Open option from the File menu or by using the Open option  icon in the tool bar. This will open the Open menu. Select the folder which contains the previously saved .nqt file and select the file. Select Open to open the file in nQuery. You can also open an nQuery (.nqt) file by double-clicking the file within Windows explorer. This will open the save file in a currently open instance of nQuery or open an instance of nQuery and open the file in the new instance.

1.8 Design Table Tools

nQuery features a number of additional tools to help users understand and summarise their calculations. The tools which will be covered in this section are the Output statement, the Notes tool, the Specify Multiple Factors tool, the plotting tools and saving and opening an nQuery table.

1.8.1 Output Statement

The output statement provides a verbal summary of the results given in a column after a solver has been activated. This statement can be used as a template for the sample size justification given in study protocol, academic paper or similar document. The output statement is shown by default in the window below the main design table in nQuery. In this example, due to the usage of the effect size side-table, that window currently shows the effect size side table for column two of the main table. To return to the Output statement, select the “Output” tab on the left of the tab bar at the bottom of the lower window. The output statement for the column two calculation will be shown below the main table in Figure 1.16.

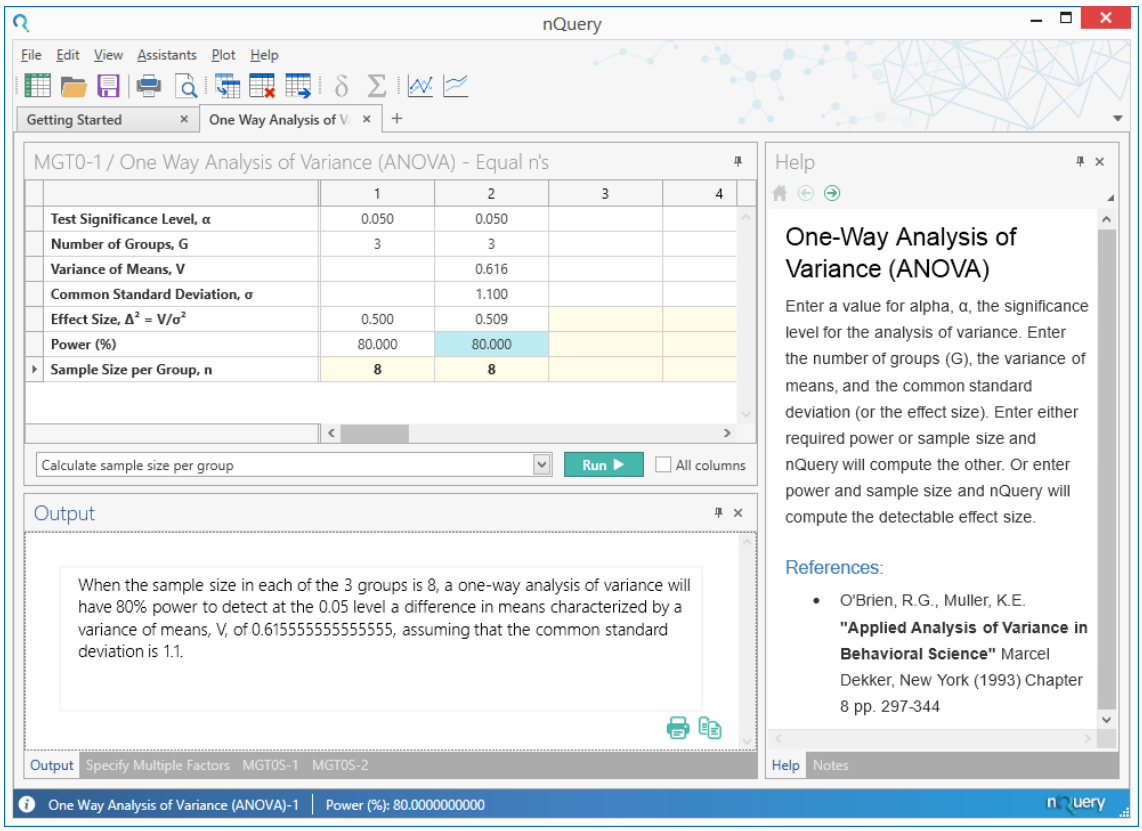

Figure 1.16: Output Statement Tutorial Example

To print the output statement select the Print icon in the bottom-right of the output statement window or use the Print or Print Preview options from the right-click context menu. The Print icon will open a Print Preview window.

To copy the output statement to the clipboard, either use the Copy All icon in the bottom-right of the output statement window or use the Copy option in the right-click context menu. The Copy option will default to Copy All if no text is selected or will copy the selected text if part of output statement text is highlighted.

Note that the output statement will change depending on whether the optional elements have been fully specified in the selected column. For example, the output statement in column 1 will remove the references to optional elements of the variance of means and the common standard deviations and be replaced with the effect size. This is shown in Figure 1.17

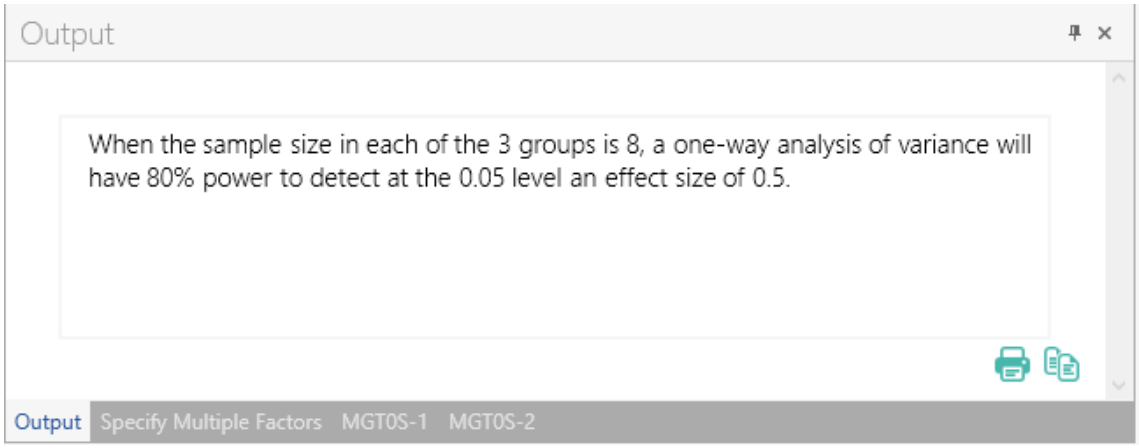

Figure 1.17: No Optional Elements Output Example

1.8.2 Notes Tool

The notes tool is a word processor facility integrated into nQuery which allows users to write and save notes about their calculations directly into their nQuery save file. By default, it is found in the window to the right of the main table. To open the Notes tool, select the "Notes" option from the bottom of the Help window in the tab bar. This Notes Tool window can be seen on the right-hand side of Figure 1.18.

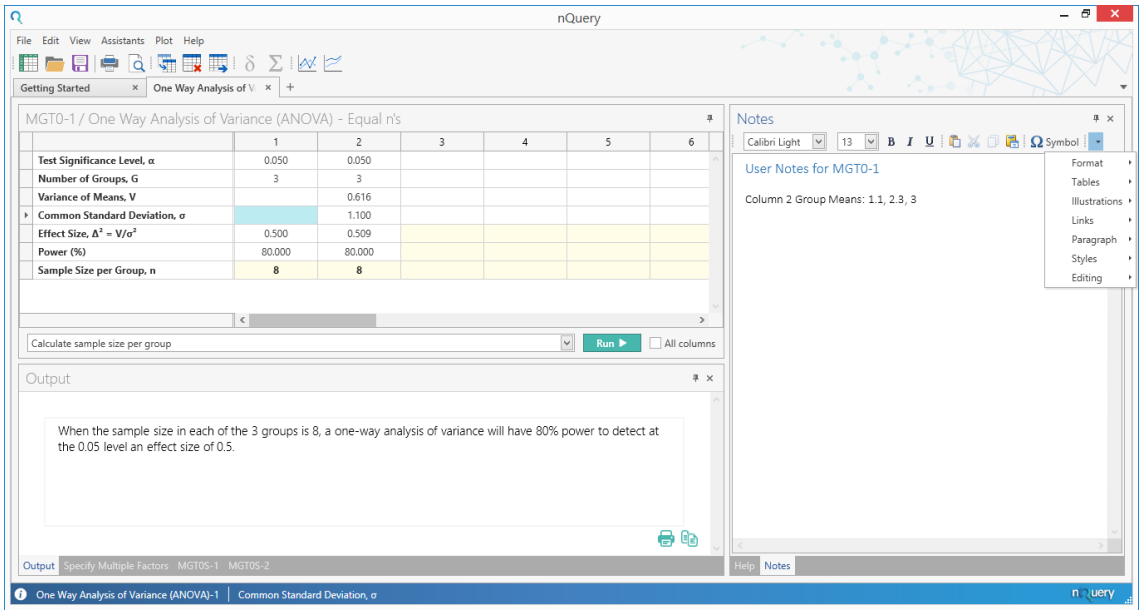

Figure 1.18: Notes Tool Example

At the top of the Notes Tool, the toolbar will contain a drop-down for the font and text size and buttons for the style (bold, italics, underline), copy, cut, paste,

paste special, print and a Symbols menu. The arrow on the right-hand side of the toolbar gives access to a large variety of advanced tools for formatting, tables, images and other options. Right-clicking within the editor will open a context menu containing the Edit tools (copy, cut, paste), indent options, font and paragraph menus, bookmark and hyperlink options.

1.8.3 Specify Multiple Factors Tool

The Specify Multiple Factors tool provides a method to quickly enter multiple values for a design parameter and to generate all combinations if more than one design parameter is varied simultaneously. The Specify Multiple Factors tool is found by default in the window below the main table. To open the Specify Multiple Factors tool, select the "Specify Multiple Factors" option from the bottom of the Output/Side-table window in the tab bar. This Specify Multiple Factors Tool window can be seen on the bottom of Figure 1.19.

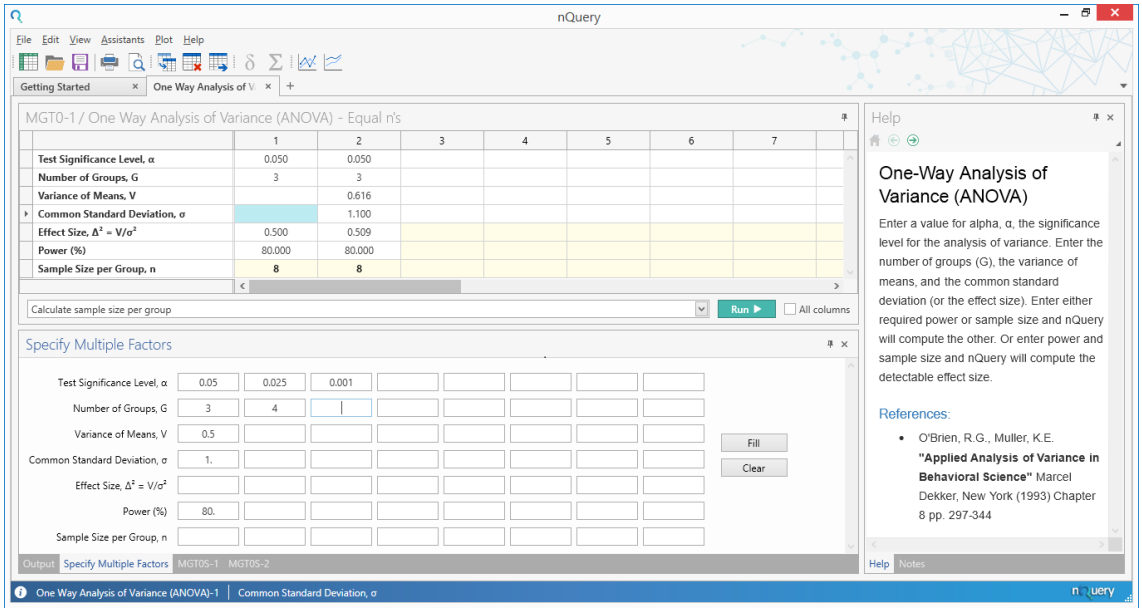

Figure 1.19: Specify Multiple Factors Tool Example

The Specify Multiple Factors tool allows a maximum of eight values per design parameter.

In this example, the significance level will be evaluated at 0.05, 0.025 and 0.001 and the number of groups will be evaluated at 3 and 4. The variance in means, common standard deviation and power will be fixed at 0.5, 1 and 80 respectively.

Note that the effect size row values are left empty since it will be auto-calculated from the variance of means and common standard deviation. The sample size is left empty as this will be solver targetted for this example.

To transfer these inputs into the main table by selecting the "Fill" button on the right-hand side of the window. The results for this set of inputs is shown in Figure 1.20.

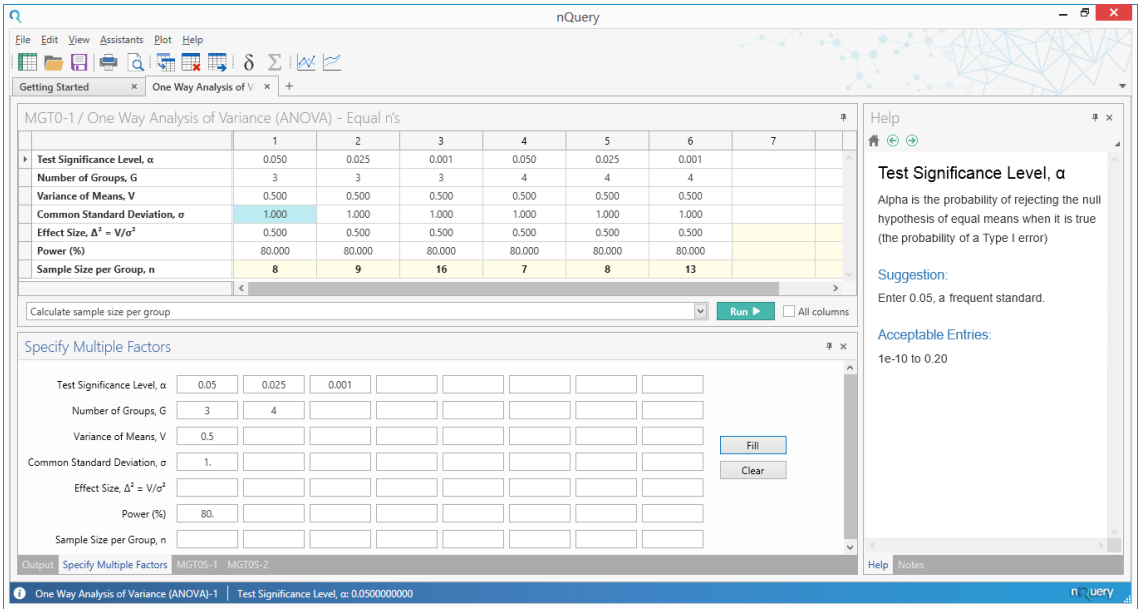

Figure 1.20: Specify Multiple Factors Output Example

Note that the number of columns filled by the Specify Multiple Factors is equal to product of the number of values entered for each row. In this case this equals 3 x 2 x 1 x 1 x 1 = 6.

To clear the inputs from the Specify Multiple Factors tool and reset the tool, select the "Clear" button on the right of the window.

If the number of column combinations exceeds the number of columns open in the table, the table will automatically create additional columns up to the maximum number of columns

If the number of column combinations exceeds the allowed maximum number of columns then a warning will be given and only the first number of combinations equal to the maximum number of columns will be shown. The maximum number of columns is configurable in Options menu.

1.8.4 Plot Power vs Sample Size

The Power vs Sample Size plot provides a quick way to assess the relationship between power and sample size over a wider range while fixing the other elements in calculation. The Power vs. Sample Size plot can be activated in two ways: using the Plot menu drop-down option "Plot Power vs Sample Size" or using the tool bar

"Plot Power vs Sample Size" button  . Both of these options will be inactive until a column is selected which has all of the mandatory elements except for the power and sample size.

After a column is filled appropriately, select the Plot menu or menu bar options and the plot will appear in the application foreground. If you want to compare up to 6 columns, select cells in the desired columns and open the plot and the results for all columns will be displayed simultaneously. Note that holding Ctrl (to select individual columns) or Shift (to select a range of columns) while clicking columns may assist in selecting multiple columns.

An example of an individual and multiple column versions of the Power vs Sample Size plot are shown in Figure 1.21. In this example, the plots are for the six columns generated from the Specify Multiple Factors example above, with the individual plot being from column 1.

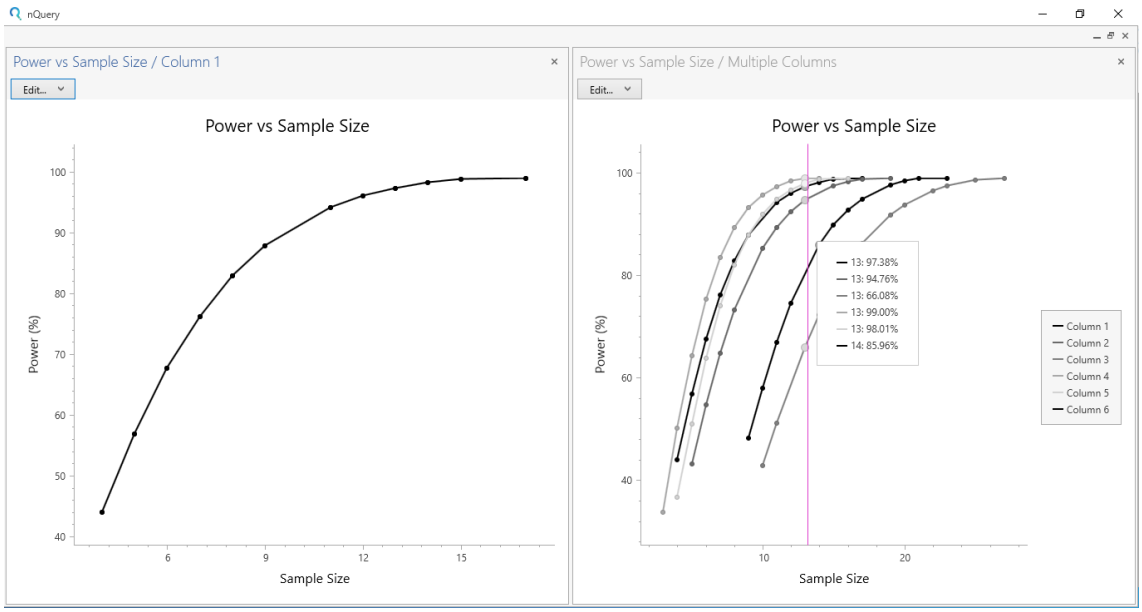

Figure 1.21: Power vs Sample Size Example

Highlighting a point in the plot will display the exact values for the X-axis and Y-axis value selected for each column.

Right-clicking on the plot will open a context menu which provides options to Print, open a Print Preview, Save the plot as an image file and options to edit the plot (seesection 1.9)

1.8.5 Plot User Selected Rows

The Plot User Selected Rows plot provides a visual and customised way to assess the relationship between any of the continuous or integer design parameters and any of the solvers given all other unselected design parameters are fixed. The Plot User Selected Rows menu can be activated in two ways: using the Plot menu drop-down option "Plot User Selected Rows" or using the tool bar "Plot User Selected Row" button  . Both of these options will be inactive until a column is filled sufficiently for at least one valid custom plot to be available.

After a column is filled appropriately, select the Plot menu or menu bar option and this will open the "Select X-axis, Y-axis" window. This menu is shown in Figure 1.22.

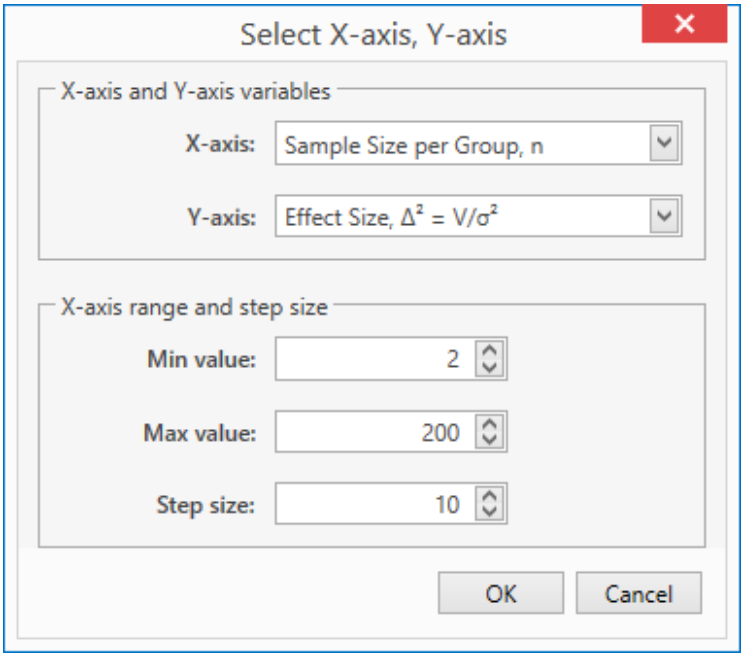

Figure 1.22: "Select X-axis, Y-axis" Plot User Selected Rows Menu Example

In this menu, the X-axis drop-down will contain the full list of design parameters which are of either a continuous or integer type. The Y-axis drop-down will contain the list of the all the solvers available in the table. This table allows us to create a combination of any of these two lists to create a custom user plot.

When an X-axis option is selected, the user can specify a range over which that design parameter will be varied in the plot by specifying a minimum value and maximum value for this range. Note that the OK button will be disabled and an error displayed if the "Min value" is below the lower limit of the acceptable entries for the selected design parameter or is greater than the current "Max value" or if

the "Max value" is above the upper limit of the acceptable entries for the selected design parameter or is lower than the current "Min value".

After a range is specified, the user can select the "Step Size". This equals the increments that the X-axis parameter will be increased in when moving from the specified Min and Max values. The number of steps can be found by dividing the range by the step size (round down if answer is a non-integer). Note that the Step Size must be positive and must give a number of steps which is between 4 and 1000. If this condition is not met then an error will be displayed and the OK button greyed out.

Default values will be given for the Max value, Min value and Step size for each X-axis option in a given table.

If you want to compare up to 6 columns for given X-axis and Y-axis option selection, select cells in the desired columns and open the plot and the results for all columns will be displayed simultaneously. Note that holding Ctrl (to select individual columns) or Shift (to select a range of columns) while clicking columns may assist in selecting multiple columns.

An example of an individual and multiple column versions of the Plot User Selected Rows plot are shown in Figure 1.23. In this example, the plots are for the six columns generated from the Specify Multiple Factors example above. The individual column plot is for column 1 where the X-axis option was the Test Significance Level, using the default 0.01 to 0.2 by 0.0095 X-axis range, and Y-axis option was the Effect Size. The multiple columns plot is for the six column where the X-axis option was the Sample Size per Group, with a custom range of 5 to 30 by 1 X-axis range, and Y-axis option was Effect Size.

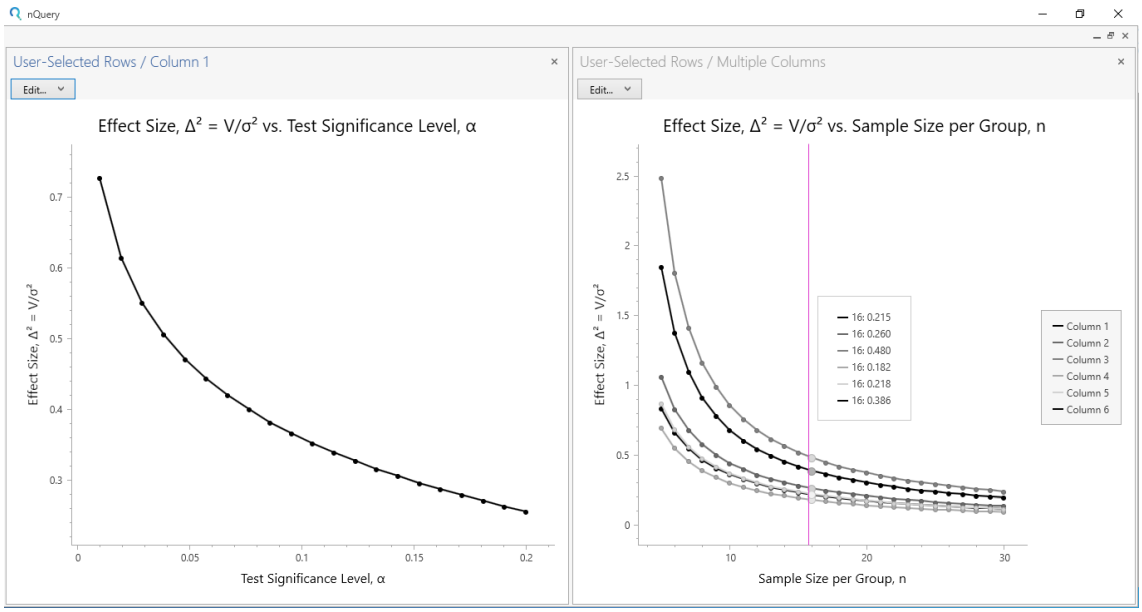

Figure 1.23: Plot User Selected Rows Example

Highlighting a point in the plot will display the exact values for the X-axis and Y-axis value selected for each column.

Right-clicking on the plot will open a context menu which provides options to Print, open a Print Preview, Save the plot as an image file and options to edit the plot (seesection 1.9)

1.9 Editing Plots

nQuery provides options to edit the default Power vs N and User Selected rows plots. To edit these plots, select the "Edit" button from the top-left corner of the plot or use the context menu by right-clicking on the plot area.

In the edit menu, there will be options for Print, Print Preview and Save As which allow the printing and saving of the current plot. There are the following options for editing a plot

1. Titles (Edit Table and Axis Titles)
2. Series (Edit Series names and colors)
3. Gridlines (Add and Edit Gridlines)
4. Legend (Add and Edit Legend)

1.9.1 Titles Options

The Titles menu contains options for the three main table titles: the Chart Title (above plot), the X-axis Title (below plot), the Y-axis Title (left of plot). Each of these titles has the same three options for editing:

1. Hide/Show Title: Hide or show selected title in plot
2. Rename: Rename title in "Rename Title" dialog window
3. Font: Edit Font, Size, Style (Bold, Italics) and Text Color ("More Colors" for custom colors) in "Update Title Font" dialog window

1.9.2 Series Options

The Series menu contains options for updating the x-axis range and editing the color and name of each series (column) in the plot. If a single series is in the current plot the the following three options will be available:

1. Update X-axis range: Edit range of x-axis displayed in plot
2. Rename: Rename series (column) in "Rename Series" dialog window

3. Color: Change the color used for series (column) in "Select Series Color" dialog window. Custom colors can be created by selecting the "More colors" option

If more than two series (columns) are selected in the current plot, this menu will have the "Update X-axis range" option and the names of all the series in the current plot. If you select a specific series, the Rename and Color options will be available for that series. These options work as above.

1.9.3 Gridlines Options

The Gridlines menu contains options to add and remove major gridlines for the X and Y axes. This menu contains the following options:

1. Hide/Show X-axis gridlines: Hide or show X-axis gridlines depending if shown at present
2. Hide/Show Y-axis gridlines: Hide or show Y-axis gridlines depending if shown at present

These gridlines will be in-line with the tick marks for the relevant axis.

1.9.4 Legend Options

The Legend menu contains options to add and remove the legend and to edit the horizontal and vertical position of the legend. This menu contains the following options:

1. Hide/Show Legend: Hide or show legend depending if shown at present
2. Horizontal Position: Selecting this will open a sub-menu for legend locations. This menu gives five options to place the legend: Left inside, Left outside, Center, Right inside, Right outside. Select the desired option to change the position
3. Vertical Position: Selecting this will open a sub-menu for legend locations. This menu gives five options to place the legend: Top inside, Top outside, Center, Bottom inside, Bottom outside. Select the desired option to change the position

"Inside" and "Outside" refer to whether the legend should be included in the plot area itself (inside) or place outside the plot area (outside). Note that changing the line colors and legend series names is done using the Series options described above. Series names can also be edited using the column name feature by double-clicking the column title in the main table before creating a plot.

2 Advanced nQuery Table Features

2.1 Manual Solver Activation

In nQuery, the easiest way to generate a result is to use the auto-activation of a solver. An example of auto-activation is given in subsection 1.7.3. In short, a solver will automatically activate if all the mandatory elements in a design table are filled except for a solver row (solver rows are orange in the design table).

However, in certain cases, a user may wish to fix a column to use a specific solver regardless of user inputs or change the solver being used in an already complete column. Both of these are covered in the following section and will make use of the **Solver Drop-down Menu and Run button** highlighted in Figure 1.9.

2.1.1 Fixing a Column's Solver

To fix the solver which will be used in a given column, open a design table and select the column or columns which will have their solver fixed. To do this, select a cell or cells in the column(s) of interest or select the full column(s) by selecting the column title(s). This can be achieved by clicking and dragging for multiple cells or by holding down Ctrl (individual cells/columns) or Shift (range of cells/columns) while selecting cells or columns.

After the cell(s) or column(s) are selected, the solver can be fixed by selecting a specific option from the Solver Drop-down menu. These options are illustrated for the One Sample Chi-Square Test in Figure 2.1.

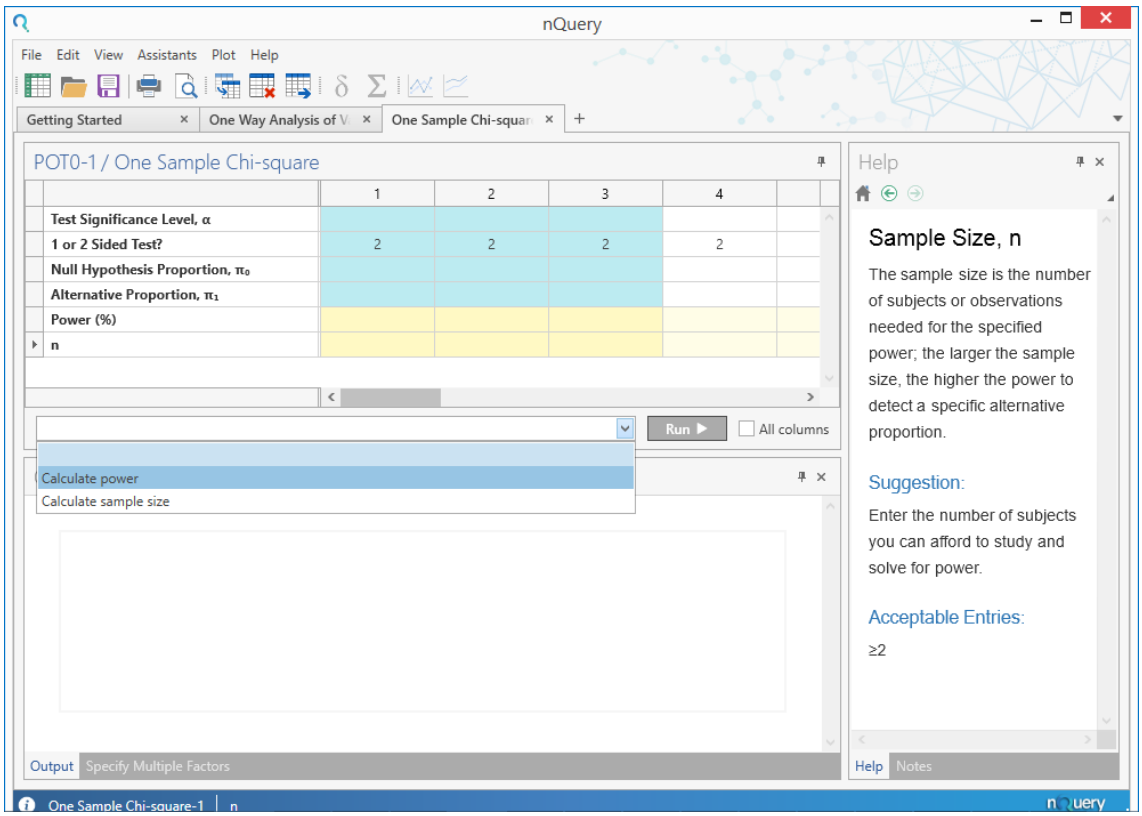

Figure 2.1: Solver Drop-down Menu Example

The One Sample Chi-Square test has two solvers available: Calculate power and Calculate sample size. In this example, the power solver will be fixed for the first three columns of the table by selecting the "Calculate power" option from the solver drop-down while the first three columns are highlighted as in Figure 2.1.

As the solver is fixed for these three columns, entering the design parameters in a column which would activate the power solver will calculate the appropriate power for those design parameters. However, entering parameter values which would activate the sample size solver does not generate any solver output. This is illustrated in Figure 2.2.

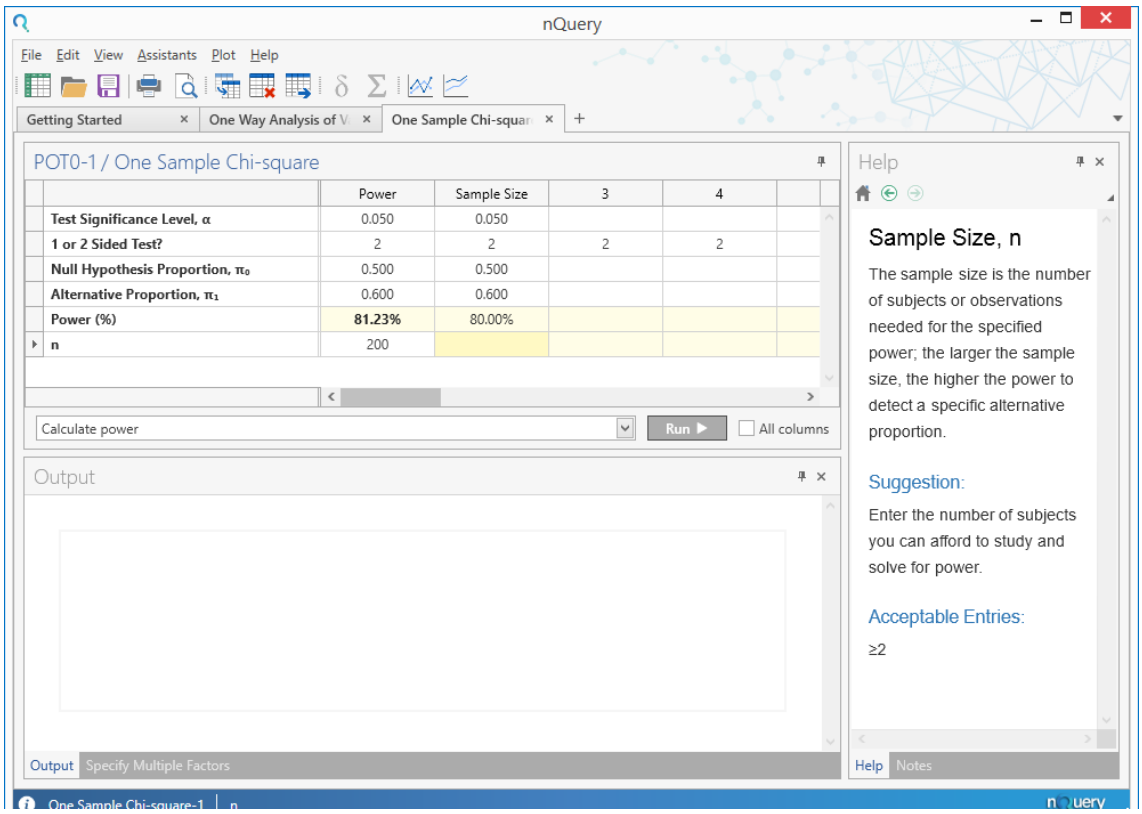

Figure 2.2: Solver Manual Activation Example

In this example, the first column has successfully generated the power but in the second column the values generate no result even though this column would generate the sample size if it was in its default state. The other solver(s) would be selected and would work in the same fashion.

Note that in every table the first option in the solver drop-down is a "blank" option. If this option is selected, the selected column(s) will reset to auto-selecting a solver based on the current column inputs.

2.1.2 Changing a Filled Column's Solver

After a solver is activated and generates a solver output result (either automatically or manually) in a column, the user may be interested in changing the solver being used in that column. There are two main methods to achieve this: Reset the column or Select an alternative solver

2.1.2.1 Reset Column

To reset a column, there are three main methods: Clear Table, Clear Column, Blank Solver Option

1. Clear Table: Selecting Clear Table from the Edit menu or the menu bar will reset all columns to their default state where any solver can be automatically or manually activated.
2. Clear Column: Highlighting a column(s) by selecting the column title or all the cells in a column and then selecting "Clear" from the right-click context menu will return the column to its default state
3. Blank Solver Option: Selecting a column and selecting the first "blank" option from the Solver Drop-down menu will reset the column to its default state

Unlike the Clear options, the Blank Solver option can be used on a partially filled or a fully filled column. If a column if fully filled and the blank option is selected, a new calculation can be activated by changing a value in the affected column(s) or by selecting the "Run" button to the right of the solver drop-down menu.

2.1.2.2 Select Alternative Solver

To select an alternative solver select the column(s) or interest, open the Solver Drop-down menu and select the solver which is desired. This works the same as the blank solver option above. After an alternative solver is selected, it will be activated if a cell is changed in the affected column(s) or by using the "Run" button to the right of the solver drop-down menu. Selecting the sample size solver and using the "Run" button is illustrated in Figure 2.3.

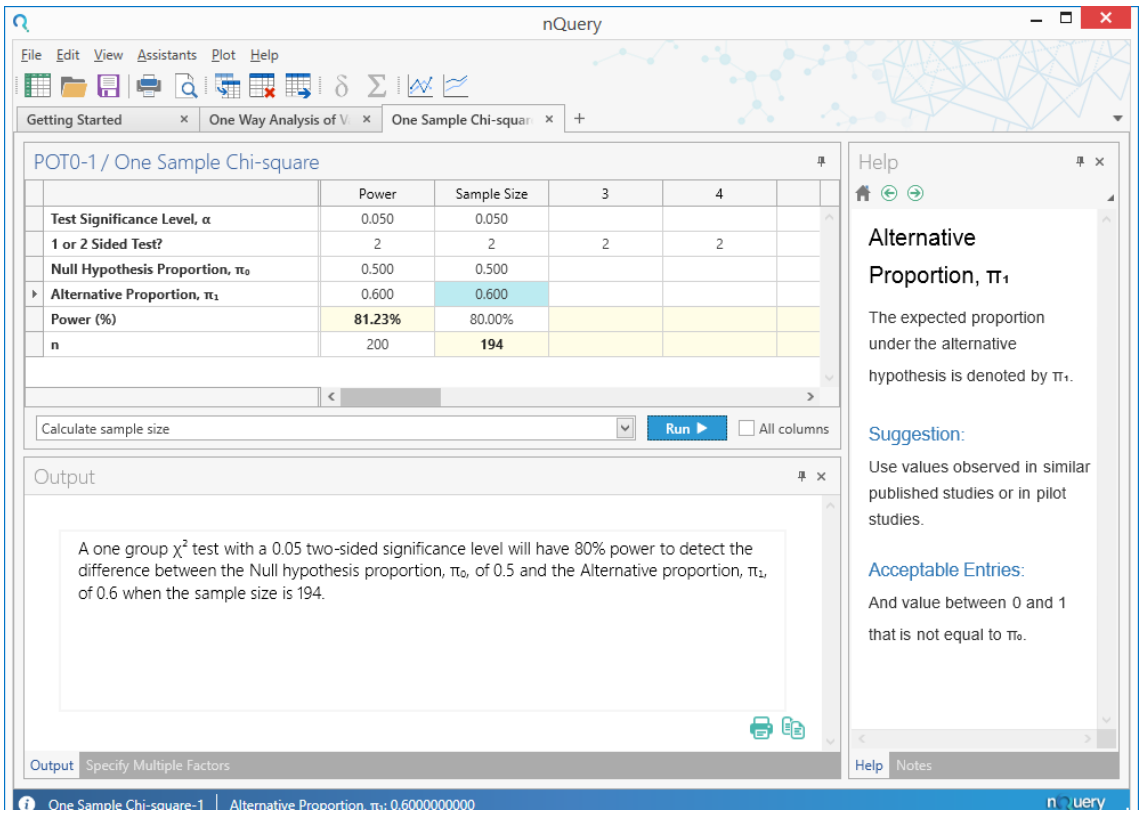

Figure 2.3: Manual Solver Change Example

2.2 Errors and Warnings

nQuery provides a number of errors and warnings to prevent incorrect values being entered or when a solver cannot find a legitimate answer. There are three primary types of errors and warnings in nQuery: Out of Range errors, Solver errors, Solver Warnings.

2.2.1 Out of Range Errors

Out of range errors occur when a value is entered into a design table or side-table cell which is outside the acceptable entries for that cell. The acceptable entries for a given cell can be found in the Help card for that cell (see subsection 1.7.2).

Note that out of range errors can occur due to table auto-calculations. For example, the auto-calculation for the Effect Size row in One Sample t-test can be outside the acceptable entries of 0.1 to 40 when derived from the optional element values of the null hypothesis mean, alternative hypothesis mean and standard deviation. This is also common in side-tables where there are read-only auto-calculated cells which

contain the sum of the writeable elements of the side-tables (e.g. Chi-Square Test of Specified Proportions in C Categories).

When an out or range error occurs, a red ⊗ symbol will appear on the left-hand side of the affected cell. Placing the mouse over the cell will provide a tooltip which gives a brief summary of the allowable values in the cell. An example of this is shown in Figure 2.4.

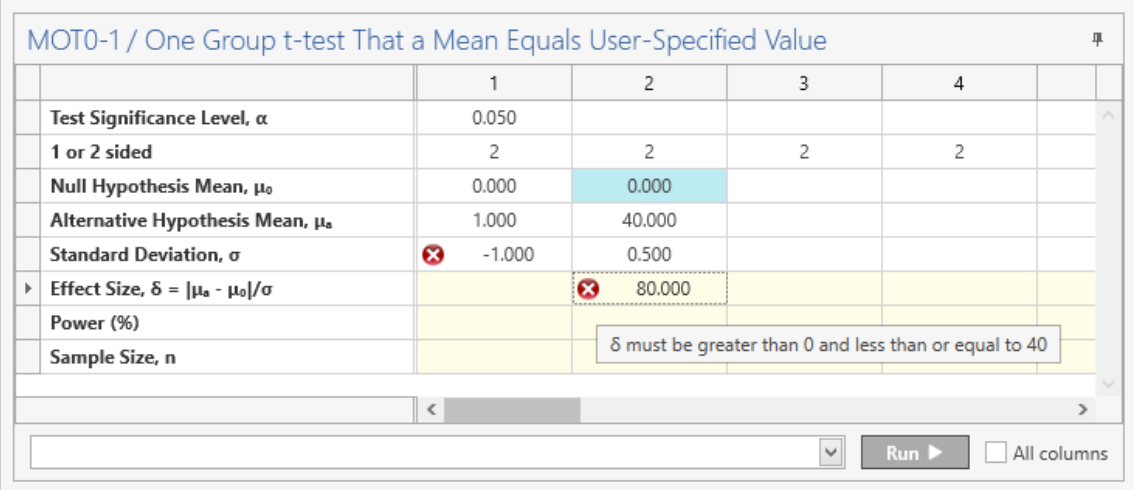

Figure 2.4: Out of Range Error Example

While an out of range error is present in a design table column, no solvers will activate in that column. In a side-table, an out of range error will prevent the Compute/Transfer buttons from activating.

2.2.2 Solver Errors

Solver errors occur when a solver cannot find a solution for the given inputs. This is most common when the design parameters used are extreme or when the solver needs the use of array search or memory intensive methods. These errors indicate that either there is no correct result for the design parameters given or that the correct result was too extreme to be found by the solver algorithm.

When a solver error occurs, the affected cell will be highlighted in red. If you hover the mouse over the affected cell, a brief description of the error will be given. An example is given in Figure 2.5 where the population size is too low causing there to be no legitimate output for the Lgamma function due to there being no value for the adjusted sample size which would achieve 80% power.

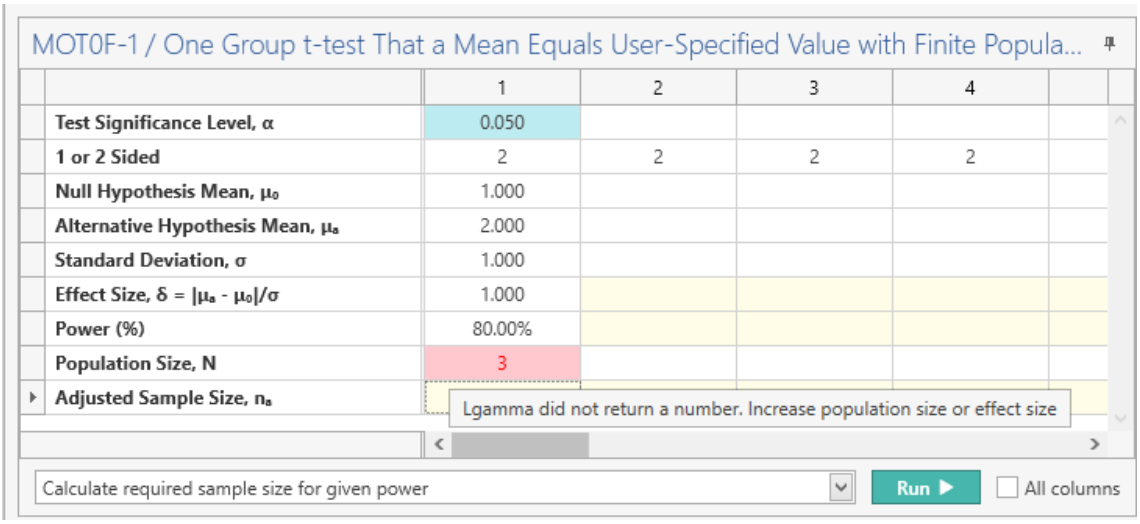

Figure 2.5: Solver Error Example

2.2.3 Solver Warnings

Solver warnings occur when a solver result is given but there are issues which would cause concern over the veracity of the result given. The two major categories of solver warnings are rounding warnings and assumption warnings.

1. Rounding Warnings: These occur if the solver output was rounded due to being too extreme originally.
2. Assumption Warnings: These occur if an underlying assumption of study design is not met by the current design parameters.

Examples of rounding warnings include rounding sample size up to two when it was below two and rounding the power down to 99% when it was greater than 99%. Examples of assumption warnings include the minimum cell count being too low for chi-square tests and the normal approximation not holding when the population size and sample size are similar when the finite population adjustment is being used.

When a solver warning occurs, the background is clear and the solver answer value is red. This is illustrated in Figure 2.6 where column one has a “Calculate power” solver warning for the minimum cell count being low and the column two has a “Calculate power” solver warning that the power was rounded down to 99% due to being “unrealistically” high.

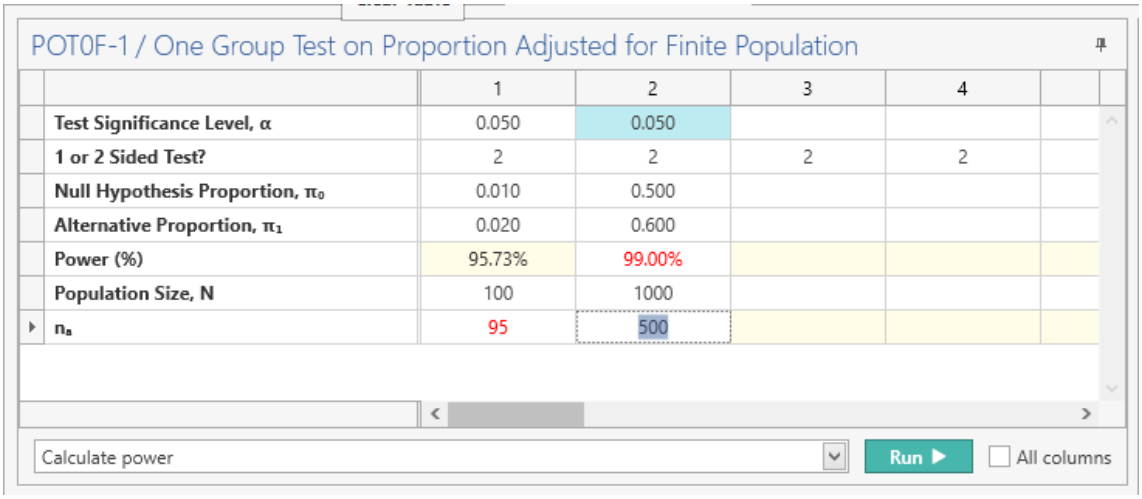

Figure 2.6: Solver Warning Example

2.2.4 Application Logging

When nQuery opens a table or encounters an error, these are logged automatically by nQuery. By default, these logs are saved in "C:/Users/<Username>/AppData/Roaming/nQuery/
The save location can be changed in the Options menu (see section 2.3). If you want to track the errors in the application or wish to provide additional support to Statsols Technical Support then these logs will be useful. An example is given in Figure 2.7.

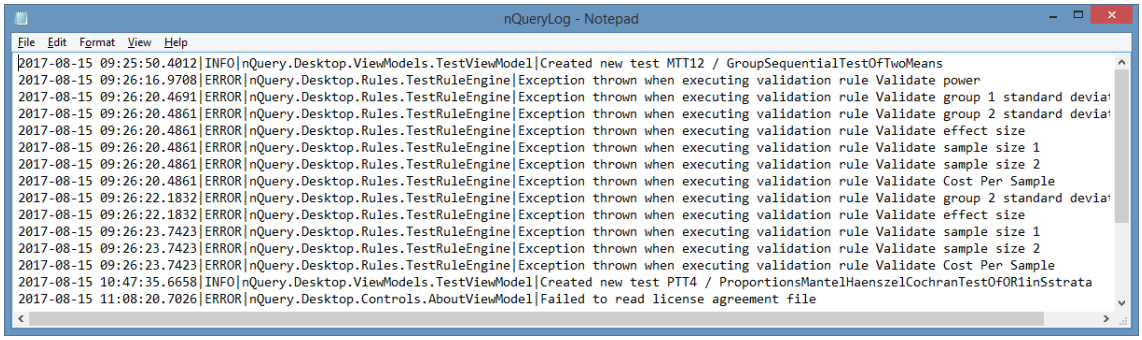

Figure 2.7: nQuery Log Example

2.3 Options Menu

The options menu gives the user the ability to change a number of nQuery features to reflect a users preferences.

2.3.1 Application-Wide Options

The application-wide options are shown in Figure 2.8.

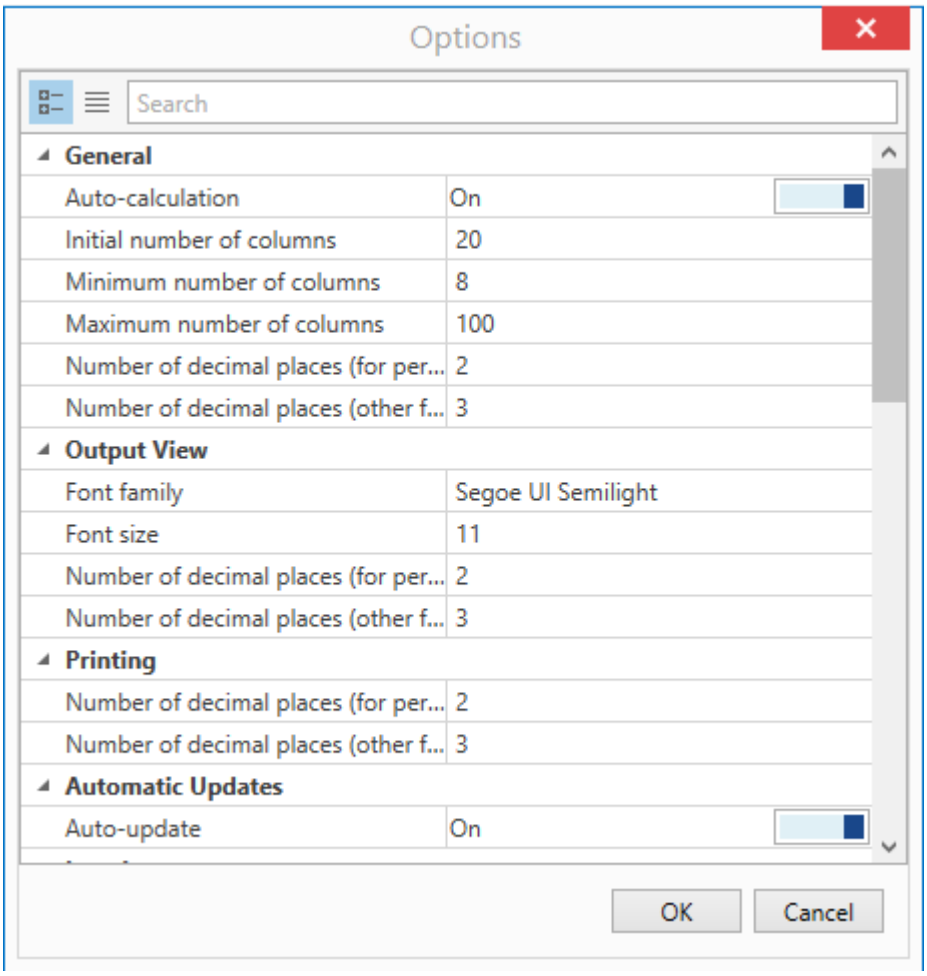

Figure 2.8: Application Options Menu

There are three main categories of application-wide settings. These are named and described below:

1. General Settings: Auto-calculation, Initial number of columns, Minimum number of columns, Maximum number of columns, Number of decimal places (percentages, other values)
    a) Auto-calculation: Toggle whether solvers are activated automatically if column is filled appropriately. If disabled, solvers are only activated via the Run button.
    b) Initial number of columns: Set the number of columns included in a design table when opened initially

   c) Minimum number of columns: Set the minimum number of columns in a design table. Resets to this value if the number of Specify Multiple Factors combinations is less than this value.

   d) Maximum number of columns: Set the maximum number of columns in a design table. Resets to this value if the number of Specify Multiple Factors combinations is more than this value.

   e) Number of decimal places (for percentages): Set number of decimals displayed for percentage rows in tables e.g. power

   f) Number of decimal places (other floating-point values): Set number of decimals displayed for all non-percentage rows in tables e.g. means, proportions

2. Output View: Font family, Font size, Number of decimal places (percentages, other values)

   a) Font Family: Select the font family used in the Output window text

   b) Font Size: Select the font size used in the Output window text

   c) Number of decimal places (for percentages): Set number of decimals displayed for percentage rows in output statement e.g. power

   d) Number of decimal places (other floating-point values): Set number of decimals displayed for all non-percentage rows in output statement e.g. means, proportions

3. Printing : Font family, Font size, Number of decimal places (percentages, other values)

   a) Number of decimal places (for percentages): Set number of decimals displayed for percentage rows when printing e.g. power

   b) Number of decimal places (other floating-point values): Set number of decimals displayed for all non-percentage rows when printing e.g. means, proportions

4. Automatic Updates: Auto-update

   a) Auto-update: Select if software should attempt to automatically download and install latest version if available on start-up

5. Logging: Logging, Log level, Log files folder.

   a) Logging: Toggle where logging occurs

   b) Log Level: Change the level of event logged by nQuery. Options are Info, Debug, Error, Warning

6. Chart: Options to edit all default options for plots. These are shown in Figure 2.9.

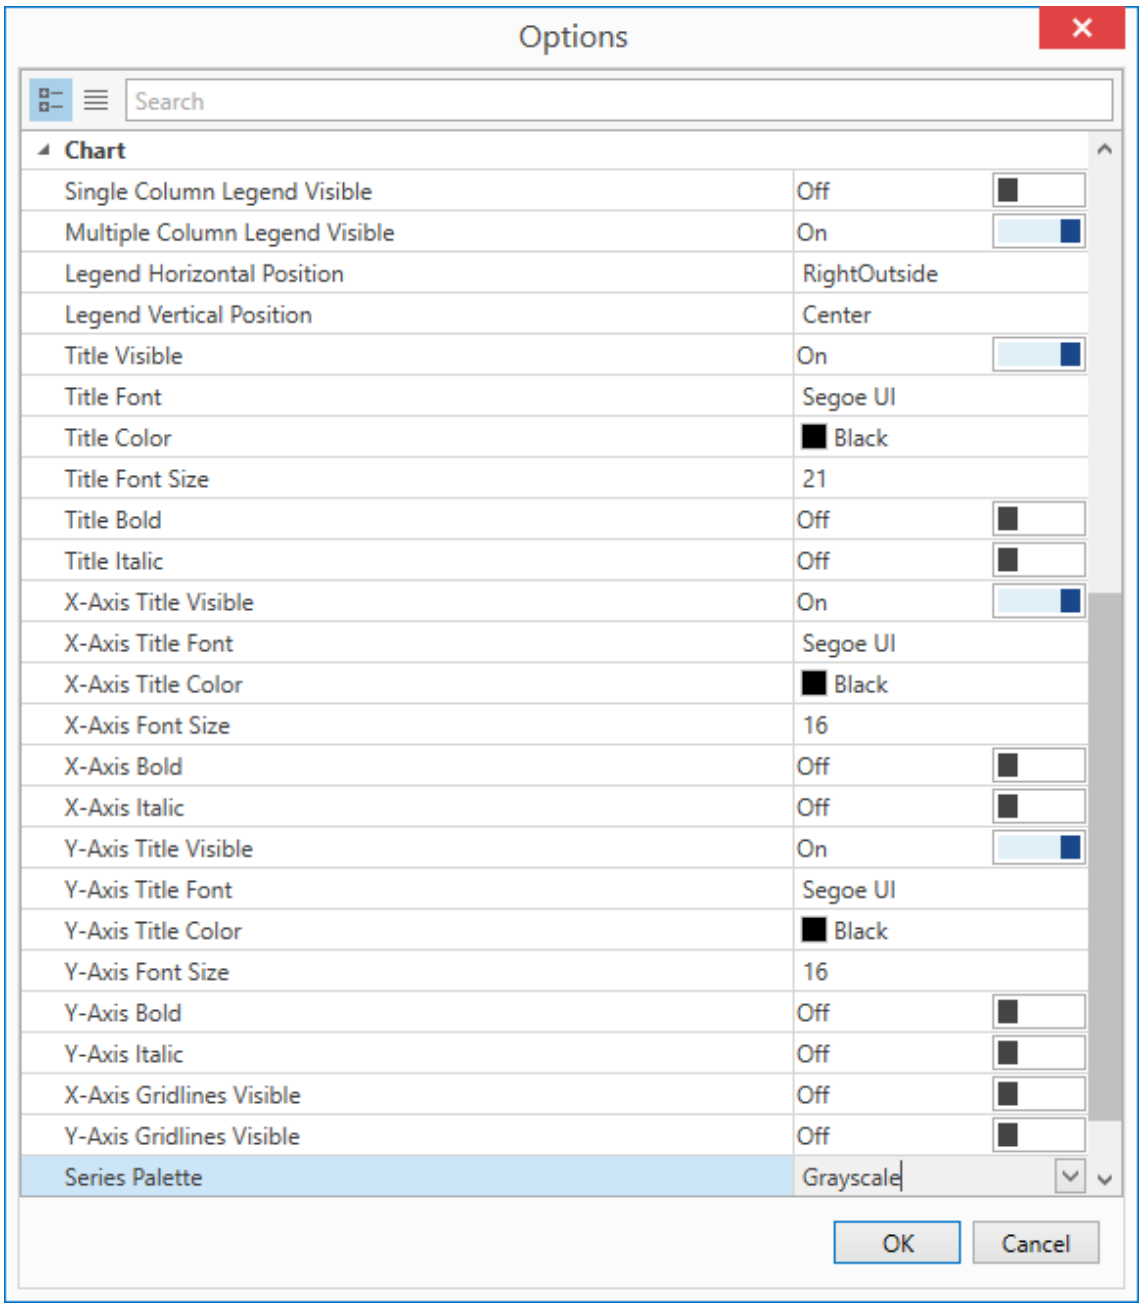

Figure 2.9: Plot Options

2.3.2 Per-Table Options

The General Settings options can also be applied on a per-table basis. To achieve this, select the arrow to the left of "Tables" under the "Per-Table Configuration Settings". Then select the arrow to the right of table code name of interest. The same settings as above for General Settings can then be edited for that specific table.

An example of this is shown in Figure 2.10.

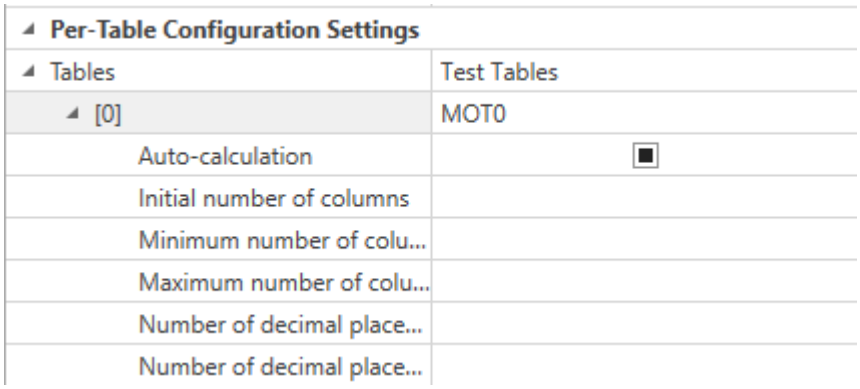

Figure 2.10: Per-Table Options Menu Example

2.4 Print Options

nQuery provides comprehensive printing options for the design tables and other table outputs.

2.4.1 Print Table

To Print the contents of the table, select either the "Print" option under the File menu or from the tool bar. This will display the default Windows print screen where you can edit the printer used, the printer preferences, the pages printed and number of copies printed.

2.4.1.1 Print Preview

To have nQuery generate a Print Preview, select either "Print Preview" option under the File menu or from the tool bar. This will display a Print Preview screen with an example shown in Figure 2.11.

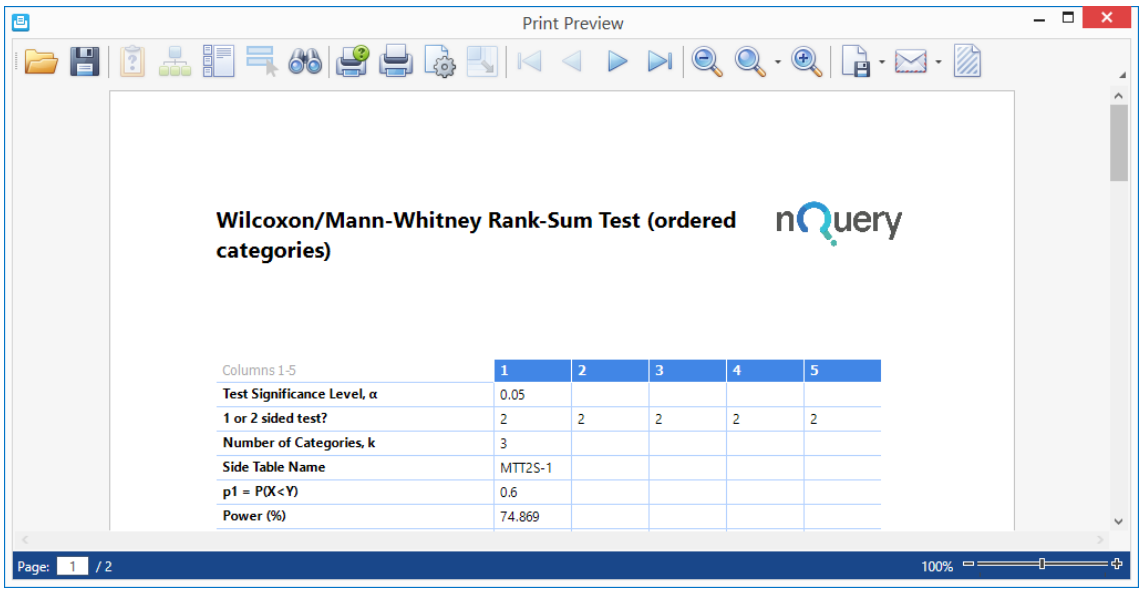

Figure 2.11: Print Preview Example

The toolbar at the top of the Print Preview screen gives options to open, save, search, print, quick print, adjust settings, page up, page down, first page, last page, zoom, zoom in, zoom out, export, send or watermark. A hand tool is also available from the right-click context menu.

2.4.2 Other Print Options

Print options are available for the following table outputs: Plots and the Output Statement.

To print plots, use the Print option in the right-click context menu used within a plot.

To print the output statement, use the Print button in the bottom-right of the Output window or in the right-click context menu.

2.5 Miscellaneous

2.5.1 Design Table Context Menu and Shortcuts

Within an nQuery design table when a cell or cells is selected, right-clicking will open a context menu for those cells. In nQuery when a cell or multiple cells are selected, the context menu will contain options for Copy, Cut, Paste, Select All, Fill Right, Clear and Copy Table.

The shortcuts for standard Windows commands are used in nQuery and are displayed to the right of the name within the context menu.

If a full column is selected either by dragging over all the cells in a column or by clicking the column title bar, selecting Clear will reset the solver in the column. It also opens a Report option which is described in subsection 2.5.3.

2.5.2 Design Table Drop-downs

In some nQuery design tables, a row is a drop-down menu rather than a numeric input. To select a value from this drop-down, select the relevant cell and click the downwards arrow on the right-hand side of the cell. This will display the full set of options available for that row. The help card will provide a description of the solver drop-down options, their meaning and usually a Suggestion for this option.

One special class of drop-down menu is for "1 or 2 Sided Test?" or "1 or 2 Sided Interval?" rows. In these rows, you have the option to either enter 1 or 2 directly into the cell or use the drop-down menu as above.

These rows are treated differently in the Specify Multiple Factors tool and the Plot User Selected rows tool to other rows. In the Specify Multiple Factors tool, these rows will have a drop-down menu similar to design table drop-down. In Plot User Selected rows, these rows will not be available as an option in the X-axis drop-down. To plot the effect of changing the options in these cells, fill multiple columns with all other values fixed and then apply the relevant plotting method to all the relevant columns in the same plot.

2.5.3 Report

nQuery provides a report function which allows users to summarise the results for specific design table column in a print-ready format. To open a report for a specific column, select that column by clicking the column name bar above the column. Then right-click the column and select the "Report" option at the bottom of the context menu or select "Report" from the Assistants file menu. This will open a column report which will include a summary of the design parameter values in that column and the output statement for that column if a solver is active. The report will be opened in Print Preview window and has the same options and functionality described in subsubsection 2.4.1.1.

An example of report is shown in Figure 2.12.

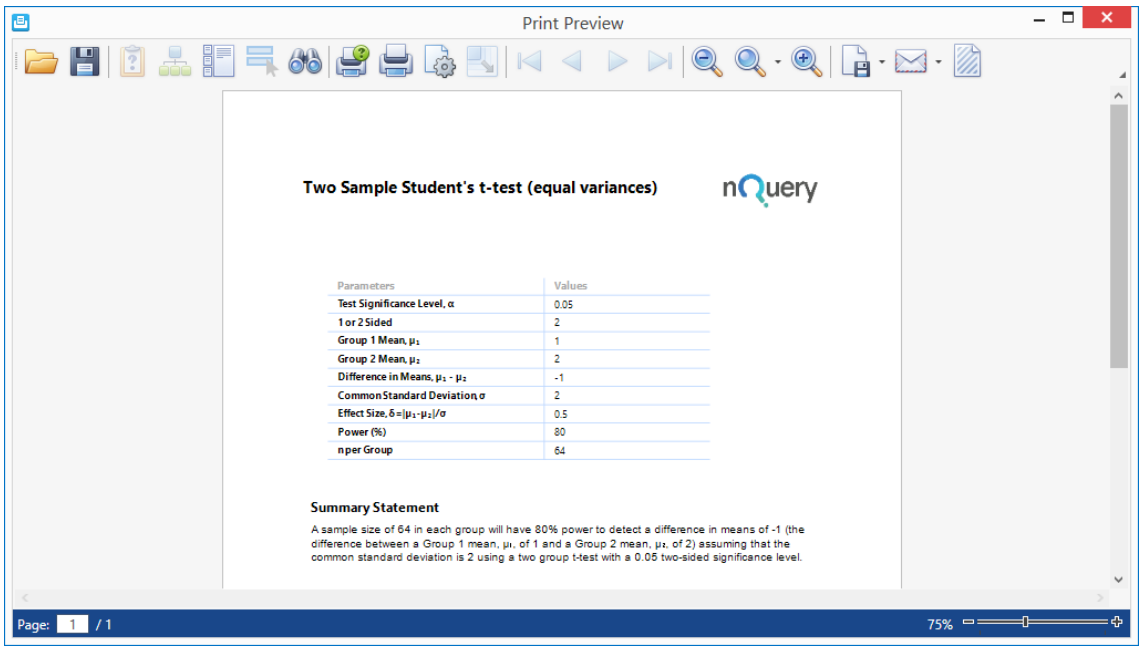

Figure 2.12: Report Example

2.5.4 Customising Table Layout

2.5.4.1 Changing Window Size and Docking

In nQuery, the default size and positioning of each window and element is editable by user to create the layout which best suits their needs. There are two main elements which can be edited: windows and tabs.

Manipulating nQuery Windows Windows are the primary user interface unit of an nQuery design table. By default the three windows displayed are the design table window (top-left), the Output/Specify Multiple Factors window (bottom-left) and the Help/Notes window (right). If a plot or plots are created, the Plot window can be manipulated in the same fashion.

There are two main actions that can be applied to a window: docking/undocking and resizing.

- Docking/Undocking

To move a window, click and drag the window bar at the top of the window of interest. This bar will contain the name of the tab currently open in that window (e.g. Help, Output etc.). While the window is being dragged, two options are given for that window: dock the window or undock the window.

To dock the window, drag the cursor while the window is selected to a "docking square" and nQuery will highlight in blue where the window will be docked. Four

edge docking squares will be available on the centre left, right, top and bottom of the nQuery window which will place the current element to the left/right/above/below all other nQuery windows. In addition, the "current" nQuery window over which the cursor is placed will contain the relational docking squares. The centre square will combine the current window and the docking window into a single window with all of their tabs combined. The other squares will place the docking window to the left/right/above/below the current window .

To undock the window, drag the cursor away from the "docking squares" and to the undocked placement of choice.

In both cases, when the desired placement is achieved let go of the left mouse button (or other "clicking" method) and the window will appear in the desired place. An example of the docking process "in action" for the Output window where it is being placed between the MTT0 Design Table and Help windows is given in Figure 2.13. The effect of this placement (after resizing) is shown in Figure 2.14.

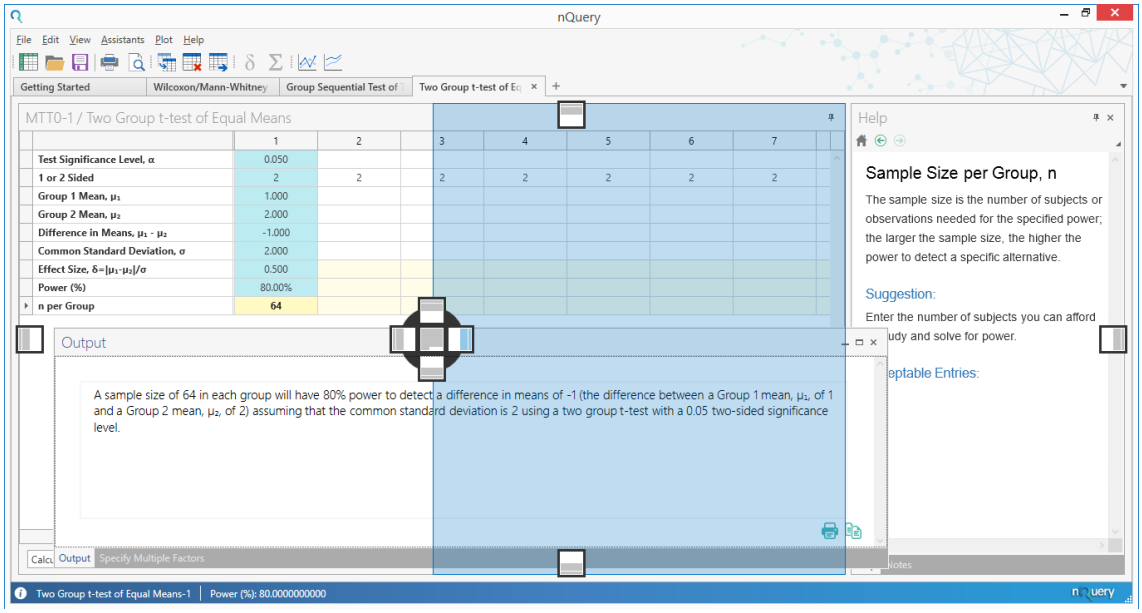

Figure 2.13: nQuery Window Docking Example

- Resizing

To resize a window, place the cursor at the edge of the nQuery window of interest. The mouse cursor will change to a ↔ symbol. Then hold down the left mouse button and drag the the cursor left/right (for a vertical edge) or up/down (for a horizontal edge) until the window is the desired size. Note that elements in the table will dynamically update to show how they will look for a given window size. An example of the output statement being resized after the docking example above is shown in Figure 2.14.

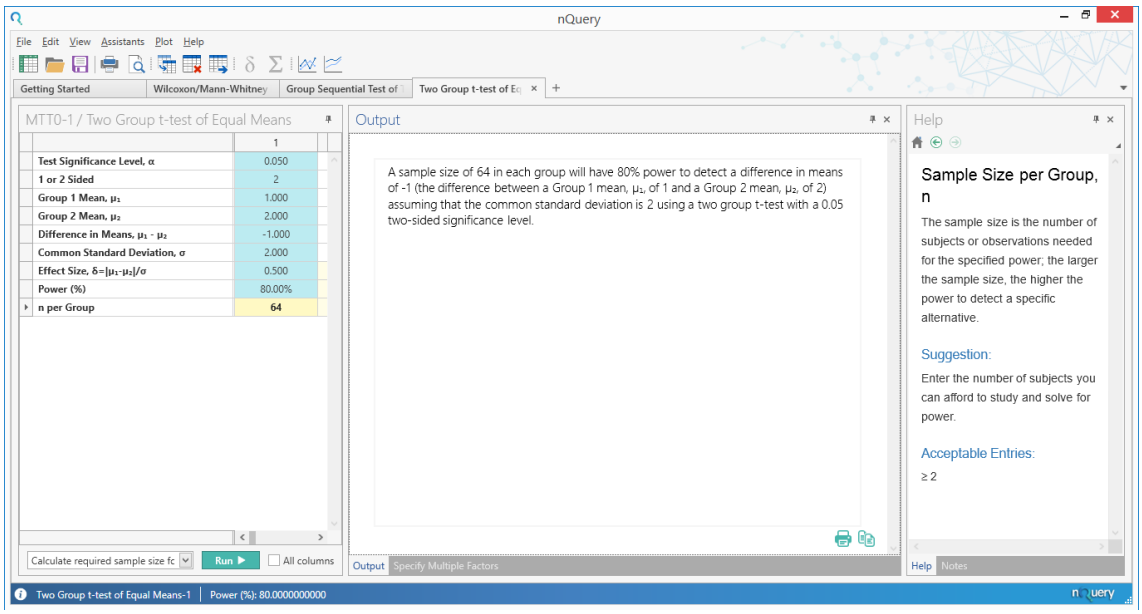

Figure 2.14: nQuery Window Resizing Example

Manipulating Window Tabs Window tabs have a single applicable operation: docking/undocking. This process works similarly to undocking/docking a window. For tabs, the window element which is clicked and dragged is the tab bar name rather than the window name. The tab names are found in the grey tab bar at the bottom of an nQuery window. When a tab is being dragged, the same options are available as for a window. The three main actions are to undock the tab, place the tab relative to the other windows or dock the tab into another window.

To undock, drag the select tab away from any of the docking squares and leave. When a tab is undocked, it will become an nQuery window.

To place relative to the other windows, select on of the left/right/up/down docking squares and place in the desired position which will be previewed in blue by hovering over a docking square while dragging.

To dock the tab into another window, drag the cursor over the window of interest and select the centre docking square in the middle of the current window.

An example where the Output tab is currently being moved and where the Notes tab has been moved into the main design table window is shown in Figure 2.15.

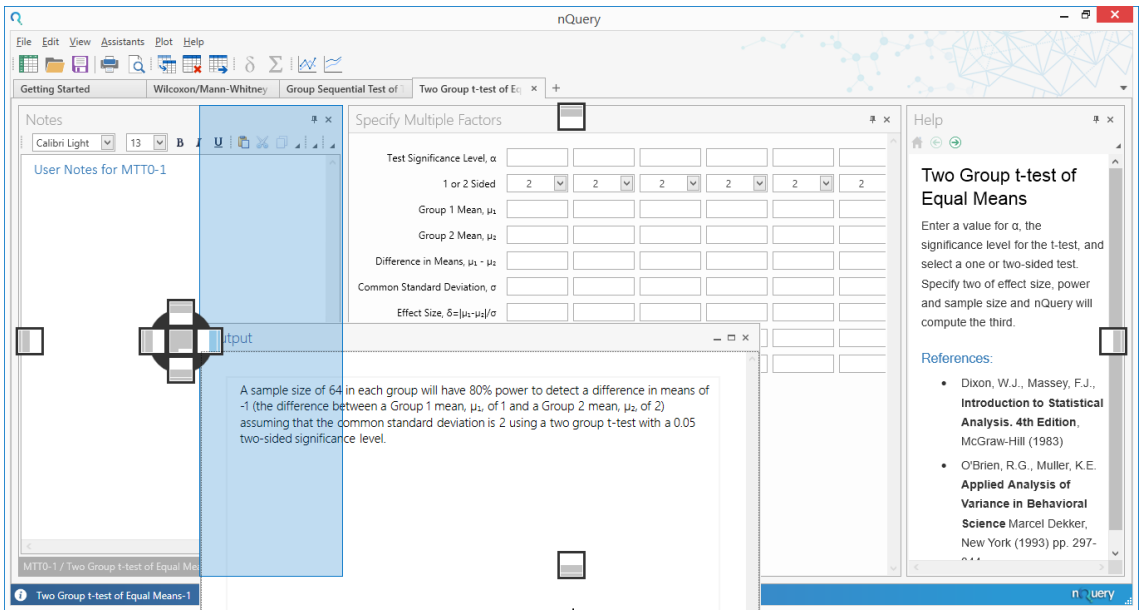

Figure 2.15: nQuery Tab Docking Example

Note that you can also drag tab names within the tab bar to place it relative to the other tab names in the same click and drag fashion.

2.5.4.2 Editing Column Names

nQuery provides the user the ability to change a column name. To change a column name, double-click the number in the column title bar above the column cells. The number will appear in a white square and will now be editable. Replace the number with the desired column name and this name will then be displayed in the design table. An example of a design table with changed column names was shown in Figure 2.2.

3 Advanced nQuery Assistants

3.1 Design Table Side-Tables

Side-tables provide table specific tools which allow users to calculate design parameters based on additional information which is commonly known or used in a particular design type. This section will go over the three primary types of side-tables: Effect Size Side-Tables, Covariance Matrix Side-Tables and Mandatory Side-Tables

3.1.1 Compute Effect Size Side-Table

For an introduction in how to open and use effect size side-tables, see subsection 1.7.4. In this section, we will briefly summarise the process of opening the side-table, using the side-table and cover some additional information not included in subsection 1.7.4.

3.1.1.1 Opening an Effect Size Side-Table

There are three main routes to open an effect size side-table: using the "Compute Effect Size" option in the Assistants menu, using the Compute Effect Size button  in the menu bar or by selecting the cell(s) in design table which are transferred from that side-table. Column cells which are the transfer target for an Effect Size side-table will be indicated in the cell's Help Card by its "Aid" section.

When a side-table is opened for a column, it will appear below the main analysis table in the same window as the Output and Specify Multiple Factors tools.

For each column in which a side-table is opened, a unique side-table will open associated with that selected column. The names of these side-tables are found in the tab bar at the bottom of the side-table window. The template for a side-table name is [Table Code]S-[Column Number]. For example, opening a side-table in MGT0 (One Way Analysis of Variance) for column one will open a side-table named "MGT0S-1".

Note that users can edit and transfer from other column side-tables into that column by selecting them in the tab menu at the bottom of the side-table window.

3.1.1.2 Using an Effect Size Side-Table

For a specific example of an effect size being used, see subsection 1.7.4. The format for an effect size side-table will be unique depending on the design table being used. For guidance on using a specific side-table refer to the help materials, in particular the Home Card and Help Card "Aid" sections of relevant cells.

In general, the effect size side-table works by the user fully specifying the required cells in the table. When the minimum amount of information required in a side-table is completed, the "Compute" button in the upper-left of the side-table window will become active (ungreyed). Selecting the Compute button will calculate the design parameter(s) which will be transferred to the main table and other parameters of interest in certain design tables. After the Compute button is used, the "Transfer" button (right of the Compute button) will become active and when selected this will transfer the relevant design parameters into the associated column in the main design table.

To return the side-table to its original state, use the "Clear" button (to the right of the Transfer button) to remove all user entries, read-only entries and computed values from the side-table.

Using the right-click context menu gives access to same edit tools and shortcuts as the main table (see subsection 2.5.1).

An example of an effect size side-table is shown in Figure 3.1.

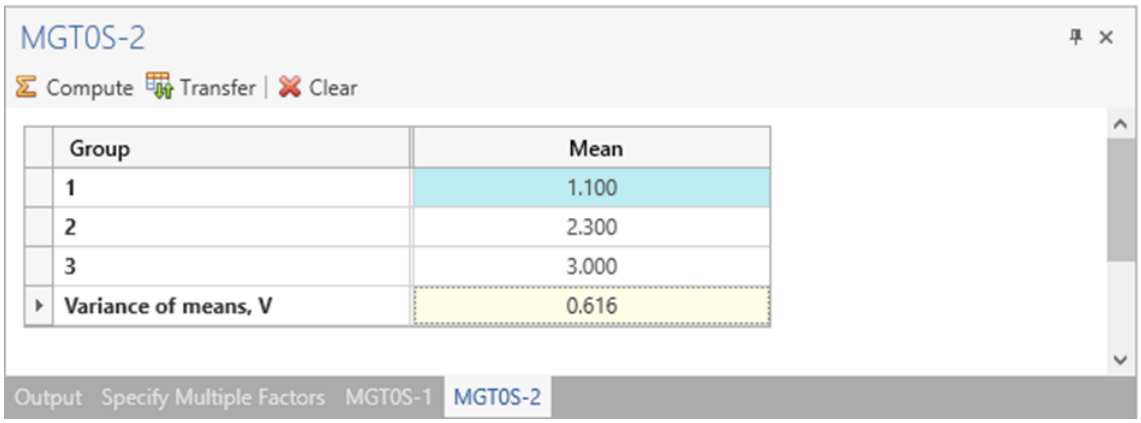

Figure 3.1: Effect Size Side-Table Example

3.1.1.3 Closing an Effect Size Side-Table

To close a specific effect size side-table, use the "x" in the top-right of the effect size side-table window. You can also hide the side-table using the pin symbol in the top-right or by right-clicking on the tab menu name.

3.1.2 Specify Covariance Matrix Side-Table

The Specify Covariance Matrix side-table provides a convenient method to calculate the error and Greenhouse-Geisser correction terms for the Repeated Measures ANOVA tables which use the Greenhouse-Geisser approximation in nQuery. These are MOT4 and MTT3.

3.1.2.1 Opening an Effect Size Side-Table

There are three main routes to open an Specify Covariance Matrix side-table: using the "Specify Covariance Matrix" option in the Assistants menu, using the Specify Covariance Matrix button  in the menu bar or by selecting the cell(s) in design table which are transferred from the side-table. These are the "error-term", sphericity and bias term rows.This will be indicated in the relevant cell's Help Card by its "Aid" section.

As per the effect size side-tables, when a covariance matrix side-table is opened for a column, it will appear below the main analysis table in the same window as the Output and Specify Multiple Factors tools.

For each column in which a side-table is opened, a unique covariance matrix side-table will open associated with that selected column. The names of these side-tables are found in the tab bar at the bottom of the side-table window. The template for a side-table name is [Table Code]C-[Column Number]. For example, opening a covariance matrix side-table in MOT4 (One Way Repeated Measures ANOVA (Greenhouse Geisser Approximation) for column one will open a side-table named "MOT4C-1".

Note that users can edit and transfer from other column side-tables (effect size or covariance matrix) into that column by selecting them in the tab menu at the bottom of the side-table window.

3.1.2.2 Using a Specify Covariance Matrix Side-Table

There are two modes for filling the Specify Covariance Matrix side-table: the Specify Standard Deviations and Correlations mode and the Specify Full Covariance Matrix mode. To switch between the modes, use arrow on the right of the mode drop-down to the right of the "Transfer" button at the top of the side-table window. This drop-down is shown in Figure 3.2.

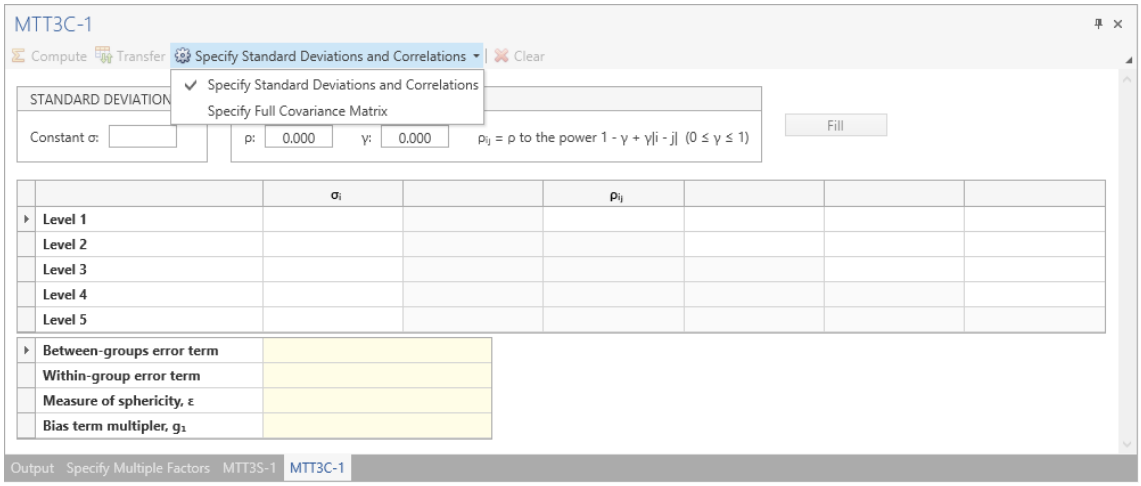

Figure 3.2: Specify Covariance Matrix Mode Drop-down Example

Specify Standard Deviations and Correlations Mode The Specify Standard Deviations and Correlations mode has two primary methods to activate the side-table: using the Constant Standard Deviations and Correlation Pattern shortcuts or directly filling the σ_i and ρ_{ij} table entries.

Note that changes made in one mode will affect the table values given in the other mode.

1) Constant Standard Deviations and Correlation Pattern

To use the Constant Standard Deviations and Correlation Patterns shortcuts you need to fill three inputs: the common standard deviation σ, the correlation ρ and the pattern term γ.

The common standard deviation value will fill σ_i in the table below with the same common standard deviation values. This value must be greater than zero.

The correlation will be the value included in the first diagonal of the ρ_{ij} table below and will adjusted downwards by the pattern term as you rightwards in the ρ_{ij} table. This value must be between 0 and 1

The pattern term defines the adjustment applied to the correlation as you move right-wards within the ρ_{ij} table. The formula for this adjustment is given the table and is $1 - \gamma + \gamma|i - j|$ where $|i - j|$ is the absolute distance between the first cell in the row in the ρ_{ij} and the cell of interest. This value must be between 0 and 1.

The common standard deviation will assign a common standard for the measurements at each measurement level and the correlation and pattern terms will represent the correlation between measurements taken on a subject at a given "measurement distance". For example, the 2nd column cell in row 1 of the ρ_{ij} table represents the average correlation between the 1st and 3rd measurement for a subject.

To transfer the relevant σ_i and ρ_{ij} values for the specified values given in the Shortcuts menu, select the "Fill" button to the right of the shortcuts menus. This will activate the Compute button which will calculate the relevant error-terms, bias and sphericity and activate the Transfer button. Selecting Transfer will place these values in the associated main design table column.

To use a compound symmetry (CS) model, set the constant correlation to the desired value and set the pattern term to zero. In the CS matrix, the correlation between measurements is assumed to be constant regardless of the distance between measurements.

To use an autoregressive(1) (AR(1)) model, set the correlation at the desired value (this will be the "distance 1" correlation) and set the pattern term to one. In the AR(1) matrix, the correlation between measurements is assumed to decay via an AR(1) process as the distance between measurement increases.

To achieve a pattern where the correlation does decay but slower than the AR(1) model, set the correlation at the desired value (this will be the "distance 1" correlation) and set the pattern term to between zero and one.

An example of using the AR(1) model computed using the shortcuts method is shown in Figure 3.3.

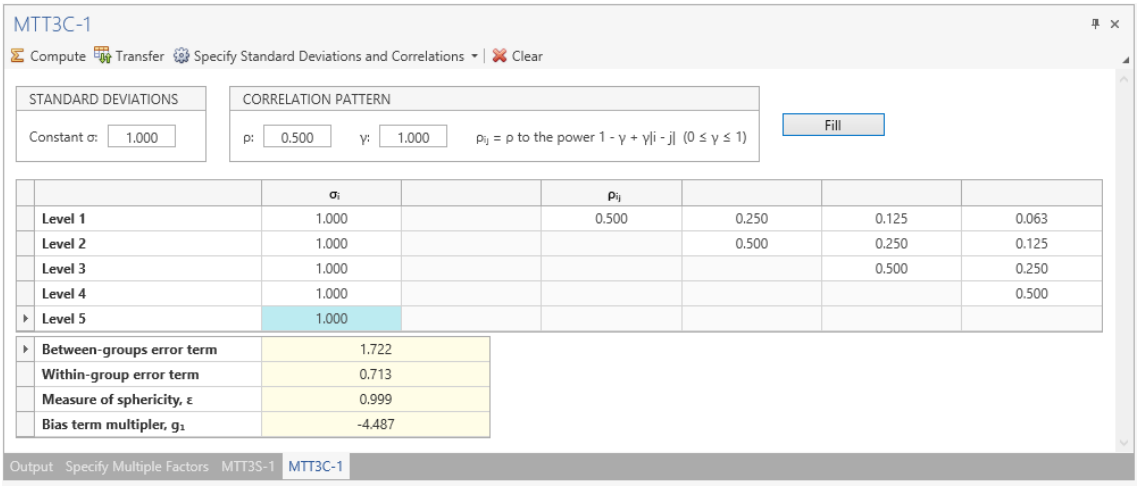

Figure 3.3: Specify Covariance Matrix SD + Correlation Example

2) Directly Filling Table

In addition to using the shortcuts method, you can fill the σ_i and ρ_{ij} in the table directly. This can be done with a blank table or the values can be edited from a shortcuts generated template. This gives the user a large amount of flexibility to set the table to their specifications. Once all σ_i and ρ_{ij} cells are filled, this will activate the Compute button which will calculate the relevant error-terms, bias

and sphericity and activate the Transfer button. Selecting Transfer will place these values in the associated main design table column.

An example where we edited from the Figure 3.3 template to have the individual level standard deviations vary and the distance 4 (i.e. between measurement 1 and 5) correlation equal to zero is shown in Figure 3.4.

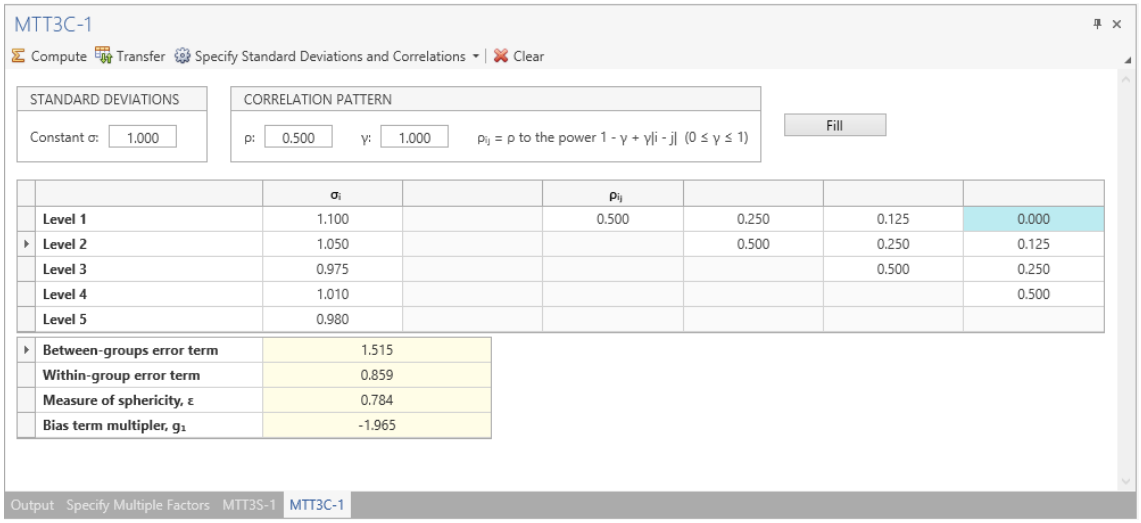

Figure 3.4: Specify Covariance Matrix Manual SD + Correlation Example

Specify Full Covariance Matrix mode The Specify Full Covariance mode works similarly to directly filling the standard deviations and correlations table. However, the covariance matrix is entered directly in this case. In the covariance matrix, the main diagonal elements are equal to within-measurement level variance and off-diagonal elements represent the covariance between main diagonal element level and the measurement at that distance (e.g. the third column entry in row one equals the average covariance between the first level subject measurement and the third level subject measurement).

There are two main methods to fill the covariance matrix: automatically fill using Constant Standard Deviations and Correlation Pattern mode or manually fill.

As noted at the start of this section, if the standard deviations and correlations table has been filled (either using the shortcuts menu or manually) in the other mode then the Covariance Matrix will be filled in this mode. Alternatively, you can fill or edit the covariance matrix manually to the study's specifications. Once all elements are filled in the covariance matrix, this will activate the Compute button which will calculate the relevant error-terms, bias and sphericity and activate the Transfer button. Selecting Transfer will place these values in the associated main design table column.

An example of an Specify Covariance Matrix side-table computed using the Specify Full Covariance Matrix mode is shown in Figure 3.5.

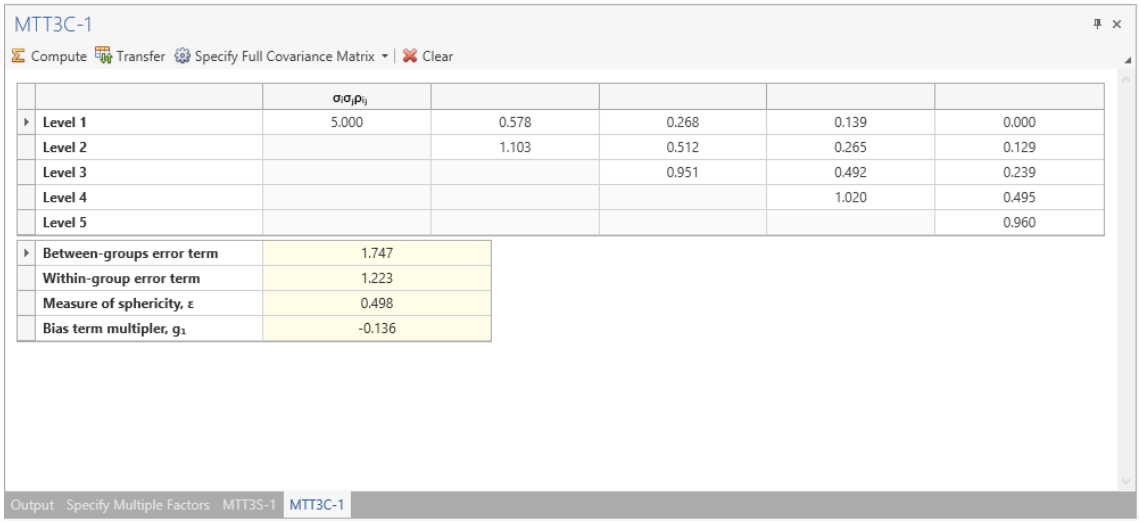

Figure 3.5: Specify Covariance Matrix Matrix Example

3.1.2.3 Closing a Covariance Matrix Side-Table

To close a specific covariance matrix side-table, use the "x" in the top-right of the effect size side-table window. You can also hide the side-table using the pin symbol in the top-right or by right-clicking on the tab menu name.

3.1.3 Mandatory Side-Tables

In the majority of nQuery design tables which use side-tables, the side-table is optional. In these cases, the role of the side-table is to provide a method for users to derive required design parameters based on other information they may have access to.

In a minority of nQuery design tables, the side-table is mandatory. In these cases, the inputs in the side-table(s) are used directly by the solvers of that table. The following is a list of the tables which have mandatory side-table(s):

- MTT2 - Wilcoxon/Mann-Whitney Rank-Sum Test (Ordered Categories)
- MTT2U - Wilcoxon/Mann-Whitney Rank-Sum Test (Ordered Categories) - Unequal n's
- PTT4 - Mantel-Haenszel (Cochran) Test of OR=1 in S Strata
- PTT4cc - Mantel-Haenszel (Cochran) Test of OR=1 in S Strata (Continuity Corrected)
- STT3 - Log-Rank Test, User Specified Survival Rates, Accrual, Dropouts (Simulation)

- STT3U - Log-Rank Test, User Specified Survival Rates, Accrual, Dropouts (Simulation) - Unequal n's
- MTT12 - Group Sequential Test of Two Means
- MTT40 - Group Sequential Test of Two Poisson Rates
- MTT42 - Group Sequential Test of Two Negative Binomial Rates
- PTT12 - Group Sequential Test of Two Proportions
- STT12 - Group Sequential Test of Two Survival Curves
- STT15 - Group Sequential Test of Two Survival Curves with Accrual
- MOT26 - Group Sequential Test for One Mean
- POT8 - Group Sequential Test for One Proportion (Alternative Variance)
- POT13 - Group Sequential Test for One Proportion (Null Variance)
- MGT3 - Multivariate Analysis of Variance (MANOVA)
- MGT5 - Multiple Comparisons Procedure - Modelling (MCP-Mod) - Discussed in chapter 5
- ROT7 - Poisson Regression *(Rate Ratio solvers only)*
- MOT25 - Bayesian Assurance for One Mean (Custom Prior)
- POT7 - Bayesian Assurance for One Proportion (Custom Prior)
- MTT21 - Bayesian Assurance for Two Means (Custom Prior)
- PTT13 - Bayesian Assurance for Two Proportions (Custom Prior)
- STT14 - Bayesian Assurance for Two Survival Curves (Custom Prior)

For all except MGT3 and the Group Sequential tables, use the same instructions as for the Effect Size side-tables above to open the side-table. In each of these design tables except MGT3 and the group sequential tables, there is a row named "Side Table Name" which will also open the side-table for a column automatically (assuming all side-table opening conditions are met).

In MGT3, there are three separate side-tables per column and these are opened by selecting the Factor Level Table, Means Matrix and Covariance Matrix read-only cells in a column. In the group sequential tables, selecting any cell in a column will open the Looks side-table.

The side-tables are the same between MTT2/MTT2U, PTT4/PTT4cc, STT3/STT3U, the Bayesian Assurance custom prior tables (MOT25/POT7 etc.) and the group sequential tables (MTT12/PTT12 etc.). These side-tables operate the same as optional side-tables except they do not have Compute/Transfer buttons. For mandatory side-tables when the table is filled sufficiently to be used in a solver, the name of the side-table in "Side Table Name" row (or MGT3 equivalent) will go from grey to black. The table otherwise will operate as per any other design table.

Note that Looks side-table for the group sequential tables does not require any editing to be active and is by default read-only. It is only required to be edited if the Information Times or Spending Function rows are set to "User Specified". See the help for the group sequential details for further details.

If a side-table has been completed in any column, you can automatically use that side-table in another column by selecting the relevant side-table name from the drop-down options in that column's "Side Table Name" row.

An example of a mandatory side-table being used in MTT2 is shown in Figure 3.6.

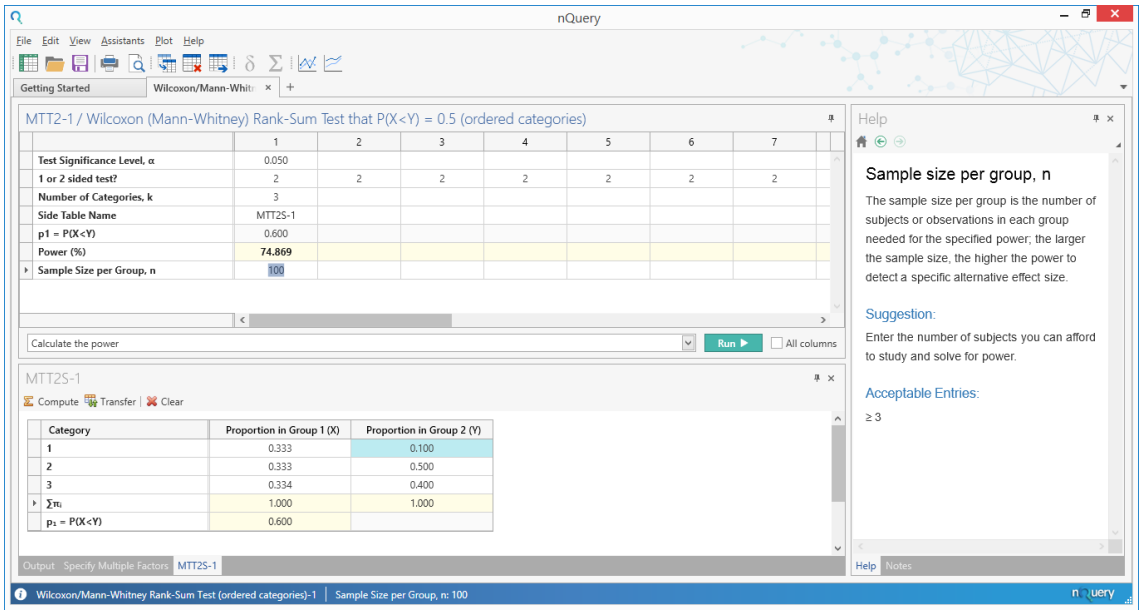

Figure 3.6: MTT2 Mandatory Side-Table Example

3.2 Compute Standard Deviation Assistants

nQuery provides a number of assistants which allow a user to derive the standard deviation for a study based on other parameters or data which the user has.

To access these assistants, select the "Standard Deviation" option from the Assistants file menu. This will open a radio button list menu containing the standard deviation assistants. Select an assistant radio button and select OK to open the relevant assistant.

There are 10 standard deviations (SD) assistants available in nQuery. These are (with a brief description):

1. From Standard Error: Derive SD from sample size (n) and standard error
2. For SD1 and SD2 (pooled SD): Derive pooled SD from group 1 and 2 standard deviations (SD1/SD2) and sample sizes (n1/n2)

3. From Range: Derive SD from range (using maximum and minimum optionally to derive range) and sample size, n
4. From Percentile: Derive the SD from the percentile and the difference in upper and lower percentile values (using upper and lower values optionally).
5. From Coefficient of Variation: Derive the log-scale SD and mean from the coefficient of variation and observed mean respectively
6. From Upper Confidence Limit: Derive the upper limit of the SD from the sample size, observed standard deviation and confidence level.
7. From SD1, SD2, Correlation: Derive the SD for the differences from the first condition SD, the second condition SD and the correlation
8. For Cluster Sampling: Derive the "cluster" SD from the between-cluster and within-cluster variances, the intra-cluster correlation and subjects per cluster.
9. For specified x values: Derive the SD from a set of x data values. See section 3.3
10. Of residuals: Derive the SD of the residuals and the dependent (Y) variable from the regression coefficient, correlation coefficient and independent (X) variable SD.

All of these assistants work in the same or similar way as nQuery design tables except for the "For specified x values" option. Use the guide given in chapter 1 and the individual Home and help cards provided for guidance on how to use each table.

Options 1 - 4 provide alternative methods to derive a standard deviation for one or two sample independent means test and Option 6 provides the alternative estimate for the standard deviation by using the upper confidence limit . Option 5, 7, 8 and 10 provide help to derive the adjusted standard deviation for specific designs or testing methods. Option 5 would be used if the analysis will be of the log-transformed data, Option 7 would be used for paired means testing or analysis, Option 8 would be used if the study will use cluster randomisation and Option 10 is used if the analysis is using linear regression.

3.3 Data Entry Assistant

nQuery provides a data entry assistant to allow users to derive commonly used design parameters based on data which may be available to the user.

The Data Entry Assistant has three main components: the Data Entry table, the Transformation menu and the Insert Values menu.

3.3.1 Data Entry Table

The data entry table is the main table on the left-hand side of the Data Entry Assistant. It consists of 20 columns and 500 rows. Each column corresponds to a data set and each row corresponds to an individual subject's data value in that column's data set.

As a user fills a column, four commonly used design parameters will be derived for the data set. These are the sample size (N), the mean, the sample standard deviation and population standard deviation, $\sigma(x)$. The difference between the sample and populations standard deviations is the denominator in the estimate for each is N-1 and N respectively.

3.3.2 Transformations Menu

The transformations menu gives the option to create a new column containing the transformed values for an entered data set. Transformations are commonly used when data is non-normal but has a distribution which would be amenable to being transformed into a normally distributed data using a common transformation.

The Transformations menu is found to the right of the Data Entry table and can be opened by selecting the Transformations tab at the top of the menu window.

The three transformations are a square root transformation (Sqrt(X)), a base 10 log transformation (log10(X)) and the natural log transformation (ln(X)). In addition, you can conduct these transformations to the data with a fixed value (A) added to each data value before transformation. In this case, the user is required to enter a value of A in the menu.

To use the transformations menu, select the relevant column in the data entry table, select the transformation required, enter a value for A if a (X+A) transformation is selected and select "Add".

This will generate the transformed data (with summary statistics) in the next empty column to the right of the selected data entry column. The column title will include the name of the transformation and the column transformed and any transformed columns will be highlighted in orange.

3.3.3 Insert Values Menu

The insert values menu gives the option to create a number column using an evenly spaced number of data values in a specified range. This can be useful to generate a large dataset easily.

The insert values menu is found to the right of the Data Entry table and can be opened by selecting the Insert Values tab at the top of the menu window.

To generate a column using the insert values menu, specify a From value and To value in the relevant cells. These correspond to the minimum and maximum values of the range of interest respectively.

The user can then activate either the Specify Increment or Specify num. values option by selecting the radio button to the left of each option.

If the Specify Increment option is activated, the entered value will be increments in which the data will increase going from the "From" to the "To" value. For example, if From = 0 and To = 10 then Specify Increment = 2 will give the five values of [0, 2, 4, 6, 8].

If the Specify num. values option is activated, the entered value will be the number of data values which will be taken between the "From" and "To" value. For example, if From = 0 and To = 10 then Specify num. values = 5 will give the five values of [0, 2.5, 5, 7.5, 10].

Once the required values are entered for "From", "To" and "Specify Increment"/"Specify num. values", select "Execute" and that data will be generated on the currently selected Data Entry Table cell.

An example of Data Entry table using the above options is shown in Figure 3.7.

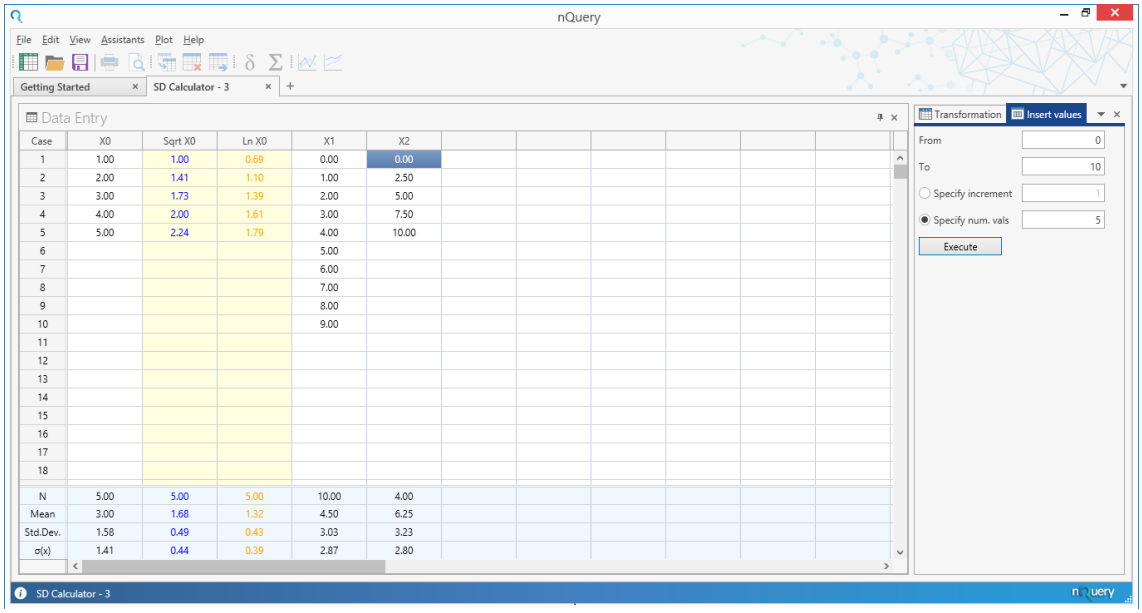

Figure 3.7: Data Entry Assistant Example

In column 1 is a user entered set of values (X0), in column 2 is Sqrt(X) transformed data for column 1, in column 3 is ln(X+1) (i.e. A=1 with natural log transform) transformed data for column, in column 4 is Insert Values menu data generated with From = 0, To = 10, Specify Increment = 1 and in column 5 is Insert Values menu data generated with From = 0, To = 10, Specify num. values = 5.

3.4 Cumulative Distribution Function Assistants

nQuery provides access to the most commonly used cumulative and inverse cumulative distribution functions used by the nQuery solver algorithms. To open a table for a specific statistical distribution, select Distribution Function from the Assistants file menu. This will open a radio button menu of the available statistical distributions. To open a distribution table, select the relevant radio button to the left of the distribution name and select "OK".

There are distribution assistants provided for the following statistical distributions:

1. z (Gaussian/Normal)
2. t (Central)
3. t (Non-central)
4. χ^2(Central) *(Chi-Square)*
5. χ^2(Non-central) *(Chi-Square)*
6. F (Central)
7. F (Non-central)
8. Cumulative binomial

The distribution tables work the same as standard nQuery design tables with the solvers being for the cumulative probability of P(X<x_p) (all tables), the percentile test statistic (e.g. F_p, t_p) (all tables except cumulative binomial) and the non-centrality parameter (non-central tables only).

3.5 Survival Parameter Conversion Assistant

The survival parameter conversion assistant provides a tool to convert between the three most commonly used survival parameters for an exponential survival curve. These are the proportion surviving at time t, the median survival and the exponential parameter.

In the survival parameter conversion assistant, the three survival parameters can be specified for a two group design individually for each group. In addition, there are rows for the Time t and the Hazard ratio. The Time t is required to calculate the proportion surviving at time t from the median survival or exponential parameter and vice-versa. The hazard ratio is calculated after the median survival/exponential parameter are specified in each group.

Note that the exponential parameter will automatically calculate the median survival when specified and vice-versa in each group.

Many nQuery survival tables contain an Effect Size side-table based on this assistant so this assistant is primarily for exploration purposes.

3.6 Bayesian Converter

4 Group Sequential Designs, Interim Monitoring & Unblinded Sample Size Re-estimation

4.1 Introduction

nQuery is designed for the calculation of power and sample size for both fixed term and group sequential design studies. nQuery will have options available for all users for planning a group sequential design for two means, two proportions and two survival curves (assuming equal follow-up).

nQuery Adapt users will have access to additional group sequential designs and an interim monitoring and unblinded sample size re-estimation tool for the most commonly used group sequential design. This section will cover the basic theory behind these concepts and will give detailed description of how to design a group sequential trial in nQuery and use the nQuery Adapt tool for interim monitoring and unblinded sample size re-estimation.

4.2 Group Sequential Test Design

4.2.1 Background

Group sequential designs are an extension of fixed period designs in which data from the trial is analyzed at one or more stages prior to the conclusion of the trial. The trial can then be stopped early if there is strong evidence for or against the proposed treatment being effective based on data up to that point. Group sequential designs are one of the most widely used types of adaptive trial and provide the opportunity to stop a trial early for efficacy (strong interim evidence against the null hypothesis) and/or futility (strong interim evidence for the null hypothesis) and thus significantly reduce the economic and ethical costs over the equivalent fixed period design.

4.2.2 Group Sequential Design Theory

In nQuery, the group sequential design tables' power and sample size calculations are performed using the Lan-DeMets alpha spending function approach [Demets and Lan, 1984, Demets and Lan, 1994]; for estimating boundary values. Building on the work of Lan and DeMets; Pampallona, Tsiatis, and Kim [Pampallona et al., 1995, Pampallona et al., 2001] later put forward the concept of using a beta spending approach to construct boundaries for futility where the evidence for the null is strong. These boundary values indicate the interim test values (test statistic, effect size, p-value) which would lead to the trial stopping early based on interim data. Boundary values can be estimated in a number of ways with nQuery providing bounds based on the O'Brien-Fleming [O'Brien and Fleming, 1979], Pocock [Pocock, 1977], Hwang-Shih-DeCani [Hwang et al., 1990] and Power Family spending functions. Each of these spending functions spend a certain proportion of the alpha error (Type I/efficacy) and/or beta error (Type II/futility) at each analysis or 'look' and then make the needed adjustments to the sample size and final errors to preserve the overall Type I and Type-2 errors. The "spent" alpha and beta values used at each look are calculated based upon the test hypothesis, the spending function chosen, the number of looks to be taken during the course of the study as well as the overall Type I and Type-2 error rates. For a full introduction to group sequential methods, we recommend *Group Sequential Methods Applications to Clinical Trials* by Jennison & Turnbull (2000) [Jennison and Turnbull, 1999].

4.2.2.1 Spending Functions

There are four spending functions available to the user in nQuery for the efficacy and futility bounds. Note that nQuery also provides the option to manually input boundary values. As standard all alpha spending functions have the properties that the error spent at the start of the trial equals zero and the error spent at the final analysis equals the original desired error level i.e. the desired test significance level for alpha spending, one minus the power (as a proportion) for beta spending. Functionally the alpha and beta spending functions are the same.

The spending functions for alpha spending are summarised in Table 4.1 where $\alpha(\tau)$ is the cumulative alpha spent at the specified look, α is overall alpha error, τ is the information time (usually the sample size up to that look as a proportion of the total sample size), $z_{1\text{-}\alpha/2}$ is the inverse standard normal cumulative distribution assessed at "1-α/2", Φ is the standard normal cumulative distribution function, ρ is the Power Family parameter and γ is the Hwang-Shih-DeCani gamma parameter.

Table 4.1: Spending Functions

Spending Function	Form
O'Brien-Fleming	$\alpha(\tau) = 2\left(1 - \Phi\left(\frac{z_{1-\alpha/2}}{\sqrt{\tau}}\right)\right)$
Pocock	$\alpha(\tau) = \alpha ln(1 + (e-1)\tau)$
Power Family	$\alpha(\tau) = \alpha\tau^{\rho}$
Hwang-Shih-DeCani	$\alpha(\tau) = \alpha\left[\frac{\left(1-e^{-\gamma\tau}\right)}{(1-e^{-\gamma})}\right]$

Most spending functions spend less error at the earlier looks, with the O'Brien-Fleming spending function being more conservative than the Pocock spending function (Power Family/Hwang-Shih-DeCani will depend on their free parameter). This is usually a desired characteristic as it means that the results of any interim analysis will only be considered significant (and thus ending the trial) at an early stage with an extreme result. It also means that the final analysis will be more comparable in terms of sample size and significance boundaries to an equivalent fixed term design.

4.2.2.2 Boundaries

The boundaries in nQuery represent the critical values at each look above or below which the trial would end early. These boundaries are usually constructed using the alpha and beta spending functions, though users are given the option of entering these manually. For the spending function approach, nQuery will automatically generate boundaries for the early rejection of the null hypothesis (if an efficacy alpha spending function is active), early finding for the null hypothesis (if a futility beta spending function is active) or both (if futility and efficacy are both active using a combination of both the alpha and beta spending functions).

Once these critical boundary statistics are generated, they can be compared during interim monitoring to the interim test statistics to decide whether to end the trial early. Essentially, if a test statistic crosses an efficacy boundary then it can be concluded that the experimental treatment shows a statistically significant effect and the trial can be stopped with rejection of the null hypothesis. If the test statistic crosses a futility boundary then this indicates with high probability that an effect will not be found, that the trial can be terminated by rejecting the alternative hypothesis. For futility bounds there are two options; either to have the boundaries binding, or non-binding. With binding boundaries, if the test statistic crosses the futility boundary, the test must be stopped, otherwise the type-1 error may become inflated. The reason for this is that there is an interaction between the efficacy and futility boundaries in their calculation that could cause the efficacy boundary to shift. In the case of non-binding boundaries; the efficacy boundaries are calculated as normal, that is, as if the futility boundaries did not exist. This eliminates the danger of inflating the Type I error if the futility boundary is overruled. The downside of

the non-binding case is that it may increase the required sample size relative to the binding case.

In nQuery, boundary values are given on the standardized scale (i.e. Z-statistic scale) and will usually equal the treatment difference divided by its standard error. Examples for two independent sample design case are given in Table 4.2.2.2.

Table 4.2: Two Sample Group Sequential Standardized Statistics

Means	Proportions (Pooled)	Survival (Approx.)
$Z = \frac{\mu_2 - \mu_1}{\sqrt{\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}}}$	$Z = \frac{\pi_2 - \pi_1}{\sqrt{\frac{\bar{\pi}(1-\bar{\pi})}{n_1} + \frac{\bar{\pi}(1-\bar{\pi})}{n_2}}}$	$Z = \frac{ln(HR)}{\sqrt{\frac{1}{E_1} + \frac{1}{E_2}}}$

For early stopping, we simply calculate the Z-statistic based on the current interim data and stop early if it is above the "upper" bound or below the "lower" (2-sided) or "futility" (1-sided) bound. Note that in nQuery; "upper", "lower" and "futility" are defined relative to the direction of the pre-specified effect size. For example, if the initial design assumes a negative difference ($\mu_1 > \mu_2$) then positive "upper" interim tests statistics would be those in which the interim difference is also negative and vice-versa. Similarly, "lower" results would be those in the opposite direction. Futility bounds can be above or below zero but the same guidance applies for the Z-statistic calculation.

4.2.2.3 Calculating Power or Sample Size

To calculate the power or sample size, we can use the drift parameter. This is the standardized Z-statistic defined above for the relevant test for the final look in the group sequential design. It also equals $\triangle\sqrt{I_x}$ where $\triangle$is the treatment effect (e.g. mean difference) and I_x is the maximum total information of the group sequential design.

To calculate the sample size, nQuery calculates the value for the drift parameter based on the values for the significance level, whether a one-sided or two-sided analysis is being used, the power, the number of looks and the spending function. The drift parameter is then set to equal to the standardized Z-statistic for the final look and via re-arrangement or iterative search, the sample size can be found.

To calculate the power, nQuery calculates the drift parameter by calculating the standardized Z-statistic based on the relevant information including the sample size. This drift is used to reverse-calculate the power via algorithm [Jennison and Turnbull, 1999] using the the significance level, whether a one-sided or two-sided analysis is being used, the number of looks and the spending function.

4.2.3 Group Sequential Design in nQuery

nQuery provides an intuitive interface to make planning a group sequential design and finding the appropriate sample size or power easy. This section will outline the steps to plan a group sequential trial in nQuery.

4.2.3.1 Background

In nQuery, the vast majority of user actions and interface options are the same as for fixed term designs. This section will focus on the major additional issues associated with group sequential design in nQuery. See previous chapters for information on the shared table elements.

In nQuery, the inputs required for group sequential design can be split into two categories: fixed term parameters, group sequential trial (GST) parameters. The fixed term parameters are those which would be required to calculate the sample size or power for the equivalent fixed term design and thus will depend on the statistical test and data type. The GST parameters are those required to define the group sequential design such as the spending functions, number of interim looks etc. and will be effectively the same across different designs and data types. We will cover how both of these are entered and how to use nQuery for group sequential design in the following example.

4.2.3.2 Group Sequential Design Example

Main Table The main table is used to enter the fixed term parameters including the significance level, effect size and power among others. These will be very similar or the same as those for an equivalent fixed term design. In this example, the group sequential design is for a two sample t-test and thus the effect size is characterised by the mean difference and per-group standard deviations.

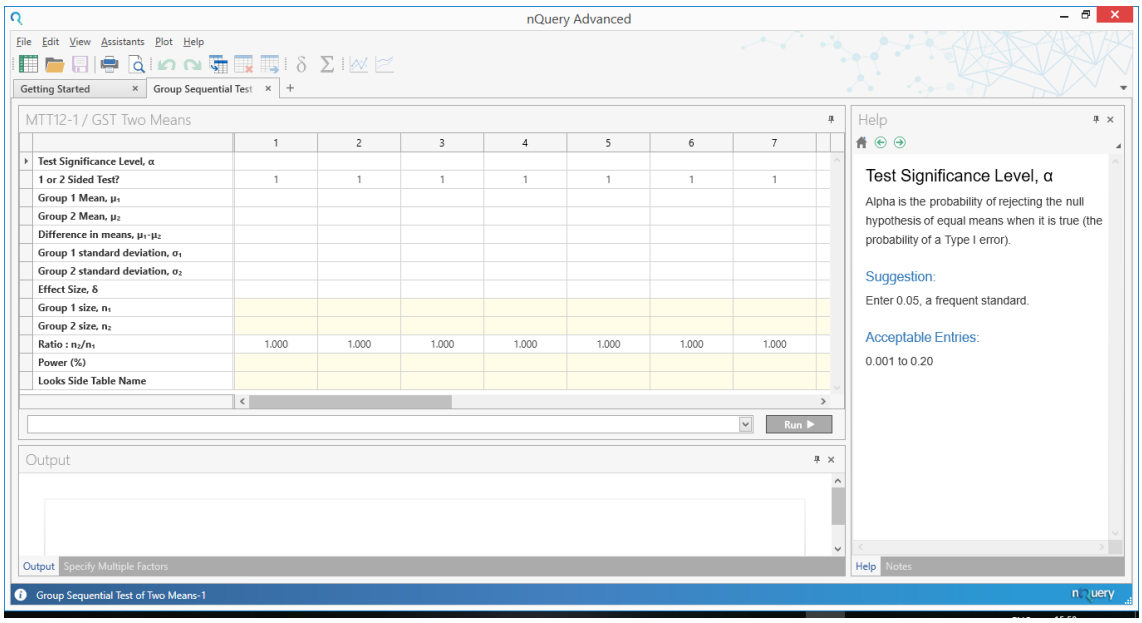

Figure 4.1: Two Means Group Sequential Design

The first thing we need to do is enter these fixed term design parameters. As these are effectively the same as for a fixed term design, please refer to previous chapters for guidance on filling these values. Here we will assume a treatment difference of -1 and per-group standard deviations of 2.5 for a 1-sided 5% significance test.

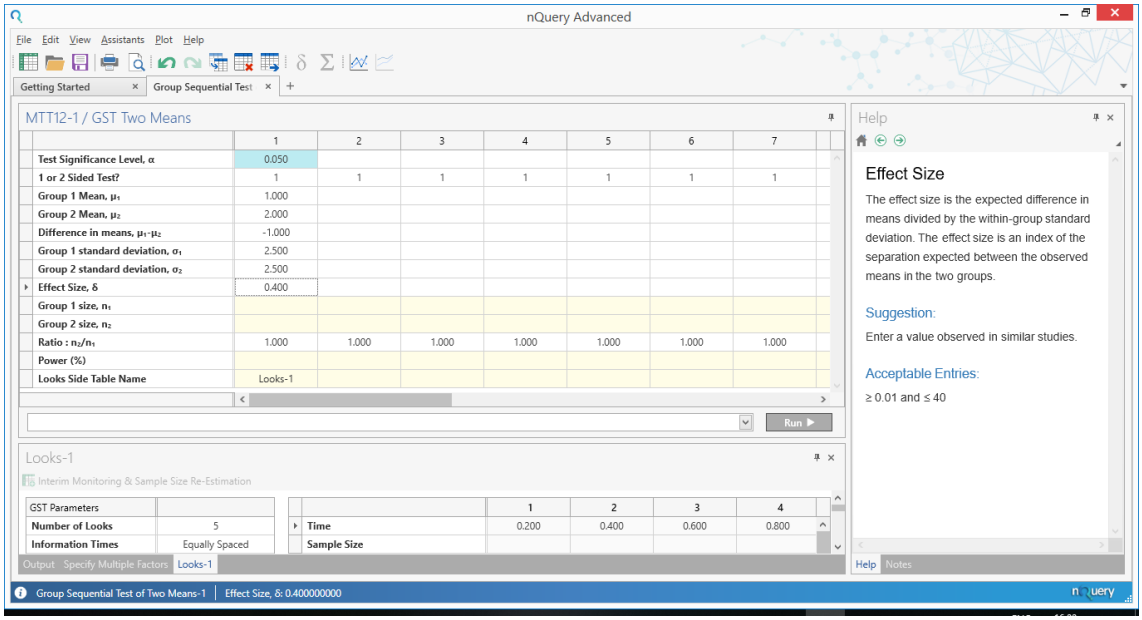

Figure 4.2: Example Fixed Term Parameters

We will leave the power and sample size empty for now as we want to define our

group sequential design before calculating the sample size. To do this we will need to use the Looks side-table.

Looks Side-table The Looks side-table will automatically open below the main table for a specific column when any cell in the column is selected. It can also be manually opened using the "Compute Effect Size" options from the toolbar or Assistants file menu. For more details on what side-tables are and how they work in general, refer to subsection 1.7.4.

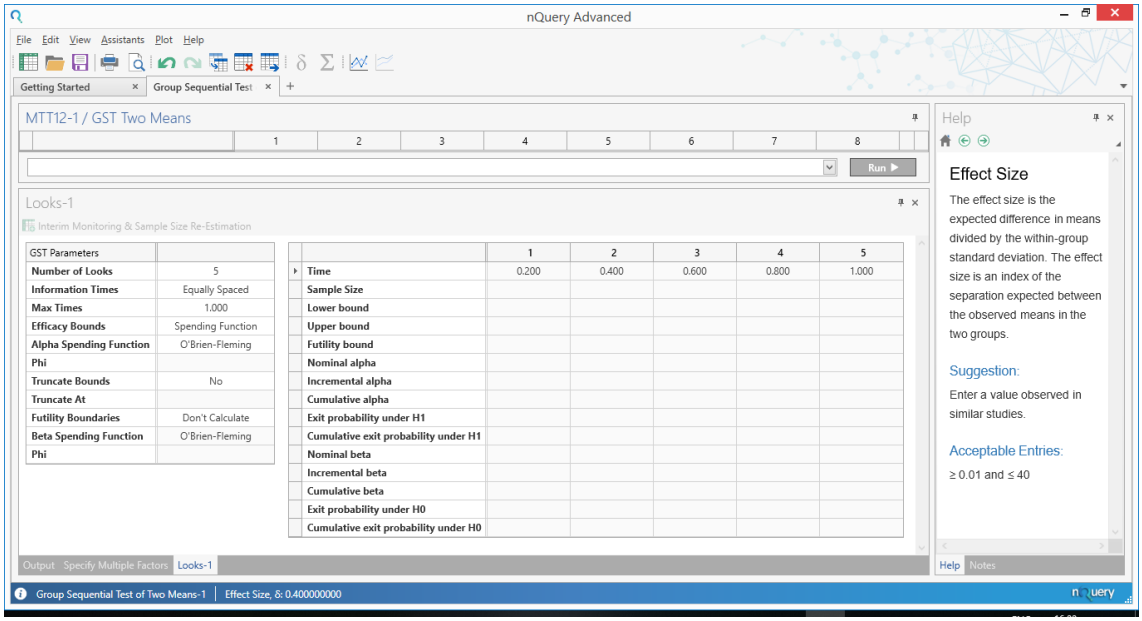

Figure 4.3: Looks Side-table

The Looks side-table can be split into two parts: the GST Parameters column (on the left) and the Looks table (on the right). The GST Parameters column defines the important features of the group sequential design. The Looks table will contain the important characteristics of the complete group sequential design as well as allowing some more complex study designs.

GST Parameters In the GST Parameters, there are 11 inputs settable to define the group sequential design operating characteristics. These can be summarised as follows:

1. Number of Looks (≥ 2): Set the total number of looks in group sequential design. As this is the total number of looks, the interim looks will equal this value minus one.

2. Information Times (Equally Spaced, User Input): Set whether the interim looks should be equally spaced or should be based on the user in-putted times (section 4.2.3.2)
3. Max Times (>0): Set the total maximum time. This will automatically re-scale the Time row in the Looks table when set. Note this is disabled when Information Times is set to User Input.
4. Efficacy Bounds (Spending Function, User Input, Don't Calculate): Set how the efficacy bounds will be calculated. Spending function will activate the Alpha Spending Function option, User Input will require editing the Upper Bound and Lower Bound rows of the Looks table (section 4.2.3.2), Don't Calculate will have no efficacy bounds and will only be usable if Futility Bounds is not set to Don't Calculate.
5. Alpha Spending Function (O'Brien-Fleming, Pocock, Power Family, Hwang-Shih-DeCani): Set the spending function used to calculate the efficacy bounds. For the Power Family and Hwang-Shih-DeCani options, the user will need to enter a value for the parameter row below this.
6. Power/HSD Parameter (>0 if Power Family, <3 if Hwang-Shih-DeCani): The free parameter for the Power Family or Hwang-Shih-DeCani efficacy spending function. Sometimes given the generic term of "Phi" in nQuery. See Table 4.1 for details.
7. Truncate Bounds (No, Yes): Set whether you want the efficacy bounds to be truncated at a specific user-defined value. If Yes is selected the Truncate at row will be active below.
8. Truncate at (>1): If Truncate Bounds is set to Yes, this defines the value at which the efficacy bounds are truncated. If an upper bound value in the Looks table is higher than this value then it will be set to this value. For lower bounds, it will be set to the minus version of this value.
9. Futility Bounds (Non-binding, Binding, Don't Calculate): Set how the futility bounds will be calculated. Non-binding and binding will activate the Alpha Spending Function option, Don't Calculate will have no futility bounds and will only be usable if Efficacy Bounds is not set to Don't Calculate. Non-binding allows futility bound contravention without error inflation, while binding does not allow this.
10. Beta Spending Function (O'Brien-Fleming, Pocock, Power Family, Hwang-Shih-DeCani): Set the spending function used to calculate the futility bounds. For the Power Family and Hwang-Shih-DeCani options, the user will need to enter a value for the parameter row below this.
11. Power/HSD Parameter (>0 if Power Family, <3 if Hwang-Shih-DeCani): The free parameter for the Power Family or Hwang-Shih-DeCani futility spending function. Sometimes given the generic term of "Phi" in nQuery. See Table 4.1 for details.

More detail on these are provided in the help cards and in subsection 4.2.2. In nQuery, the default values are for 5 look design with an efficacy bound using an O'Brien Fleming spending function. In this example, we will set these to have a two look equally spaced design (i.e. one interim analysis at 50% of the total sample size), an O'Brien-Fleming spending function for the efficacy bound and a Power Family spending function with a parameter equal to 1 (which is similar to Pocock spending function) for the non-binding futility bound.

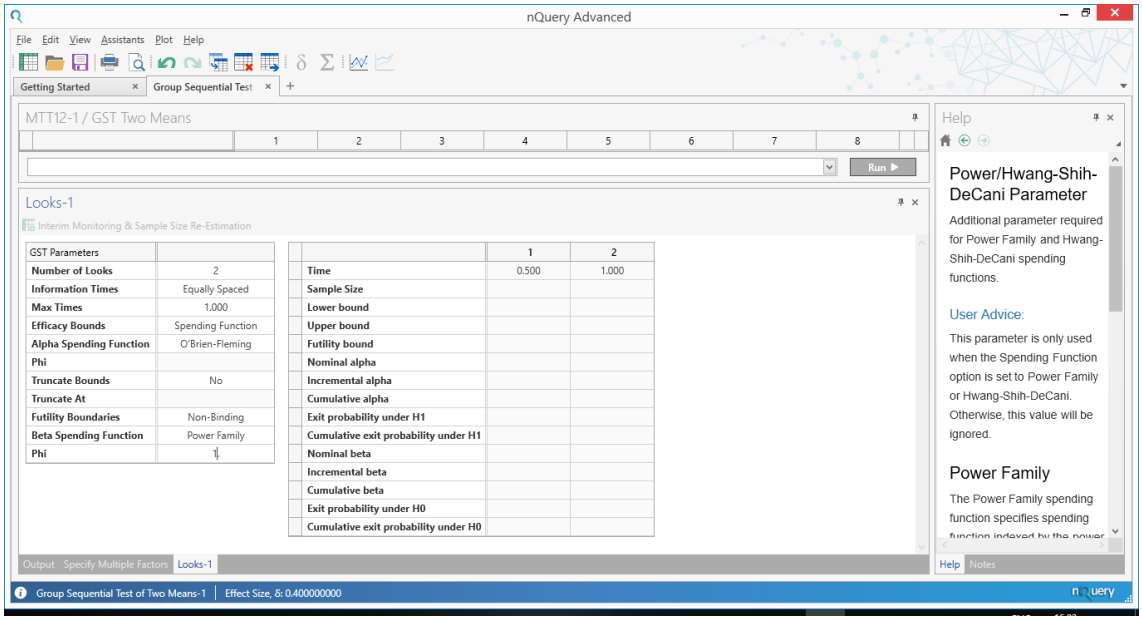

Figure 4.4: Example GST Parameters

Looks Table The Looks Table provides the information on the group sequential bounds and number of other useful parameters and additional information that characterises the complete group sequential design. In this table, each column corresponds to that Look. In this example, column 1 is the single interim look and column 2 is the final analysis. The Looks table contains the following rows:

1. Time: The Time at which each interim analysis will occur. By default, this will be the proportion of the total information at that look. The values in Time can be changed using the Max Times option (where each of these will be multiplied by the value in Max Times) or can be entered manually if Information Times is set to User Input
2. Sample Size: The cumulative sample size required at each look. The final entry will equal the calculated total sample size.
3. Lower Bound: The Lower Bound for efficacy below which we would end the trial early. The bounds are on the standardized test statistic scale. This will

only be used for 2-sided designs with an active efficacy bound. These will be entered manually if Efficacy Bounds is set to User Input.

4. Upper Bound: The Upper Bound for efficacy above which we would end the trial early. The bounds are on the standardized test statistic scale. This will be only used with an active efficacy bound. These will be entered manually if Efficacy Bounds is set to User Input.
5. Futility Bound: The Futility Bound below which we would end the trial early for futility (can ignore if non-binding). The bounds are on the standardized test statistic scale. This will only be used for 1-sided designs with an active futility bound.
6. Nominal Alpha: This is the value of alpha (i.e. p-value) for these boundaries if they were used in a single stand-alone test.
7. Incremental Alpha: This is the amount of alpha (type I error) that is spent at this interim test since the last look. It is close to nominal alpha but differs slightly due to being adjusted for multiple testing.
8. Cumulative Alpha: This is total amount of alpha (type I error) that has been spent up to and including this look.
9. Exit Probability under H1: This is the percentage chance that the trial will stop, given the specified alternative hypothesis in the main table is true. This is chance for both the efficacy and futility bounds.
10. Cumulative Exit Probability under H1: This is the cumulative percentage chance that the trial will stopped at this or any previous look, given the specified alternative hypothesis in the main table is true. This is chance for both the efficacy and futility bounds.
11. Nominal Beta: This is the value of beta for these boundaries if they were used in a single stand-alone test.
12. Incremental Beta: This is the amount of beta error (type II error) that is spent at this interim test since the last look. It is close to nominal beta but differs slightly due to being adjusted for multiple testing.
13. Cumulative Beta: This is total amount of beta error (type II error) that has been spent up to and including this look.
14. Exit Probability under H0: This is the percentage chance that the trial will stop, given the specified null hypothesis (usually difference of zero) in the main table is true. This is chance for both the efficacy and futility bounds.
15. Cumulative Exit Probability under H0: This is the cumulative percentage chance that the trial will stopped at this or any previous look, given the specified null hypothesis (usually difference of zero) in the main table is true. This is chance for both the efficacy and futility bounds.

More information on these can be found in the table help cards when these rows are selected or in subsection 4.2.2. Unless a User Input option is set, this table will be read-only and will fill automatically once a power or sample size calculation occurs in the main table. In this example, set power to 80% in the main table. This will give a sample size of 87 in each group and following Looks table (Figure 4.5)

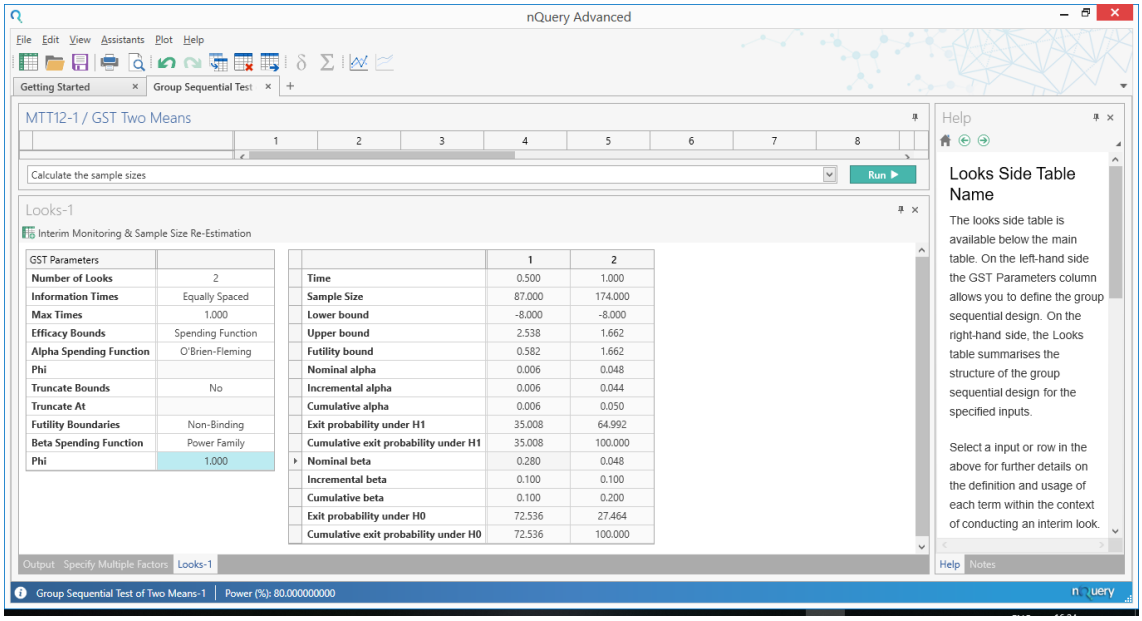

Figure 4.5: Example Looks Table

For practical purposes, the most important values are the Upper Bound, Futility Bound and Nominal Alpha. These are the values we will compare the interim test statistics (or p-values for nominal alpha) to decide whether to end the trial early or at the final analysis whether to reject the null hypothesis. The additional information provides useful context on the amount of error spent at each look and how likely the trial will end under the specified alternative hypothesis (i.e. the effect size in main table) or the null hypothesis.

As mentioned previously, certain elements of the Looks table will have to be set manually if Information Times or Efficacy Bounds are set to User Input. An example of these being set manually is given in Figure 4.6.

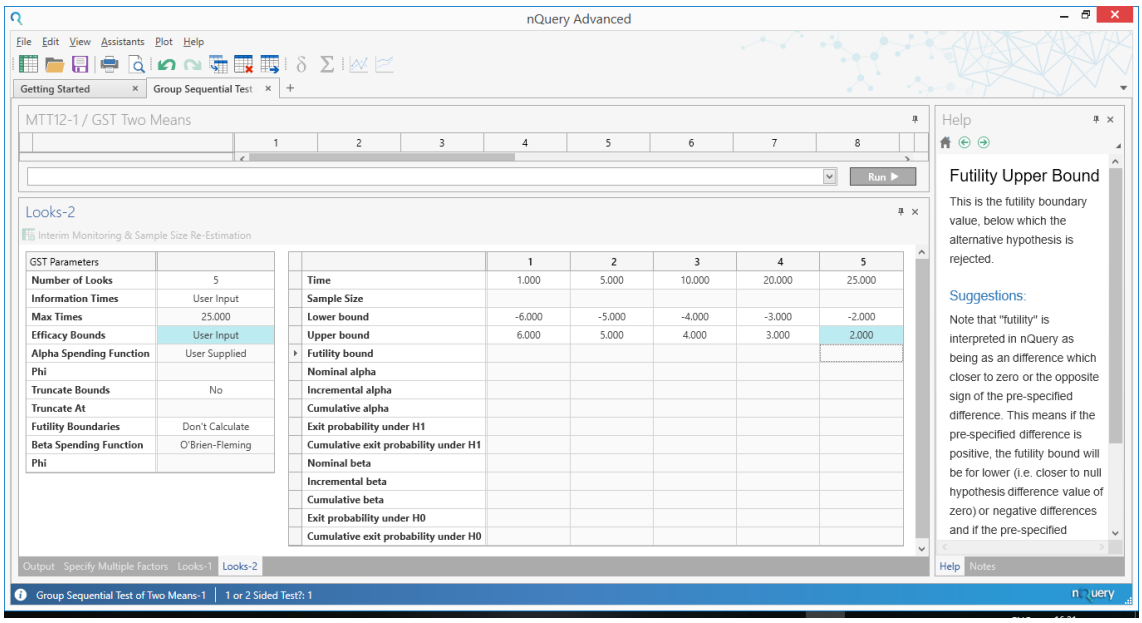

Figure 4.6: Custom Looks Table

4.3 Interim Monitoring and Unblinded Sample Size Re-estimation (nQuery Adapt)

4.3.1 Background

With nQuery Adapt, users will have access to the interim monitoring and sample size re-estimation tool. This tool will provide the opportunity to increase the sample size at an interim analysis if the interim information suggests a "promising" result but which is under-powered to find the current interim effect size if it were the true effect size. In group sequential designs and other similar designs, access to the interim data provides the opportunity to improve a study to better reflect the updated understanding of the study. One way to improve a group sequential design would be to use the interim effect size estimate not only to decide to whether to stop a trial early but to increase the sample size if the interim effect size is promising. This optionality gives the trialist the chance to power for a more optimistic effect size, thus reducing up-front costs, while still being confident of being able to find for a smaller but clinically relevant effect size by increasing sample size if needed.

4.3.2 Unblinded Sample Size Re-estimation Theory

nQuery provides unblinded sample size re-estimation methods based on two main approaches: Chen-DeMets-Lan method [Chen et al., 2004] and Cui-Hung-Wang

method [Cui et al., 1999]. The differences between this will be explored below but both of these methods are based on using the conditional power to define whether a result is "promising" and whether to increase and how much to increase the sample size.

4.3.2.1 Conditional Power

Conditional power is the probability that the trial will reject the null hypothesis at a subsequent look given the current test statistic and the assumed parameter values, which are usually assumed to equal their interim estimates. For "promising" trials where the conditional power falls between a lower bound, a typical value would be 50%, and the initial target power the sample size can be increased to make the conditional power equal the target study power.

The conditional power is calculated by assuming a set of "true" values for the fixed term parameters, the interim standardized test statistic and the proportion of information (i.e. sample size) used at the interim analysis. The generalized formula for the conditional power at look "k" for ending at look "k+1" is as follows [Jennison and Turnbull, 1999]:

$$CP\left(k\right) = \Phi\left(\frac{Z_k\sqrt{I_k} - z_{k+1}\sqrt{I_{k+1}} + \theta(I_{k+1} - I_k)}{\sqrt{I_{k+1} - I_k}}\right)$$

where Z_k is the interim standardized test statistic at time k, I_k is the information at time k, $z_{k+1}$is the target test statistic at look k+1, I_{k+1} is the information at time k+1 and θ is the "true" parameter value of interest

For unblinded sample size re-estimation, we will assume k is set to specified sample size re-estimation look and k+1 is the final look. This means the conditional power in nQuery is an approximation for the chance of rejecting the null at any subsequent look if k is not the penultimate look but practically there will be a negligible difference between this value and the "true" conditional power.

For unblinded sample size re-estimation, the information at a given look will equal the reciprocal of the squared standard error based on the "true" parameter values (e.g. for two means, the within-group standard deviations and sample sizes: $\left[\frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2}\right]$) and θ is the "true" parameter of interest (e.g. for two means, the mean difference). These "true" parameters are assumptions made by the user. Two common suggestions would be set these to either the interim estimates of these parameters or to set these to the assumed "true" values from the original design. It will be discussed later but it should be noted that the value of the interim test statistic, Z_k, will be treated the same for the unweighted (Wald) statistic and weighted (Cui-Hung-Wang) statistic.

4.3.2.2 Unblinded Sample Size Re-estimation Methods

Unblinded sample re-estimation increases the sample size in response to a "promising" result (i.e. effect size) given the interim data. Note that sample size decreases are also allowable but practically are greatly disfavoured by most sponsors and regulators. A "promising" result is one which is clinically relevant but which the current sample size is unlikely to be found significant given the current study design and trajectory. The definition of "promising" is tied typically to conditional power, though alternative metrics such as Bayesian Predictive Power have been put forward. The sample size re-estimation is tied to the conditional power (and corresponding effect sizes) falling within a range defined as "promising". This range is defined by the study design but a typical range would be between 50% and the target power (e.g. 80% for our group sequential example). If the conditional power is based on the interim analysis, we can increase the sample size until the conditional power reaches the target level (typically the original target power level).

Thus by increasing the sample size, this "promising" result can be adequately powered for given the current available information to give the appropriate chance for finding significance at the final analysis (or any subsequent). However by allowing the study sample size to be changed based on the unblinded interim data there is the chance of increasing the overall error rate or introducing bias into a study. Two strategies for ensuring that the error rate is maintained will be discussed in this section: Chen-DeMets-Lan and Cui-Hung-Wang. Note that these methods change the method for calculating the standardized test statistics but that the group sequential bounds are unaffected and can be used as for a classic group sequential design.

Chen-DeMets-Lan The Chen-DeMets-Lan method [Chen et al., 2004] is a method of unblinded sample size re-estimation that requires minimal change from a group sequential trial. This method works by setting out a number of criteria related to the allowable conditions for increasing the sample size that ensures that the normal group sequential methods (bounds, test statistics etc.) can be used while ensuring the overall error rate is not inflated versus the original group sequential design. The two primary conditions which are required to ensure no error inflation are:

1. Sample Size Re-estimation is at the penultimate look
2. The conditional power at the penultimate look must lie between 50% and the original target power.

Chen-DeMets-Lan show that if these two conditions are met then the Type I error rate after increasing the sample size until the conditional power equals the original target power will not exceed the original target Type I error. For some conditions, the Type I error rate is lower than the original target Type I error but the conservativeness is very minor over this range. In addition, Gao, Ware and Mehta [Gao et al., 2008] and Mehta and Pocock [Mehta and Pocock, 2011] extended the work of Chen-DeMets-Lan to derive the exact lower bound for the conditional power which ensures

Type I error control for any given group sequential or equivalent fixed term analysis. In basic terms, this is done by iteratively searching over the full range of conditional powers (and corresponding effect sizes), calculating the required sample size increase for each conditional power and then calculating the appropriate weighted test statistic (Cui-Hung-Wang) that ensures error control. For the range where this weighted statistic is lower than the original critical unweighted test statistic, the sample size can be increased while maintaining the Type I error.

After the sample size increase occurs, the final analysis is conducted as if it were a standard group sequential design and uses the pre-existing group sequential bound values and Wald test statistic calculation (subsection 4.2.2).

Cui-Hung-Wang The Cui-Hung-Wang method [Cui et al., 1999] is an extension of the generalized adaptive design approach of Muller and Schaefer for the case of sample size re-estimation. Under the Cui-Hung-Wang method, the user has full flexibility over which look sample size increases can occur, how many sample size increases occur, how large the sample size increase can be and the range of "promising" conditional powers allowable for sample size increases. This robustness to in-study changes gives the greatest operational flexibility possible regarding sample size. However, this method requires the usage of a weighted test statistic which differs from the classic standardized test statistic used in group sequential designs.

The Cui-Hung-Wang method uses a weighted test statistic based on the weighted combination of the incremental test statistics. Previously when discussing test statistics, we have been referring to the cumulative test statistics which are based on all the interim data up to that point. Incremental test statistics are those calculated based only on the data between the last look and the current look. However, the cumulative and incremental test statistics can be easily related for the traditional group sequential trial as follows:

$$Z_{Cumulative} = \frac{\sqrt{w_1}Z_1+\sqrt{w_2}Z_2+...+\sqrt{w_{k-1}}Z_{k-1}+\sqrt{w_k}Z_k}{\sqrt{w_1+w_2+...+w_{k-1}+w_k}}$$

where $Z_{1,2...,k-1,k}$ are the respective incremental test statistics at each look and $w_{1,2...,k-1,k}$ are the weights for each test statistics. The weights must sum to one.

For the classic group sequential test statistic, the weights will equal the proportion of the total sample size used for each incremental test statistic (e.g. for equally spaced 4 total look design they would equal 0.25, 0.25, 0.25, 0.25). The Cui-Hung-Wang and Chen-DeMets-Lan method differ in how these weights are affected by an increase in sample size during the study. Note that for both statistics, the bounds for the classic group sequential trial are used as before for accepting and rejecting the null hypothesis or for ending the trial early for efficacy or futility.

For the Chen-DeMets-Lan method, the weights used in this calculation are based on the actual (i.e. updated/increased) sample size for each

incremental test statistic. Practically, this means that if a sample size occurs at the penultimate look, the weight placed on the final incremental test statistic will be increased so that the weight for the final cohort of data is proportional to the increased number of subjects between the penultimate and final look. Note that since the cumulative test statistic is equivalent to this version of the weighted sum of incremental test statistics, we can typically ignore the incremental tests statistics for practical purposes when using the Chen-DeMets-Lan or classic group sequential methods.

For the Cui-Hung-Wang method, the weights used in this calculation are based on the original proposed sample size for each incremental test statistic. This means that if a sample size increase occurs, subjects which are recruited after the sample size increase will be weighted less than subjects before the sample size re-estimation. However, this formulation ensures that the Type I error is retained regardless of the size of the sample size increase, when the sample size increase occurs and what range of conditional powers are considered "promising".

This robustness means that the Cui-Hung-Wang statistic provides a far greater level of operational control over the sample size re-estimation procedure. However, this comes at the cost of having to use a different test statistic than for the classic group sequential design and the down-weighting of subjects after the sample size increase. Note the weighting issue is a source of debate and we refer to relevant papers for further details.

4.3.3 Unblinded Sample Size Re-estimation and Interim Monitoring in nQuery

This section will show how to conduct an unblinded sample size re-estimation and interim monitoring for a group sequential design in nQuery. We will cover two examples: one for the Chen-DeMets-Lan method and one for the Cui-Hung-Wang method.

4.3.3.1 Opening the Interim Monitoring and Sample Size Re-estimation Table

To open the Interim Monitoring and Sample Size Re-estimation table, we first must complete a group sequential design sample size or power calculation in a group sequential table with this feature available (see `www.statsols.com` for details on which tables have this feature in nQuery Adapt). In this section, we will take the completed example from subsubsection 4.2.3.2.

In short, this was a group sequential design with fixed term parameters of a mean difference of -1, common within-group standard deviations of 2.5, a one-sided 5% significance level and 80% power. The group sequential design was for a 2 look

design (1 interim analysis), the interim look at 50% of subjects analysed, O'Brien-Fleming efficacy bounds and Power Family futility bounds with the Power Family parameter set to 1. Given this the overall sample size required was 174 (87 per group).

To open the Interim Monitoring & Sample Size Re-estimation tool, we select the "Interim Monitoring & Sample Size Re-estimation" button at the top of the Looks window. This option will be greyed out if a sample size or power calculation has not been completed in that column or if you do not have an nQuery Adapt license (see `www.statsols.com` and section 7.2 for details on purchasing and activating nQuery Adapt). This button is highlighted in Figure 4.7.

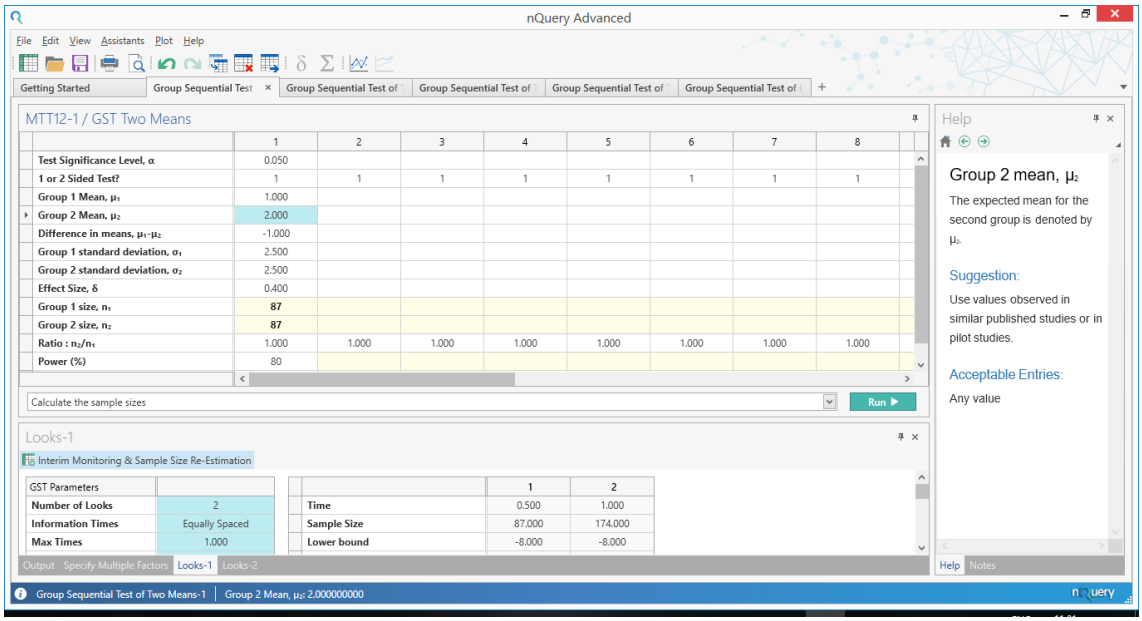

Figure 4.7: Opening Interim Monitoring and Sample Size Re-estimation Table

4.3.3.2 Interim Monitoring and Sample Size Re-estimation Table Introduction

The interim monitoring and sample size re-estimation table provides an intuitive approach for setting up a proposed sample size re-estimation and then allowing interim monitoring with that sample size re-estimation in mind. The basic structure of the interim monitoring and sample size re-estimation tables can be split into two main parts: the sample size re-estimation rules column and the interim monitoring tool. We will refer to the former as the "SSR Rules" and the latter as the "Monitoring Table" for short. We will summarise each of these in the following section.

When first opened the table will look as per Figure 4.8 for a two means group sequential design.

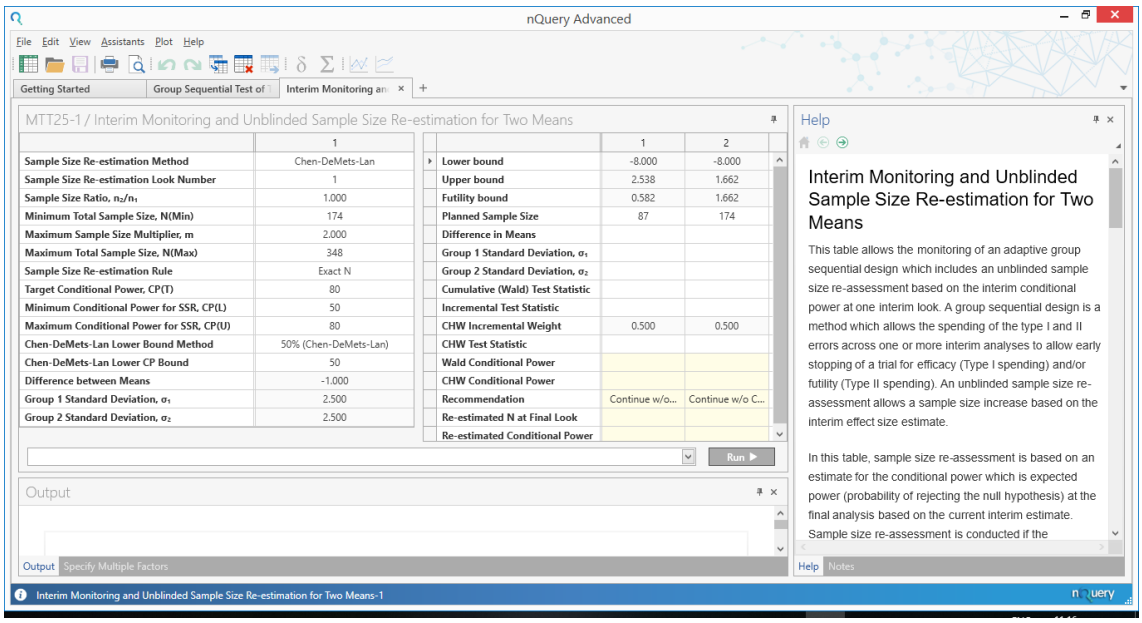

Figure 4.8: Interim Monitoring and Sample Size Re-estimation Table for Two Means GST

Setting up the Sample Size Re-estimation Rules To set up the rules that will define when and where a sample size re-estimation will occur, nQuery provides a large number of options to tailor the sample size re-estimation to the requirements of the user. These are set in the SSR Rules column (on the left in the main window). There are 12 inputs that can be edited by the user to define the sample size re-estimation rules. These are summarised below:

1. Sample Size Re-estimation Method (Chen-DeMets-Lan, Cui-Hung-Wang, None): Define the sample size re-estimation method desired for this study. The Chen-DeMets-Lan and Cui-Hung-Wang methods are summarised in subsubsection 4.3.2.2. In short, the Chen-DeMets-Lan method uses the standard group sequential test statistic but restricts sample size re-estimation to the penultimate look and for a certain range of conditional powers. The Cui-Hung-Wang method uses a different weighted test statistic from the group sequential design, based on the incremental test statistics weighted by the original planned sample size, and can be conducted for any look or conditional power. None will disable any sample size re-estimation (all subsequent options in this column can be ignored) and the tool will operate as a standard group sequential design monitoring table. This will default to the Chen-DeMets-Lan method.
2. Sample Size Re-estimation Look Number (If Chen-DeMets-Lan: Looks-1, if Cui-Hung-Wang: < Looks): The look at which a sample size re-estimation will occur. Note that nQuery Adapt allows only a single sample size re-estimation. For the Chen-DeMets-Lan method, this must equal the total number of looks (including the final analysis) minus one. For the Cui-Hung-Wang method, this

must be less than the total number of looks. Defaults to the number of looks minus one.

3. Sample Size Ratio (>0): The ratio between the group 1 and group 2 sample sizes. This will be assumed to be constant across all looks. Defaults to the sample size ratio from the original group sequential design.
4. Minimum Total Sample Size (>2): The minimum total sample size allowable in this study. Changing this will automatically update the Planned Sample Size row in the Monitoring Table to reflect this value. Defaults to the sample size from the original group sequential design.
5. Maximum Sample Size Multiplier (>1): A multiplier for the minimum total sample size versus the maximum total size. Changing this will automatically update the Maximum Total Sample Size to reflect this value. Defaults to 2 (maximum sample size double minimum sample size).
6. Maximum Total Sample Size (>Minimum Total Sample Size): The maximum total sample size allowable in this study. If a sample size re-estimation occurs, this will be maximum allowable total sample size that will occur even if the target conditional power is not reached with this value. Defaults to twice the minimum total sample size.
7. Sample Size Re-estimation Rule (Exact N, Maximum N): The rule for how much the sample size should be increased by if the interim statistic gives a conditional power within the Minimum and Maximum Conditional Power for SSR value for the Sample Size Re-estimation Look Number column. Exact N will increase the total sample size until either the conditional power equals the Target Conditional Power or the Maximum Total Sample Size is reached. Max N will always increase the sample size to the Maximum Total Sample Size. The Max N rule may over-power the study but reduces the operational risk of study participants reverse-calculating the effect size using the "Exact N" calculated to restore the conditional power to the original target power.
8. Target Conditional Power (>Minimum Conditional Power for SSR, <100): This is the conditional power targetted by the sample size re-estimation increase if the Exact N rule is active in the Sample Size Re-estimation Rule row. The sample size is increased until this value is reached or the maximum sample size is reached. Defaults to the target power from the original group sequential design.
9. Minimum Conditional Power for SSR (<Target Conditional Power & Maximum Conditional Power for SSR): This is the minimum conditional power that would be considered "promising" and lead to a sample size re-estimation at the specified Look Number. For the Chen-DeMets-Lan method, this needs to be greater than the Chen-DeMets-Lan Lower CP Bound. Defaults to 50%.
10. Maximum Conditional power for SSR (>Minimum Conditional Power for SSR, <100): This is the maximum conditional power that would be considered

"promising" and lead to a sample size re-estimation at the specified Look Number. Conditional powers above this value would have sufficient power and then a sample size increase would not be necessary. For the Chen-DeMets-Lan method, this should not be set higher than the power from the original group sequential design. Defaults to the target power from the original group sequential design.

11. Chen-DeMets-Lan Lower Bound Method (50% (Chen-DeMets-Lan), Derived (Mehta-Pocock)): Select the method used to the derive the allowable Minimum Conditional Power for SSR which can be used without leading to error inflation. Chen-DeMet-Lan sets this to 50% which holds for all group sequential designs. Derived (Mehta-Pocock) derives the lower bound which ensures the Type I error never exceeds the level for the equivalent Cui-Hung-Wang test statistic and this value will always fall between 0% and 50%. Defaults to 50% (Chen-DeMets-Lan).

In addition to these parameters, the SSR Rules column will also include the fixed term parameters used to define the study effect size. These values can be useful for reference and if the user wishes to base conditional power calculations on these original design parameters rather than the interim estimates for these parameters. For the two means case, the relevant parameters are the difference in means and group 1 and group 2 standard deviations. Once we have set our SSR Rules, we can use the Monitoring Table to start entering our interim statistics and stop our trial early, implement a sample size increase if needed or find for or against the null hypothesis at our final look. More detail on all these rows is given in the Home and Help cards of the table.

4.3.3.3 Interim Monitoring Table

After the SSR Rules are set we need to enter our interim results to decide what action should be made at each look based on those results. To do this nQuery provides the Monitoring Table (on the right in the main window) which contains 16 rows and a number of columns equal to the total number of looks in the original group sequential design.

In this table each column corresponds to the respective interim look (e.g. column 1 is look 1, column 2 is look 2 etc.). To use this table we enter the required inputs in each column sequentially (i.e. complete column 1 before moving onto column 2 and so on) until the conditional power is calculated and Recommendation is made based on those inputs. We fill in each column in turn until either a recommendation is made to stop the trial early (ignorable if below non-binding futility bound) or the final column is filled and a recommendation is made to find for efficacy or futility.

There are 16 rows per column in the Monitoring Table for the Two Means case. These are summarised below:

1. Lower Bound (Read-only): The Lower Bound for efficacy below which we would end the trial early. The bounds are on the test statistic scale. These are inherited from the original group sequential design.
2. Upper Bound (Read-only): The Upper Bound for efficacy above which we would end the trial early. The bounds are on the standardized test statistic scale. These are inherited from the original group sequential design.
3. Futility Bound (Read-only): The Futility Bound below which we would end the trial early for futility (can ignore if non-binding). The bounds are on the test statistic scale. These are inherited from the original group sequential design.
4. Planned Sample Size (>2, >Planned Sample Size in Current Column - 1): The cumulative sample size analyzed at the current look. By default, this will equal the sample size at each look from the original group sequential design. This value can be changed in each column manually using the Minimum Total Sample option in the SSR rules or can be edited manually at any point by the user. This value will also be over-written by the cumulative test statistic side-table if the default value for the Cumulative Interim Total Sample Size (which defaults to this value) is changed and the Transfer Estimates option is set to Yes. If a sample size re-estimation is recommended these values will automatically update such that the sample size increase will happen between the Sample Size Re-estimation Look and subsequent look and all other Planned Sample Sizes will increase accordingly.
5. Difference in Means ($\neq$0, Unique to Two Means): The "true" assumed value for the difference in means used for the conditional power calculation. Usually set to either the interim estimate for this parameter or the assumed parameter value from the original design. The Cumulative Test Statistic side-table will automatically set this to the Interim Difference Between Means if the Transfer Estimates option is set to Yes. The original design value for this parameter is available at the bottom of the SSR Rules column.
6. Group 1 Standard Deviation (>0, Unique to Two Means): The "true" assumed value for the group 1 standard deviation used for the conditional power calculation. Usually set to either the interim estimate for this parameter or the assumed parameter value from the original design. The Cumulative Test Statistic side-table will automatically set this to the Interim Group 1 Standard Deviation if the Transfer Estimates option is set to Yes. The original design value for this parameter is available at the bottom of the SSR Rules column.
7. Group 2 Standard Deviation (>0, Unique to Two Means): The "true" assumed value for the group 2 standard deviation used for the conditional power calculation. Usually set to either the interim estimate for this parameter or the assumed parameter value from the original design. The Cumulative Test Statistic side-table will automatically set this to the Interim Group 2 Standard

Deviation if the Transfer Estimates option is set to Yes. The original design value for this parameter is available at the bottom of the SSR Rules column.

8. Cumulative (Wald) Test Statistic (Any Value): The standardized test statistic based on the interim data (see subsection 4.2.2) This is equivalent to the standard Z statistic for this data and is alternatively called the unweighted statistic (although its weighting is simply based on the empirical sample size rather than the initial sample size). This is the test statistic used for conditional power when the Chen-DeMets-Lan method or None options are selected from the Sample Size Re-estimation Method row in the SSR Rules column. This can be calculated using the Cumulative Test Statistic side-table based on the interim estimates for relevant effect size parameters (for two means, the mean difference and group 1 and 2 standard deviations). If this value falls above the Upper Bound or below the Lower Bound (2-sided), the trial should be ended early for efficacy. If this value falls below the Futility Bound, the trial can be ended early for futility.
9. Incremental Test Statistic (Any Value): The standardized test statistic (see subsection 4.2.2) based on the incremental data i.e. for the data between the last look and the current look. This is equivalent to the standard Z statistic for this incremental data. This incremental test statistics for all columns up to and including the current column are used to derive the Cui-Hung-Wang (CHW) Test statistic in combination with CHW Incremental Weights. The CHW test statistic is then used for conditional power when the Cui-Hung-Wang method is selected from the Sample Size Re-estimation Method row in the SSR Rules column. This can be calculated using the Incremental Test Statistic side-table based on the incremental estimates for relevant effect size parameters (for two means, the mean difference and group 1 and 2 standard deviations).
10. CHW Incremental Weight (Read-only): The incremental weights used in combination with the incremental test statistics to derive the Cui-Hung-Wang (CHW) Test Statistic. These will equal the original proposed incremental sample size used between each look and its prior look.
11. CHW Test Statistic (Any Value): The Cui-Hung-Wang Test weighted test statistic used for conditional power calculations when the Cui-Hung-Wang method selected from the Sample Size Re-estimation row of the SSR Rules. This will typically be auto-calculated using the incremental test statistic and CHW Incremental Weight column values (see section 4.3.2.2). However, these values can be entered manually if desired. If this value falls above the Upper Bound or below the Lower Bound (2-sided), the trial should be ended early for efficacy. If this value falls below the Futility Bound, the trial can be ended early for futility.
12. Wald Conditional Power (Read-only): The conditional power based on the "true" parameter estimates and the Cumulative (Wald) Test Statistic value

based on the interim data. This will only be calculated if the Chen-DeMets-Lan and None options are selected from Sample Size Re-estimation Method row in SSR Rules. For the sample size re-estimation look column, if the conditional power falls between the Minimum Conditional Power for SSR and Maximum Conditional Power for SSR then a sample size re-estimation will automatically activate.

13. CHW Conditional Power (Read-only): The conditional power based on the "true" parameter estimates and the CHW Test Statistic value based on the interim data. This will only be calculated if the Cui-Hung-Wang option is selected from Sample Size Re-estimation Method row in SSR Rules. For the sample size re-estimation look column, if the conditional power falls between the Minimum Conditional Power for SSR and Maximum Conditional Power for SSR then a sample size re-estimation will automatically activate.

14. Recommendation (Read-only): This gives the recommendation for this look based on the relevant test statistic and conditional power. There are four recommendations that can be made at a given interim look: a) Stop for Efficacy - Stop trial early due to strong evidence for efficacy due to test statistic being beyond Upper (1-sided or 2-sided) or Lower Efficacy Bounds (2-sided only) b) Stop for Futility - Stop trial early due to strong evidence for futility due to test statistic being below Futility bound (1-sided only) c) Continue without (w/o) change - Continue until next look with no changes due to test statistic not being in the early stopping region or the conditional power falling within the lower and upper SSR conditional power bounds d) Add N & Continue - Add required additional sample size for "promising" design and continue trial until next look as conditional power fell between lower and upper SSR conditional power bound. At the final look, two recommendations can be made: "Find for Efficacy" (for 1-sided, test statistic above upper bound; for 2-sided, test statistic outside lower to upper bound range) or "Find for Futility" (for 1-sided, test statistic below upper bound; for 2-sided, test statistic inside lower to upper bound range) .

15. Re-estimate N at Final Look (Read-only): The recommended total sample size after a sample size re-estimation occurs. This will be empty unless a sample size re-estimation has been recommended in this column (Add N & Continue in Recommendation row). The increase in sample size will automatically place the additional sample size between the current look and the next look, with the Planned Sample Size row automatically updated to reflect this. This pattern can be edited manually if desired in the Planned Sample Size row in the subsequent columns.

16. Re-estimated Conditional Power (Read-only): The conditional power based on the interim test statistic and "true" parameter values based on the sample size increase calculated in the Re-estimate N at Final Look row. This will be empty unless a sample size re-estimation has been recommended in this

column (Add N & Continue in Recommendation row). For Exact N, this will equal the target conditional power approximately unless the maximum sample size is insufficient to restore the conditional power to the target conditional power.

More detail on these is given in the Home and Help cards of the table. Note that the difference in means and standard deviation rows are unique to the two mean design. Other designs will have the appropriate parameters instead but the basic workflow and assumptions are the same for other designs.

Note that in nQuery; "upper", "lower" and "futility" are defined relative to the direction of the pre-specified effect size. For example, if the initial design assumes a negative difference ($\mu_1 > \mu_2$) then positive "upper" interim tests statistics would be those in which the interim difference is also negative and vice-versa. Similarly, "lower" results would be those in the opposite direction. Futility bounds can be above or below zero but the same guidance applies for the Z-statistic calculation.

4.3.3.4 Monitoring Side-tables

In addition to these tables, nQuery provides two side-tables which make calculating and entering the required inputs easier for the "true" parameter values and the test statistics required in the Monitoring Table. These two side-tables are the Cumulative Test Statistic side-table and the Incremental Test Statistic Side-Table.

These side-tables allow the user to enter relevant interim/incremental effect size estimates and calculate the relevant test statistic and transfer this into the Monitoring Table. The Cumulative Test Statistic side-table will automatically open in the window if the "true" parameter values or Cumulative (Wald) Test Statistic rows are selected in a column and the Incremental Test Statistic side-table will automatically open in the window below the main table if the Incremental Test Statistic row is selected. Functionally, these work the same as side-tables seen in other tables (see section 3.1).

The calculations in both side-tables are identical and based on the relevant standardized test statistic calculations from subsection 4.2.2. They differ in that the Cumulative Test Statistic side-table requires the user input the cumulative interim estimates (i.e. based on all the data up to and including the current look) and the Incremental Test Statistic side-table requires the user input the incremental estimates (i.e. based only on the data from the last look to the current look). In addition, the Cumulative Test Statistic side-table gives the option to transfer the relevant interim effect size parameters and sample size into the relevant rows of the Monitoring Table by setting the Transfer Estimates option to Yes.

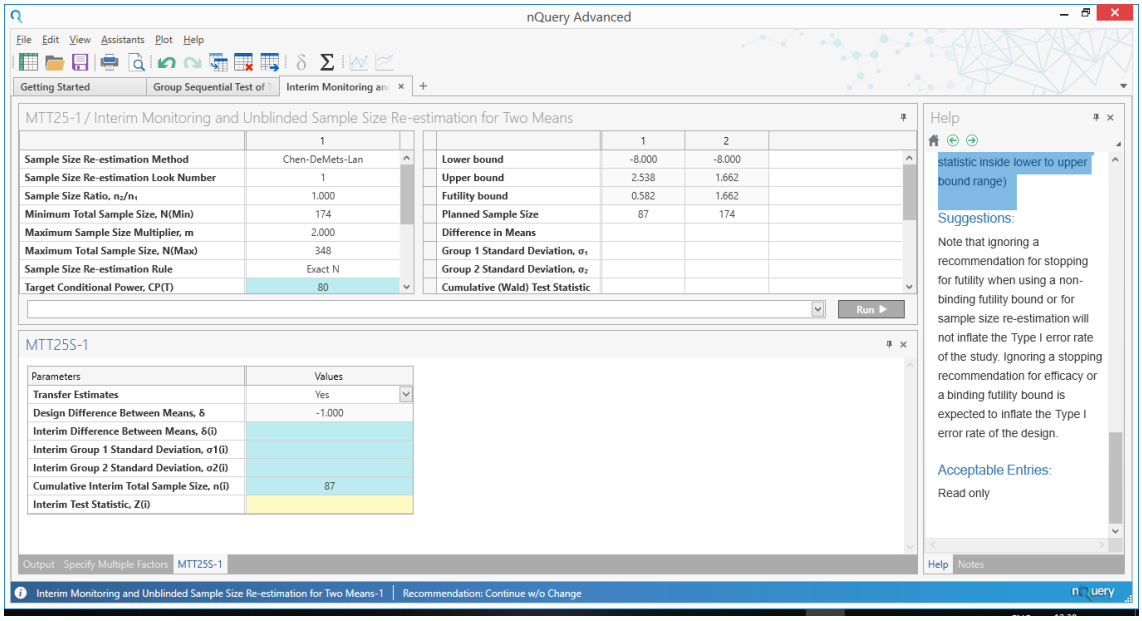

Figure 4.9: Cumulative Test Statistic Side-Table for Two Means

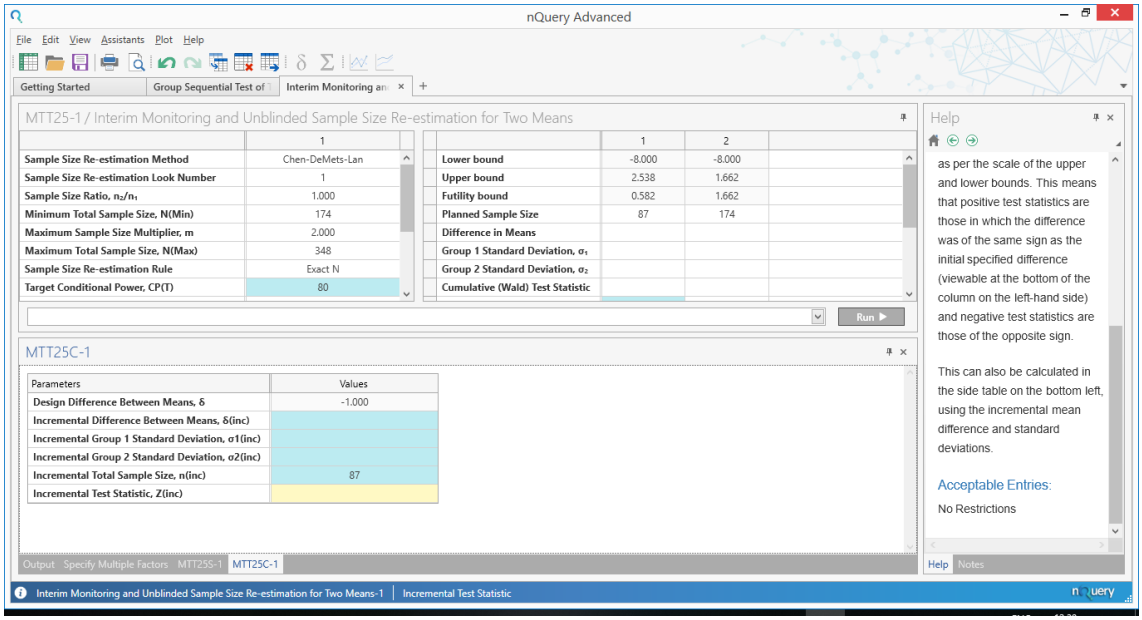

Figure 4.10: Incremental Test Statistic Side-Table for Two Means

The Cumulative Test Statistic side-table can help calculate the cumulative (Wald) test statistic for the Chen-DeMets-Lan method while also automatically assigning the "true" parameter values for conditional power as being equal to their interim estimates. The Incremental Test Statistic side-table can help calculate the incremental test statistics needed to calculate the Cui-Hung-Wang (CHW) test statistic

in combination with the pre-set CHW Incremental Weights. The functionality of both will be explored further in the worked example that follows.

Note that in nQuery; "upper", "lower" and "futility" are defined relative to the direction of the pre-specified effect size. For example, if the initial design assumes a negative difference ($\mu_1 > \mu_2$) then positive "upper" interim tests statistics would be those in which the interim difference is also negative and vice-versa. Similarly, "lower" results would be those in the opposite direction. Futility bounds can be above or below zero but the same guidance applies for the Z-statistic calculation. Thus these side-tables will provide the Design effect (e.g. difference in means). Cumulative/Incremental effects (e.g. differences) which have the same sign as the design effect will give a positive test statistic and those of the opposite sign will give a negative test statistic.

4.3.4 Interim Monitoring and Sample Size Re-estimation Worked Examples

4.3.4.1 Chen-DeMets-Lan Worked Example

For this example, we will use the group sequential design from subsubsection 4.2.3.2 and have a Chen-DeMets-Lan sample size re-estimation based on the default rules used by the Interim Monitoring and Sample Size Re-estimation table. For reference, this was a group sequential design with fixed term parameters of a mean difference of -1, common within-group standard deviations of 2.5, a one-sided 5% significance level and 80% power. The group sequential design was for a 2 look design (1 interim analysis), the interim look at 50% of subjects analysed, O'Brien-Fleming efficacy bounds and Power Family futility bounds with the Power Family parameter set to 1. This gave an overall sample size of 174 (87 per group).

These default rules will give a sample size re-estimation at the penultimate look (i.e. Look 1 for two look design), the minimum sample size will equal our original sample size (174), the maximum sample size will be twice our original sample size (348), we will only increase sample size if needed until the target conditional power (equalling our initial target power of 80%) is reached and will only increase the sample size if the conditional power is between 50% (Chen-DeMets-Lan lower bound) and 80% (target power) at Look 1. Before we enter our interim data, our table will look as per Figure 4.8.

We will now enter the relevant inputs in column 1 of the monitoring table. Let us assume that the interim data suggests a 25% reduction from our initial estimate of the effect size (i.e. -0.75 mean difference, assuming standard deviations are the same). We could enter the relevant mean difference and standard deviations manually and also calculate the cumulative test statistic manually and enter it here. However, we will use the Cumulative Test Statistic Side-Table to calculate these automatically instead. When we select the "Difference in Means" row in column 1, this side-table

will open automatically in the window below the main table. In this side-table, we leave the Transfer Estimates option as Yes so that our entered interim estimates will be set to be the true values for the conditional power calculations. We will also leave the Cumulative Sample Size as its default value of the Planned Sample Size for this column. We then enter our interim estimates for the difference in means and the group 1 and 2 standard deviations. In this case, these are -0.75, 2.5 and 2.5 respectively. The side-table will automatically calculate the Cumulative Test Statistic and transfer it and the interim parameter estimates into the main table.

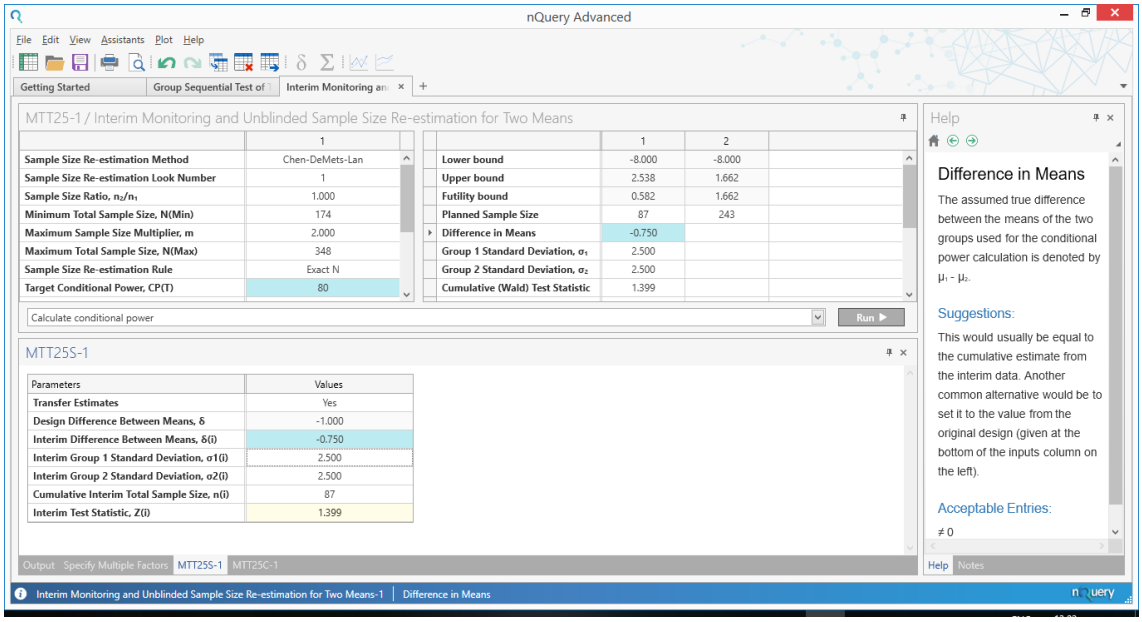

Figure 4.11: Chen-DeMets-Lan Side-Table Example

Note that in nQuery; "upper", "lower" and "futility" are defined relative to the direction of the pre-specified effect size. For example, if the initial design assumes a negative difference ($\mu_1 > \mu_2$) then positive "upper" interim tests statistics would be those in which the interim difference is also negative and vice-versa. Similarly, "lower" results would be those in the opposite direction. Futility bounds can be above or below zero but the same guidance applies for the Z-statistic calculation. Thus these side-tables will provide the Design effect (e.g. difference in means). Cumulative/Incremental effects (e.g. differences) which have the same sign as the design effect will give a positive test statistic and those of the opposite sign will give a negative test statistic. Thus, in this case entering -0.75 returns an interim test statistic of 1.399 but a value of 0.75 would have returned a test statistic of -1.399.

Once these values have been transferred, the conditional power is automatically calculated in this column in the Monitoring Table. In this case, these interim results give a conditional power of 67.28% which falls between our Conditional Power Sample Size Re-estimation range of 50% to 80%. Thus nQuery gives a recommend-

ation of "Add N & Continue" in the Recommendation row in column 1 and nQuery automatically calculates the required total sample size to increase the conditional power to 80%. In this case, this corresponds to an increase in the total sample size from 176 to 243, giving an updated conditional power of 80.06%. The Planned Sample Size in column 2 is updated automatically to reflect this.

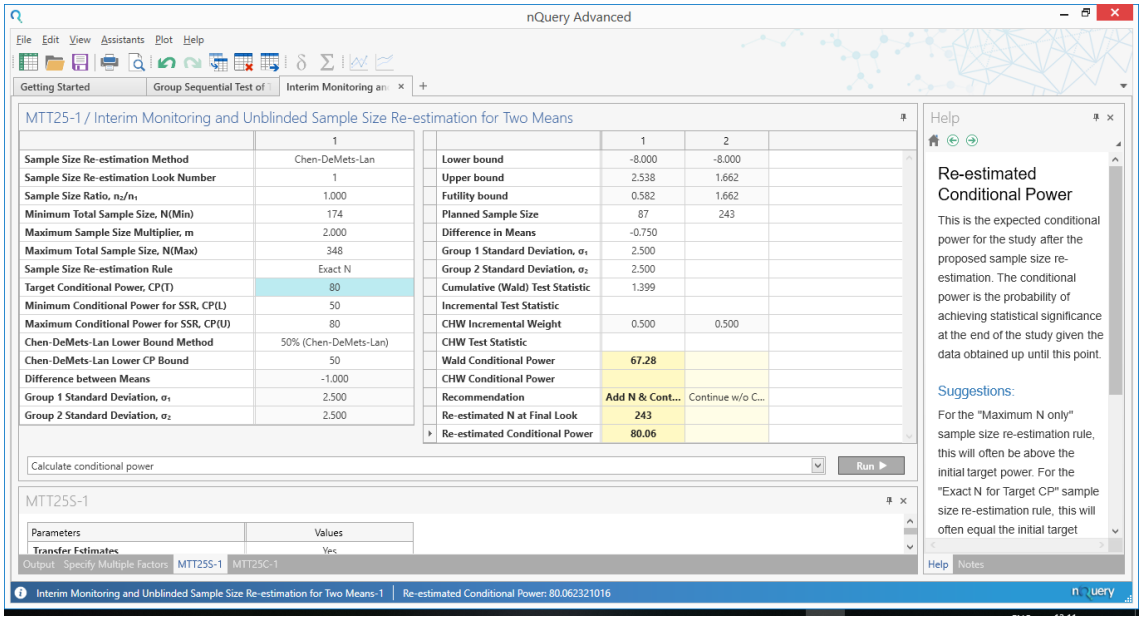

Figure 4.12: Chen-DeMets-Lan Sample Size Re-estimation Example

We can now complete our interim monitoring by entering our estimates at the end of our (now larger) study. Assume that the final analysis has the same difference in means (-0.75) and per-group standard deviations (2.5). We can use the Cumulative Test Statistic Side-table as before and calculate a test statistic of 2.338. As this test statistic value is greater than the Upper Bound (1.662) at the final analysis, we find for efficacy for this study at the final look.

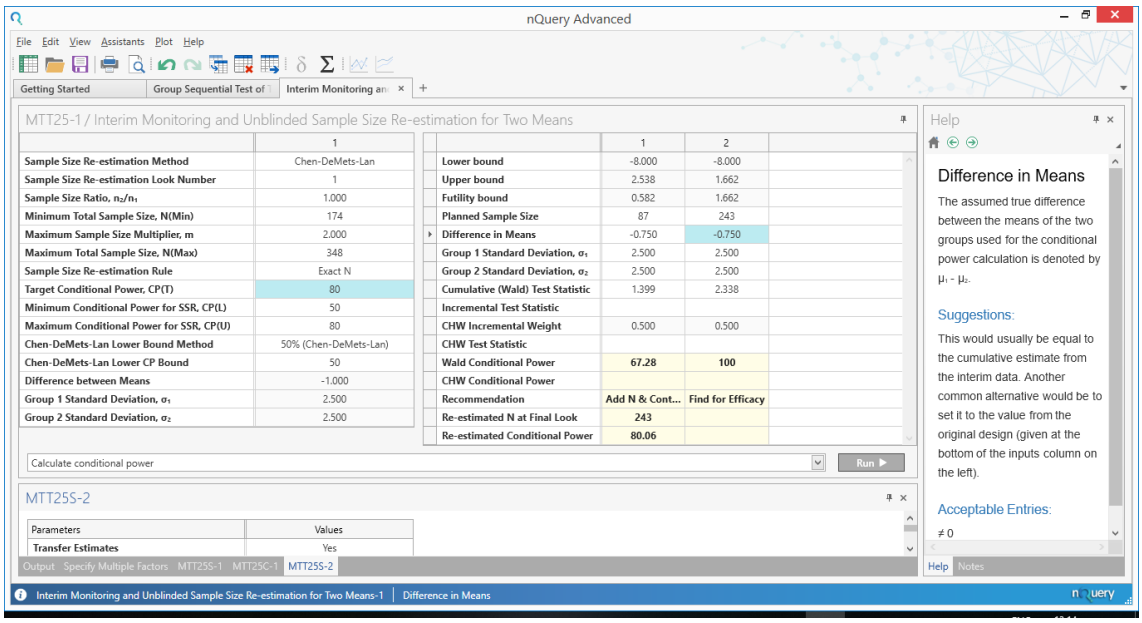

Figure 4.13: Chen-DeMets-Lan Completed Design Example

4.3.4.2 Cui-Hung-Wang Worked Example

For this example, we will assume the same group sequential design.

For reference, this was a group sequential design with fixed term parameters of a mean difference of -1, common within-group standard deviations of 2.5, a one-sided 5% significance level and 80% power. The group sequential design was for a 2 look design (1 interim analysis), the interim look at 50% of subjects analysed, O'Brien-Fleming efficacy bounds and Power Family futility bounds with the Power Family parameter set to 1. This gave an overall sample size of 174 (87 per group).

In this case we assumed a 50% reduction in our difference in means versus the original study design (-0.5 difference in means).

For this sample size re-estimation we will also make some additional changes to the previous example. In this example we will increase the sample size multiplier to 3 (maximum N equal to 522), use a Max N rule for sample size increases (i.e. N always increased to the maximum N) and decrease the lower bound for the conditional power to 30%. This will give SSR Rules inputs as in Figure 4.14.

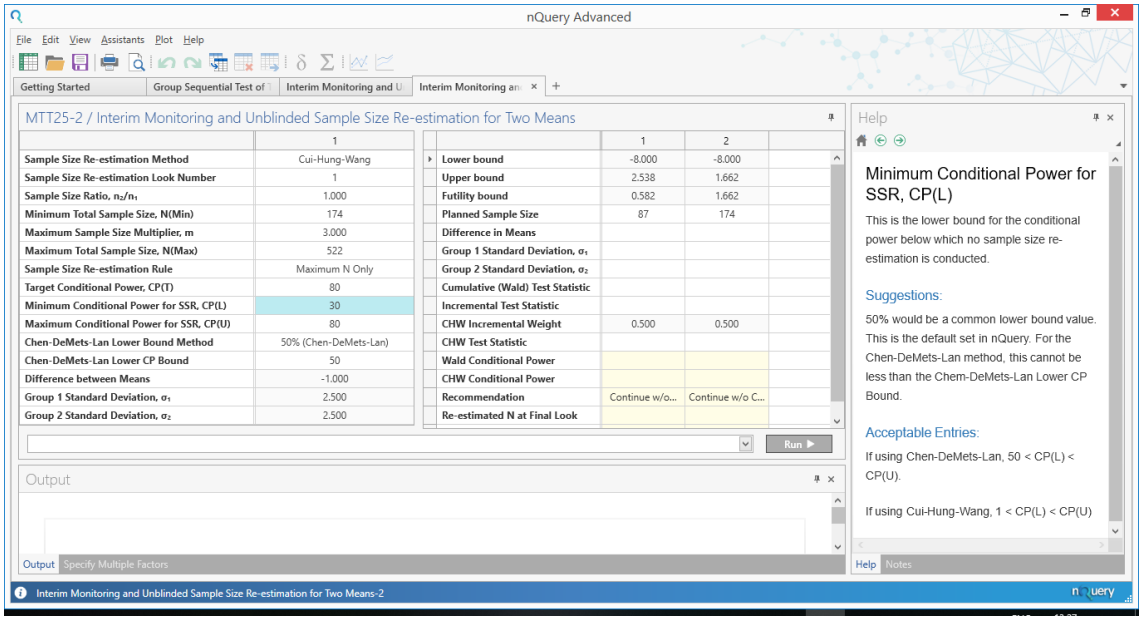

Figure 4.14: Cui-Hung-Wang Example Rules Setup

We will now enter the relevant inputs in column 1 of the monitoring table. In this case, we will enter our interim estimates for the design parameters in column 1 (difference in means of -0.5, per-group standard deviations equal to 2.5). We will then use the Incremental Test Statistic Side-Table to calculate the Incremental Test Statistic. When we select the "Incremental Test Statistic" row in column 1, this side-table will open automatically in the window below the main table. In this side-table, we will leave the Incremental Sample Size as its default value of the Planned Sample Size for this column. For designs with a greater number of columns this would equal the Planned Sample Size in the current column minus the Planned Sample Size in the prior column. We then enter our incremental estimates for the difference in means and the group 1 and 2 standard deviations. In this case, these are -0.5, 2.5 and 2.5 respectively. The side-table will automatically calculate the Incremental Test Statistic and transfer it into the main table.

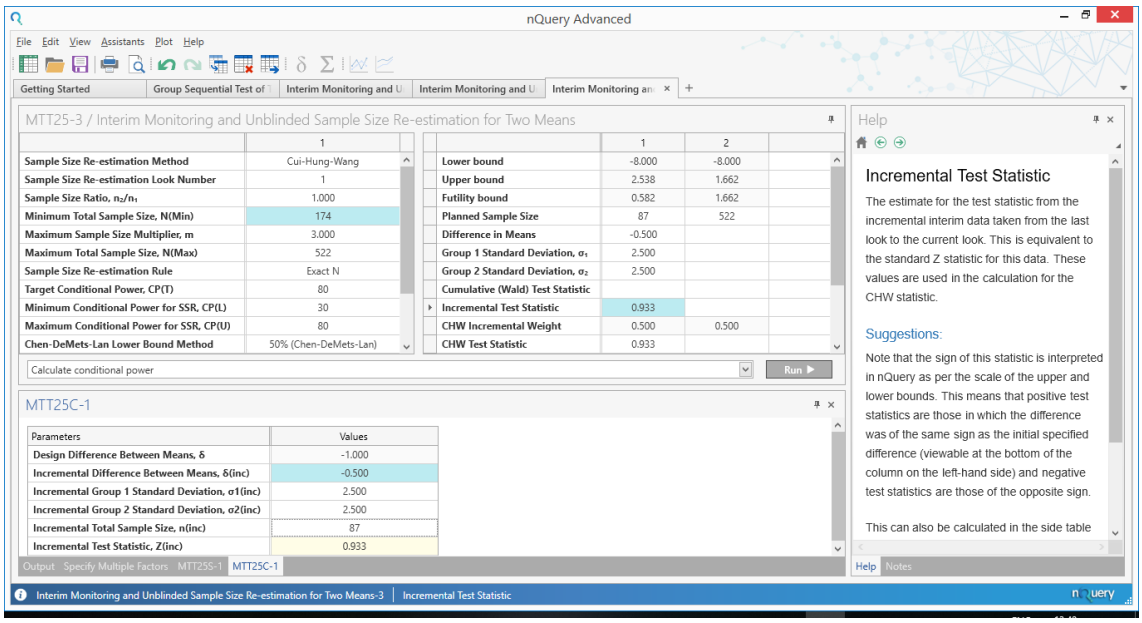

Figure 4.15: Cui-Hung-Wang Side-Table Example

Note that in nQuery; "upper", "lower" and "futility" are defined relative to the direction of the pre-specified effect size. For example, if the initial design assumes a negative difference ($\mu_1 > \mu_2$) then positive "upper" interim tests statistics would be those in which the interim difference is also negative and vice-versa. Similarly, "lower" results would be those in the opposite direction. Futility bounds can be above or below zero but the same guidance applies for the Z-statistic calculation. Thus these side-tables will provide the Design effect (e.g. difference in means). Cumulative/Incremental effects (e.g. differences) which have the same sign as the design effect will give a positive test statistic and those of the opposite sign will give a negative test statistic. Thus, in this case entering -0.5 returns an interim test statistic of 0.933 but a value of 0.5 would have returned a test statistic of -0.933. It is important to note that for the looks before the sample size re-estimation look that the Cumulative and CHW tests statistics should be identical. You may want to calculate both at each look to confirm this.

Once these values have been transferred, the conditional power is automatically calculated in this column in the Monitoring Table. In this case, these interim results give a conditional power of 31.38% which falls between our Conditional Power Sample Size Re-estimation range of 30% to 80%. Note that no sample size re-estimation would have occurred with the default range. Thus nQuery gives a recommendation of "Add N & Continue" in the Recommendation row in column 1 and nQuery automatically sets this sample size to the maximum sample size of 522. In this case, this gives an updated conditional power of 75.24%. Note that this conditional power is lower than the target conditional power. With the Exact N rule, the result would have been the same in this case as the target conditional power was

not reached with this maximum sample size. The Planned Sample Size in column 2 is updated automatically to reflect this.

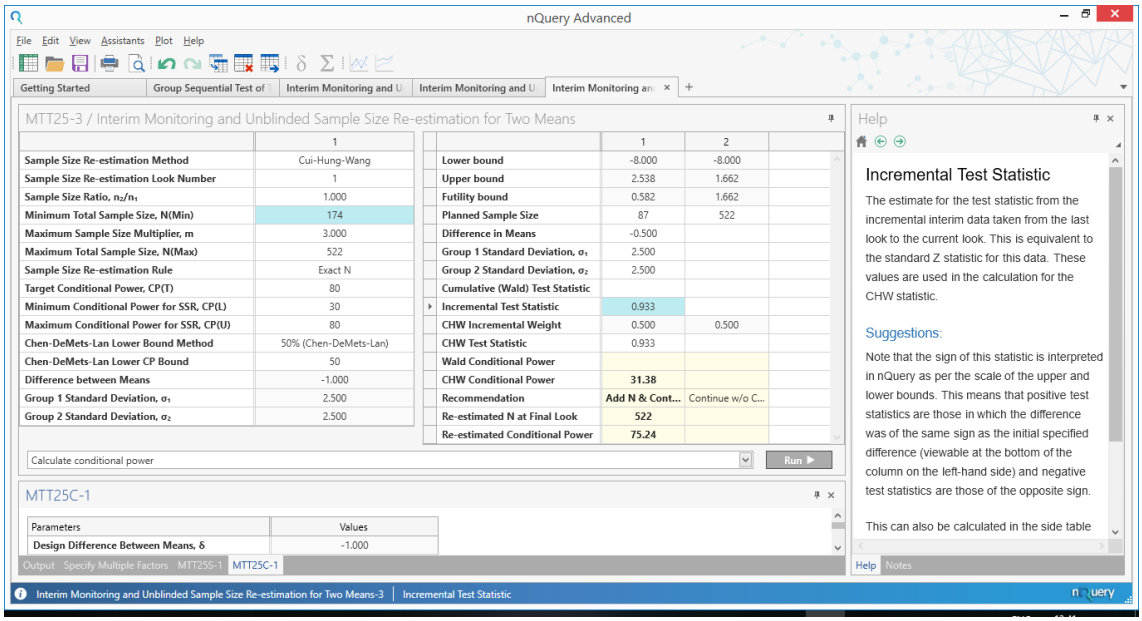

Figure 4.16: Cui-Hung-Wang Sample Size Re-estimation Example

We can now complete our interim monitoring by entering our estimates at the end of our (now larger) study. Assume that the final analysis has the same difference in means (-0.5) and per-group standard deviations (2.5). We can enter these parameter values in column 2 and use the Incremental Test Statistic Side-table as before and calculate a CHW test statistic of 2.134 based on an incremental test statistic of 2.086. As this test statistic value is greater than the Upper Bound (1.662) at the final analysis, we find for efficacy for this study at the final look.

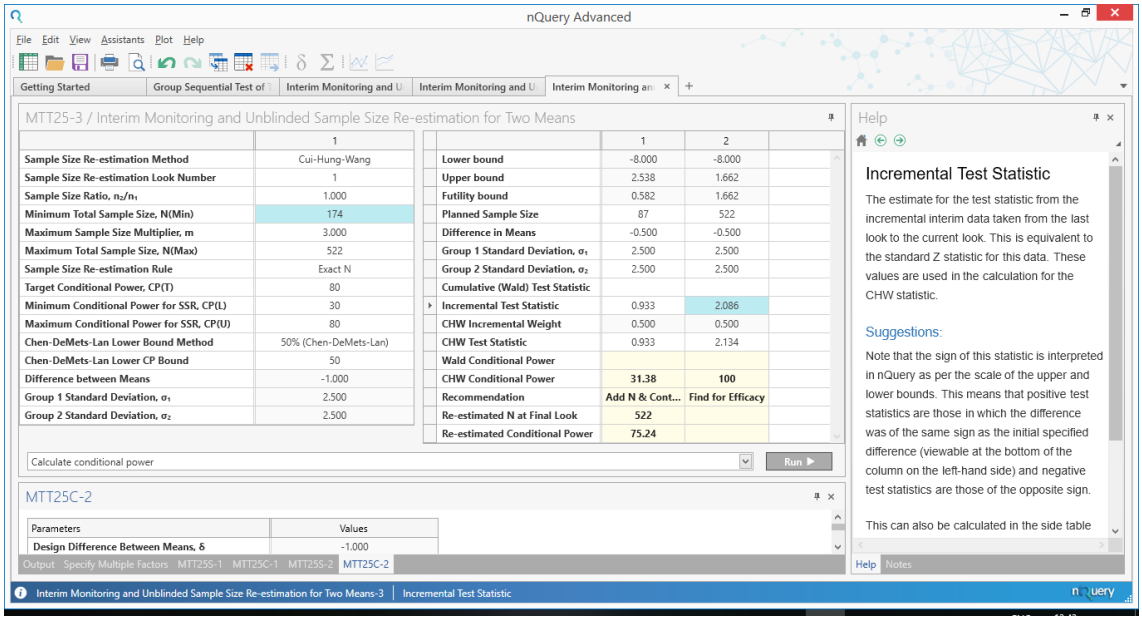

Figure 4.17: Cui-Hung-Wang Completed Design Example

5 Multiple Comparisons Procedure - Modelling (MCP-Mod)

nQuery has the facility to calculate the sample size or power needed for the "Proof-of-Concept" (PoC) stage of an MCP-Mod trial. As the MCP-Mod methodology has two main purposes (PoC, dose-finding (DF)) different criteria can be used to calculate the required sample size or power.

One approach is to concentrate on the PoC step of the procedure, calculating the sample size needed to ensure a minimum power level for detecting PoC under an assumed placebo and maximum dose-response and an assumed set of dose-response models. The focus here and in nQuery is on this approach.

Another strategy is to focus on the dose-estimating part of the method. This would find the sample size needed for a prespecified minimum precision (i.e. confidence interval width) for the dose-response estimates. This approach is not explored here.

This section will give a background to MCP-Mod, the underlying theory and derivations and the methods employed in nQuery. A demonstration of how to design an MCP-Mod trial in nQuery is also included.

5.1 Multiple Comparison Procedures and Modelling (MCP-Mod) Design

5.1.1 Background

MCP-Mod is a type of Phase II trial design methodology used in dose-response studies. Phase II studies have the following two main goals:

1. Proof-of-Concept (PoC): The goal here is to establish that changes in dose lead to significance changes in the efficacy (and/or safety) endpoint and that at least one dose is of clinical interest.
2. Dose-finding (DF): The second goal is to select one or more effective (and safe) doses for evaluation in confirmatory (Phase III) clinical trials or equivalent studies.

Historically, the proof-of-concept and dose finding objectives have often been evaluated in separate trials with proof-of-concept established first and then a dose-finding

study used to select doses of interest for further evaluation. In this framework, the proof-of-concept trial was referred to as the Phase IIa trial and the dose-finding trial was referred to as the Phase IIb trial, with the proof-of-concept and dose-response studies using different statistical methods. For proof-of-concept Phase IIa trials a common approach was the multiple comparison procedure (MCP) and for dose-finding Phase IIb trials a common approach was parametric dose-response modelling (Mod).

Under the Multiple Comparison Procedure (MCP) strategy, one evaluates the statistical significance of a contrast test between doses while preserving the family wise error rate (FWER). PoC is established when at least one contrast is statistically significant. If PoC is established, the minimum effective dose (MED) is found. The MED is the lowest dose that is considered statistically (significant) and clinically superior to placebo. The MCP approach regards the dose as a qualitative factor and generally makes few, if any, assumptions about the underlying dose response relationship. However, inferences about the target dose are restricted to the discrete, possibly small, set of doses used in the trial.

The modeling approach is primarily used to estimate the true dose-response curve. Doses that achieve desired clinical effects are estimated via inverse regression methods, which can also be used to evaluate the precision (i.e. confidence intervals) of the estimated doses. The dose is taken to be a quantitative factor, allowing greater flexibility for target dose estimation including potentially doses not evaluated directly in the study. However, the validity of the modelling approach strongly depends on choosing the appropriate parametric dose-response model (e.g. linear, Emax, Beta). Note that modelling can also be used for PoC by testing the significance of the parametric fitted dose-response model versus a flat dose-response null model. However, this approach is less robust than MCP due to the sensitivity to model choice as mentioned previously.

The MCP-Mod methodology was designed to provide a unified strategy to the analysis of data from dose-response studies which combines strengths of the multiple comparison procedure and modelling approaches by having the flexibility of modeling for dose estimation, while preserving the robustness to model misspecification associated with MCP.

The basic idea behind MCP-Mod was first proposed in [Tukey, 1985]. They proposed the use of several trend tests simultaneously and to adjust the resulting p-values for multiplicity. [Bretz et al., 2005] formalized the MCP-Mod procedure and so users are recommended to see this paper for more details on the basic methodology. A more detailed description of the method and it's practical implementation, including power analysis, can be found in [Pinheiro et al., 2006].

5.1.2 MCP-Mod Theory

5.1.2.1 MCP-Mod Model

The notation used to represent the dose-response model is presented here. Let Y denote the observed response for a given set of patients assigned to one of a set of k doses which we will denote as d where $d = d_1, \ldots, d_k$. We will consider the following dose-response model:

$$Y_{ij} = \mu_{d_i} + \epsilon_{ij}, \epsilon_{ij} \sim N(0, \sigma^2), i = 1, \ldots, k.j = 1, \ldots, n_i$$

Where $\mu_{d_i} = f(d_i, \theta)$ denotes the mean response at dose d_i, θ represents the vector of model parameters, i represents the dose group and j represents the patient within dose group i. Let $\mu = \mu_{d_1}, \ldots, \mu_{d_k}$ denote the mean dose response vector.

5.1.2.2 MCP-Mod Test (MCP)

Instead of prespecifying a single dose-response model, MCP-Mod uses a set of candidate models covering a suitable range of dose-response shapes. Each of the models in the candidate set is tested using appropriate contrasts and employing an MCP that preserves the FWER. PoC is established if at least one of the model tests is significant. Once PoC is verified, the "best" model(s) among the candidate set is chosen to fit the data and to produce estimated doses using modeling techniques, while still ensuring the appropriate FWER. The "best" model can be evaluated using a model fit statistic such as the lowest MCP p-value or the AIC.

The procedure requires that a set of M candidate models be chosen. Let $f_m(d, \theta_m)$ and $f_m^0(d, \theta_m^0), m = 1, \ldots, M$ denote the model and standardized model functions respectively, where θ_m and θ_m^0 are the parameters of model m under the unstandardized and standardized model functions respectively.

The hypothesis of interest for each of the dose-response models is $H_0^m : c_m^{'}\mu = 0$, where $c_m = (c_m 1, \ldots, c_m k)$ is the optimal contrast vector for model m, and $\sum_{i=1}^k c_m i = 0$. Therefore each candidate models is tested using the following contrast test,

$$T_m = \frac{\sum_{i=1}^k c_{mi}\bar{Y}_i}{S\sqrt{\sum_{i=1}^k c_{mi}^2/n_i}}, m = 1, \ldots, M.$$

where $\bar{Y}_i = \sum_{i=1}^k Y_{ij}/n_i$ is the arithmetic mean response in dose group i, $S^2 = \sum_{i=1}^k \sum_{j=1}^{n_i} (Y_{ij} - \bar{Y}_i)^2/v$ is the pooled variance estimator, with $v = \sum_i n_i - k$ degrees of freedom.

Each contrast test can thus be considered as a decision procedure determining whether the given dose-response shape is statistically significant, based on the observed data. The contrast coefficients, c_m are optimal in the sense that they maximize the power to detect the underlying model.

The final detection of a significant dose-response signal (i.e demonstrating PoC), is based on the maximum contrast test statistic, $T_{max} = maxT_1, \ldots, T_M$. Under the null hypothesis of no dose effect the mean response at each dose should be equal. Under this assumption and the dose-response model formulation, it follows that the contrast test statistics follow a multivariate t-distribution with v degrees of freedom. PoC is hence established if

$$T_{max} \geq q_{1-\alpha}$$

where $q_{1-\alpha}$ is the α significance level multiplicity adjusted critical value (i.e., the relevant quantile of the central multivariate t distribution).

5.1.2.3 MCP-Mod Power and Sample Size

Sample size is based on MCP contrast test for the models in the candidate set outlined in subsubsection 5.1.2.2. This method is outlined in Pinheiro et al (2006). Under this sample size procedure, the power of the MCP procedure is determined by the distribution of the maximum contrast test statistic under the alternative hypothesis that the m^{th} dose-response model is the true response.

Under the alternative hypothesis, these contrast test statistics have a non-central multivariate t-distribution. Given the contrast test outlined previously the power for each candidate model m in MCP procedure is defined as

$$(1-\beta)_m = P\left(max_l T_l \geq q_{1-\alpha} | \mu = \mu_m\right) = 1 - P\left(T_1 < q_{1-\alpha} \ldots T_m < q_{1-\alpha} | \mu = \mu_m\right)$$

For sample size determination, it is useful to consider the vector of power values per candidate model. It follows from the model specification and the properties of the multivariate t distribution that the contrast test statistics T_l are jointly distributed as non-central multivariate t distribution with $k(n-1)$ degrees of freedom, the correlation matrix derived from the contrast coefficients (see Pinehiero et al 2006) and per model non-centrality parameters of $\delta_{ml} = \sqrt{n} c_l' \mu_m / \sigma$.

The calculation of quantiles of the non-central multivariate t-distribution is required and in nQuery the calculation of these quantiles is done using the Separation-of-Variables Method as discussed in Genz and Bretz (2002).

For practical purposes, a single summary measure of power is needed for sample size determination. Multiple options are available but the chosen function should be monotonically increasing. In nQuery, three choices of summary function are given. These are the Mean (unweighted), Minimum, and Median. nQuery provides the per-model power so the manual calculation of other summary functions is possible.

5.1.2.4 Dose-Response Models

The list of available dose-response models is given below. Note that most of these models require the pre-specification by the user of 1-3 additional parameters to fully define the model. The definition of each model including these additional parameters is given below with additional parameters given in brackets with model name e.g. **Emax*(ED50)*** indicates ED50 is required unique input for Emax model. These additional parameters alongside the placebo effect, maximum effect, standard deviation and specified dose-levels given by the user elsewhere (see section 5.2) define the model effects.

Note some common terminology including d (dose response at specific dose), E_0 (placebo dose response for most models, basal effect for logistic model), E_{max} (maximum dose response, for models other than the Beta and Quadratic models this will equal response at last dose). $f(d, \theta)$ is the unstandardized model and $f^0(d, \theta^0)$ is standardized model.

- **Linear**: This model assumes a positive linear relationship between E_0 and E_{max} on the original dose scale (i.e. between response at placebo and largest dose).δ is the linear slope parameter and this is derived automatically from the model definitions via simultaneous equations.

$$f(d, \theta) = E_0 + \delta d$$

$$f^0\left(d, \theta^0\right) = d$$

- **Linear Log-Dose (*Off*):** This model assumes a positive linear relationship between E_0 and E_{max} on the log-dose scale (i.e. between response at placebo and largest dose). The Off parameter is provided to prevent issues with doses equal to zero. This is typically set to a small value such as 0.01 times the maximum dose. δ is the linear slope parameter and this is derived automatically from the model definitions via simultaenous equations.

$$f(d, \theta) = E_0 + \delta \log(d + Off)$$

$$f^0\left(d, \theta^0\right) = \log(d + Off)$$

- **Emax (*ED50*)**: A positive monotonic concave dose-response curve model. Also known as the hyperbolic Emax model, in contrast to the more flexible sigmoidal Emax model (see below). ED_{50} is the expected dose at which we would expect 50% the maximum dose response to occur.

$$f(d, \theta) = E_0 + E_{max}\frac{d}{ED_{50} + d}$$

$$f^0\left(d, \theta^0\right) = \frac{d}{ED_{50} + d}$$

- **Sigmoidal Emax (*ED50, h*)**: A positive monotonic concave dose-response curve model. It is similar to the Emax model (see above) but has an additional parameter, h, (often called the Hill parameter/coefficient) which characterizes the slope of the dose-response curve at the ED_{50} dose. A Hill parameter value of one is equivalent to the (hyperbolic) Emax model with values lower than one implying a shallower slope than the Emax model and values higher than one implying a steeper slope than the Emax model.

$$f\left(d,\theta\right)=E_0+E_{max}\frac{d^h}{ED_{50}+d^h}$$

$$f^0\left(d,\theta^0\right)=\frac{d^h}{ED_{50}+d^h}$$

- **Logistic (*ED50, δ*)**: A positive monotonic concave dose-response curve model. It is closely related to the sigmoidal Emax model which is equivalent to a logistic model on the log(dose) scale, with this Logisitic model being on the original dose scale. It requires the specification of ED50 (the dose giving half the maximum dose response effect) and δ which controls for the slope of the curve. Higher δ imply a steeper dose-response curve.

$$f\left(d,\theta\right)=E_0+E_{max}/\left\{1+\exp\left[\left(ED_{50}-d\right)/\delta\right]\right\}$$

$$f^0\left(d,\theta^0\right)=1/\left\{1+\exp\left[\left(ED_{50}-d\right)/\delta\right]\right\}$$

- **Exponential (*δ*)**: A positive monotonic convex dose-response curve model. It requires the specification of the parameter δ which controls the convexity of the model with higher values implying a faster increase in response with increase in dose. E_1 is a slope parameter and this is derived automatically from the model definitions via simultaenous equations.

$$f\left(d,\theta\right)=E_0+E_1\left(\exp\left(\frac{d}{\delta}\right)-1\right)$$

$$f^0\left(d,\theta^0\right)=\exp\left(\frac{d}{\delta}\right)-1$$

- **Beta (*a, b, Scale*)**: A flexible model that can be used to model non-monotonic dose-response curves. It requires the specification of three parameters. "a" and "b" are the first and second shape parameters for the Beta model and the distributional assumptions of the Beta distribution in terms of shape can be applied here. The scale parameter is used to characterise the range of the Beta function (e.g. scale of 100 would re-scale Beta Distribution to [0,100] rather than the standard [0,1]). The scale parameter should be greater than the largest dose with a value of 1.2 times the maximum dose

being a useful default.

$$f\left(d,\theta\right)=E_0+E_{max}B\left(a,b\right)\left(\frac{d}{Scale}\right)^a\left(1-\frac{d}{Scale}\right)^b,B\left(a,b\right)=\frac{\left(a+b\right)^{a+b}}{a^ab^b}$$

$$f^0\left(d,\theta^0\right)=B\left(a,b\right)\left(\frac{d}{Scale}\right)^a\left(1-\frac{d}{Scale}\right)^b$$

- **Quadratic (δ)**: A flexible model which assumes a quadratic dose-response relationship and can be used to model non-monotonic dose-response curves. It requires only the specification of the parameter δ. δ represents a ratio of the quadratic model parameters, $\beta_1$and β_2,and is defined below. Note that the concave version of the quadratic model has been implemented in nQuery which imposes the restriction that $\beta_2$is negative. Note that, through the formulation below, specifying a negative value for δ will result in a positive value being found for $\beta_1$and vice versa.

$$f\left(d,\theta\right)=E_0+\beta_1d+\beta_2d^2$$

$$f^0\left(d,\theta^0\right)=d+\delta d^2,\delta=\frac{\beta_2}{|\beta_1|}\&\beta_2<0$$

β_2 here represents the curvature of the dose-response relationship and $\beta_1$represents the slope of the curve. The values for the$\beta_1$and $\beta_2$parameters of the quadratic model are found through the following equations:

$$\beta_2=-|4\delta^2(ME-PE)|$$

$$\beta_1=\frac{\beta_2}{\delta}$$

where ME represents the maximum response effect expected to be observed in the dose-range and PE represents the effect at placebo.

5.2 MCP-Mod Demonstration

5.2.1 Background

In nQuery the layout for MCP-Mod is similar to that of most other tables in nQuery. See chapter 1 for an introduction to using nQuery. This section will focus on the additional issues associated with the MCP-Mod table by demonstrating a detailed example of using the table in nQuery.

In nQuery, the inputs required for the MCP-Mod design can be split into three separate categories.

The first of these categories consists of the general test design parameters such as power and the placebo effect. These are equivalent to the main table inputs seen in standard nQuery tables.

The other two categories are the dose levels being assessed in the study and the candidate models selected and their per-model inputs (see subsubsection 5.1.2.4) and outputs e.g. power. These are both contained in the mandatory side-table. See section 3.1 for general guidance on side-tables.

In this section we will consider an example which relates to a real Phase II design in the development program of a drug for the indication of generalized anxiety disorder (GAD). Five active doses are to be used in the study: 10, 25, 50, 100, and 150mg, with an additional placebo arm (corresponding to a 0mg dose). This is the same example considered in section 2 of Pinheiro et al. (2006). Further details will be provided as we complete the example.

5.2.2 Main Table (General Design Parameters)

The main table (as discussed in previous chapters of this manual) is used to enter the common parameters which apply to each model. These are the test significance level, whether to use a one or two sided test, the placebo effect, the maximum treatment effect and common standard deviation (equivalent to the residual standard deviation in our model and practically the standard deviation of responses within a dose-level).

Alonside these common parameters, the number of doses and number of models need to be specified. This will set the number of columns in the doses input table and candidates models specificiation table contained in the side table. See subsection 5.2.3 for details.

Figure 5.1 illustrates how the main table will appear on the initial opening of the table in nQuery.

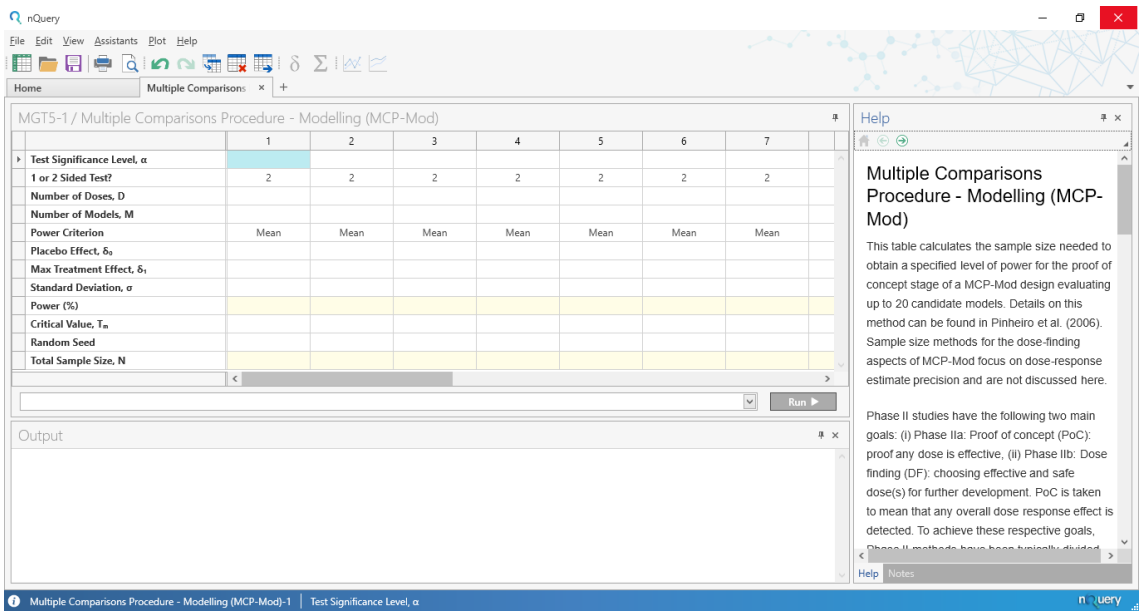

Figure 5.1: MCP-Mod Design Main Table

The first step required is to specify the desired significance level (α), whether a one or two sided test is required and the number of doses and candidate models of interest. Select the power criteria (mean, minimum or median power) and then specify the expected placebo effect, the maximum effect, and the common standard deviation. If one wishes to solve for the total sample size then specify the target power. If one wishes to solve for power, then specify the total sample size here or specify the individual sample sizes per dose group in the Dose Level Side table, as discussed in section 5.2.3.1 below.

In this example, considered here we will assume a significance level of 0.05, 80% power, a residual standard deviation of 1, a one sided test, a placebo Effect of 0, a maximum treatment effect of 0.4, 6 doses (5 doses and the placebo dose), a candidate set of 6 models and use the "Mean" power criteria. The power value will be left empty for now and entered after the side-table section is complete.

Optionally, the user can also enter a random seed and critical value. The random seed is used for the simulations used for the multivariate t-distribution method implemented in nQuery. If left blank this will default to a value based on the system time. Setting this option to a specific value allows the user to generate the exact same results for a given set of inputs. In this example, we will set the seed to 1234 to allow exact replication of this example.

The critical value is the critical multivariate t-statistic for the multivariate t-distribution CDF ($q_{1-\alpha}$). To use the optimal t-statistic set the test significance level to desired level and this will make the Critical Value read-only. To use a custom test statistic do not enter the significance level and enter the critical value. As we have entered the significance level, the optimal contrasts will be used here.

For more information on these parameters, see the help cards in the software. The Critical Value parameter will be calculated automatically for us. The completed main table is shown in Figure 5.2.

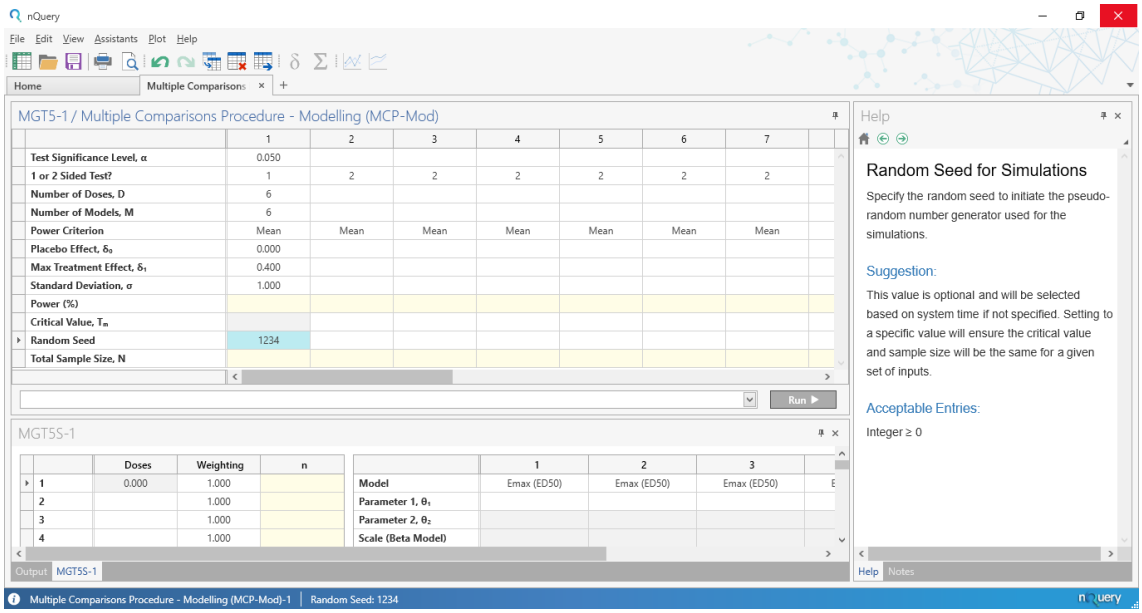

Figure 5.2: Completed MCP-Mod Main Table

5.2.3 Side-Table (Dose Levels and Candidate Models)

After the number of doses and number of models is specified in a column, the MCP-Mod side-table will automatically appear below the main table window. See subsection 5.2.3 for general guidance on side-tables. In this example, upon opening the side-table will be as per Figure 5.3.

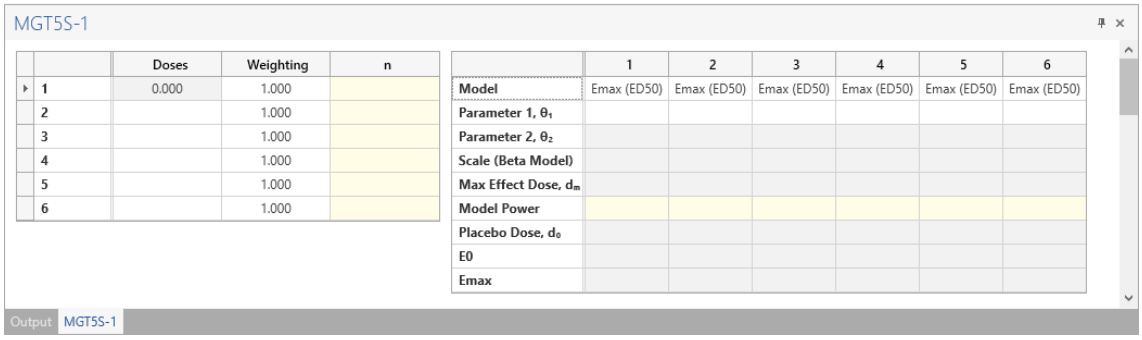

Figure 5.3: MCP-Mod Side-Table

The table on the left-hand side is the Doses Table where the dose levels of interest

will be entered. On the right-hand side is the Model Specification Table where candidate models are selected and model specific inputs and outputs are provided.

5.2.3.1 Doses Table

On the left-hand side of the side-table, the planned dose levels table is provided as per Figure 5.4. This table specifies which doses will be considered in this trial and sample size weightings applied to each dose level.

		Doses	Weighting	n
▸	1	0.000	1.000	
	2		1.000	
	3		1.000	
	4		1.000	
	5		1.000	
	6		1.000	

Figure 5.4: Dose Levels Table

In this table, enter the proposed study dose levels (how much/concentration of the proposed treatment given to each dose level) in the first "Doses" column. Note that the first row entry in the Doses column is automatically set to 0 as this represents the placebo dose (i.e. no treatment given) which is required for this method. It is recommended that doses are entered in increasing order down the column but nQuery will automatically adjust if a non-increasing set of dose levels is given. In this example we will assume the following set of doses: 0, 10, 25, 50, 100, 150 (mg).

Next, enter the sample size weightings per dose level in the second "Weighting" column. By default all entries in this column will equal one, corresponding to equal sample size per dose level. The proportion of the total sample size in a given dose level is given by $W_i / \sum W_i$ where W_i is the weighting per dose level.

If solving for power, one may enter the individual dose group sample sizes in the third column. Otherwise, once a sample size determination is complete the third column, "n", will automatically be filled with the per-dose sample sizes. For this example, assume equal sample size per column and leave the default "1's" inputs. The final dose levels tables will be filled as per Figure 5.5.

		Doses	Weighting	n
	1	0.000	1.000	
	2	10.000	1.000	
	3	25.000	1.000	
	4	50.000	1.000	
	5	100.000	1.000	
✎	6	150.	1.000	

Figure 5.5: Completed MCP-Mod Dose Levels Table

5.2.3.2 Model Specification Table

On the right-hand side of the side-table, the model specification table is provided as per Figure 5.6. This table specifies which candidate set of models will be considered and where any additional parameters required per-model are inputted. After a sample size determination, it contains the per-model outputs including model power, minimum effect dose level (d_0), maximum effect dose level (d_m), the minimum dose effect (E_0) and maximum dose effect (E_{max}). Note that the maximum and minimum dose levels will automatically be calculated for most models after the Dose Levels table is filled. The major exceptions are the Beta model and the Quadratic model where the maximum effect dose level is not known pre-calculation.

	1	2	3	4	5	6
Model	Emax (ED50)	Emax (ED50)	Emax (ED50)	Emax (ED50)	Emax (ED50)	Emax (ED50)
Parameter 1, θ_1						
Parameter 2, θ_2						
Scale (Beta Model)						
Max Effect Dose, d_m	150.000	150.000	150.000	150.000	150.000	150.000
Model Power						
Placebo Dose, d_0	0.000	0.000	0.000	0.000	0.000	0.000
E0						
Emax						

Figure 5.6: Model Specification Table

In this table, the desired candidate models are selected in each column using the dropdown menu in the Model row of the model specification table. After a model is selected, 0 to 3 (equal to the number of parameters given in brackets beside model name) of the rows will become available to edit depending on the model selected. Column cells that require inputs for calculations will be white and those not required

for the given model will be grey. Note that the same model can be selected multiple times, characterized with different model specific parameters, except for the Linear model which has no additional inputs. Details on the models available and the model specific parameters required are given in subsubsection 5.1.2.4.

In this example we will consider the following set of candidate models, with the corresponding parameter values assumed for each model given in brackets after each model:

- Linear
- Emax (ED50 = 25)
- Logistic (ED50 = 50, δ = 10.88111)
- Exponential (δ = 85)
- Beta (a = 0.33, b = 2.31, Scale = 200)
- Beta (a = 1.39, b = 1.39, Scale = 200)

The completed model specification table is given in Figure 5.7.

	1	2	3	4	5	6
Model	Linear	Emax (ED50)	Logistic (ED50, δ)	Exponential (δ)	Beta (a, b, Scale)	Beta (a, b, Scale)
Parameter 1, θ_1		25.000	50.000	85.000	0.330	1.390
Parameter 2, θ_2			10.881		2.310	1.390
Scale (Beta Model)					200.000	200.000
Max Effect Dose, d_m	150.000	150.000	150.000	150.000	150.000	150.000
Model Power						
Placebo Dose, d_0	0.000	0.000	0.000	0.000	0.000	0.000
E0						
Emax						

Figure 5.7: Completed Model Specification Table Example

In this example, after the final parameter which is required in the model specification table is entered (e.g. "Scale" in column 6), the calculation will not run as the Power row has not been filled in the main table. However, if the main table had been fully filled the calculation would occur automatically. Note that if you do not want every edit to this side-table to automatically start the calculation, leaving one or more of the mandatory inputs in the main table empty is recommended. Once a calculation is complete the per-model power, placebo dose, maximum effect dose, E0 and Emax will be calculated automatically in this table. Note that some models will also calculate additional parameters such as the slope parameter for the linear and exponential models.

5.2.4 Results

In this example, the last input is 80% in the Power row in the main table After this is entered, the calculation will start. After the calculation is complete the entire table in nQuery will be as per Figure 5.8. Note that this method (due to the computation of multivariate t-distribution) may take some time to run.

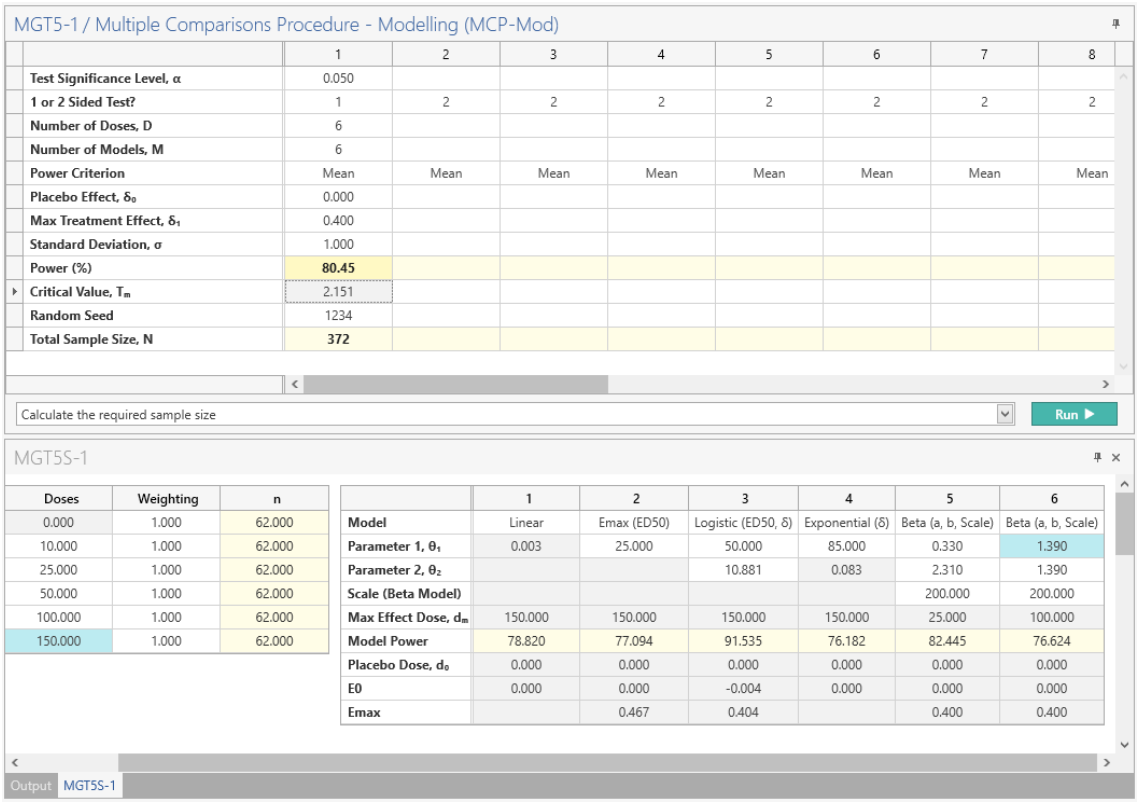

Figure 5.8: Completed MCP-Mod Example

This example results in a sample size of 372 and thus a per dose group sample size of 62 as per [Pinheiro et al., 2006].The exact mean power is automatically updated to 80.44% for this sample size. The per-model powers and other parameters are given in the model specification table in the side-table. The per-model powers ranged from 76.139% for the exponential model to 91.522% for the logistic model. Note the maximum effect dose is the final dose for all models except the Beta models, with the first Beta model having its maximum effect (i.e. the maximum dose effect specified in the main table) at the 25mg dose and the second Beta model having its maximum effect at the 100mg dose. These are found automatically by evaluating the Beta model at each specified dose level and finding the maximum effect dose from the results.

Using a custom critical value will have a shorter calculation time so the optimal

critical value can be used in other examples directly if the significance level and model specification were unchanged.

6 nQuery Qualification Tools

nQuery provides automated qualification scripts to quickly allow a user to verify that their nQuery application is installed and operating to the manufacturers specifications. nQuery specifically has two separate tools for installation qualification (IQ) and operation qualification (OQ) to assist you in verifying nQuery in regulatory industries.

Note that operational and performance qualification (PQ) are considered interchangeable based on our correspondence with the relevant stakeholders. However, you may wish to perform within-application testing for PQ purposes.

6.1 Installation Qualification (IQ) Tool

The nQuery IQ tool assists you in demonstrating that nQuery has been installed and maintained to the manufacturer's specifications. nQuery IQ verifies the integrity of each file in the nQuery 8 system and provides the user a set of reports detailing the results.

The nQuery Installation Qualification (IQ) Tool validates an nQuery installation by verifying that each installed file is correct with a report generated detailing all file results. This determination is made using the SHA-1 algorithm to create a hash value for each file. This thus checks that each file is present and that its integrity has been maintained.

Important: IQ validation represents the expected state of the system as it is at the time an installation is completed. If changes are made to the original files subsequently then these may show as failed when IQ is run on the system at a later date.

6.1.1 Running nQuery IQ

To run nQuery IQ, select "Installation Qualification" from the Help file menu. This will automatically run the nQuery.Tools.InstallationQualification.exe application (found in your installation folder, default of C:/Program Files (x86)/Statistical Solutions Ltd/nQuery) and output a HTML report which will display the results automatically in the machine's default HTML viewer (e.g. default internet browser).

6.1.2 nQuery IQ Results

The IQ report file will automatically be saved on the machine when the IQ tool runs. It will be saved to the C:\Users\<User Account>\AppData\Local\nQuery\Reports folder where <User Account> is the named User which currently logged into on the machine. An example of an IQ report is shown in Figure 6.1.

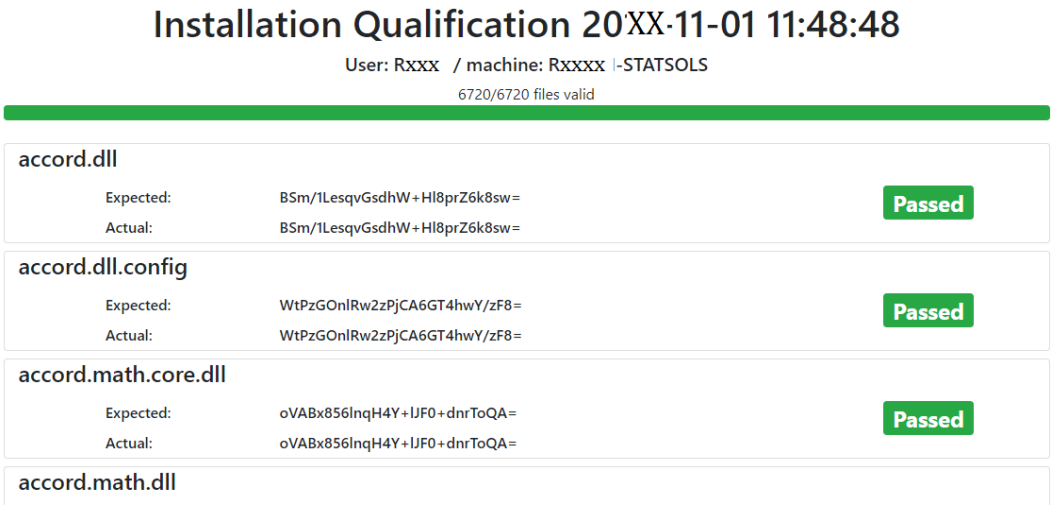

Figure 6.1: Installation Qualification Report Example

6.2 Operational Qualification (OQ) Tool

The nQuery OQ tool assists you in demonstrating that nQuery 8 is operational. nQuery 8 uses the nQuery solver functions over each design table for a set of validated design parameter inputs and will execute, process, and report the solver results.

Note that the files used for OQ are available in the TestData folder in your installation folder (default is C:/Program Files (x86)/Statistical Solutions Ltd/nQuery)

Important: All nQuery results have been validated exactly against the results from nQuery + nTerim 4.0 and the original validation documentation and references. All results should be identical to nQuery + nTerim 4.0 except for simulation tables. There are slight difference with simulation tables due to an upgraded random number generator being used in nQuery. For those tables a 5% precision rule was used for validation. Results were also compared to EAST 6.5, PASS 2019 and SAS 9.5 Proc Power where the same or similar methods were implemented in those software for a design table.

6.2.1 Running nQuery OQ

To run nQuery OQ, select "Operational Qualification" from the Help file menu. This will open the OQ Validation Tool menu which is shown in Figure 6.2.

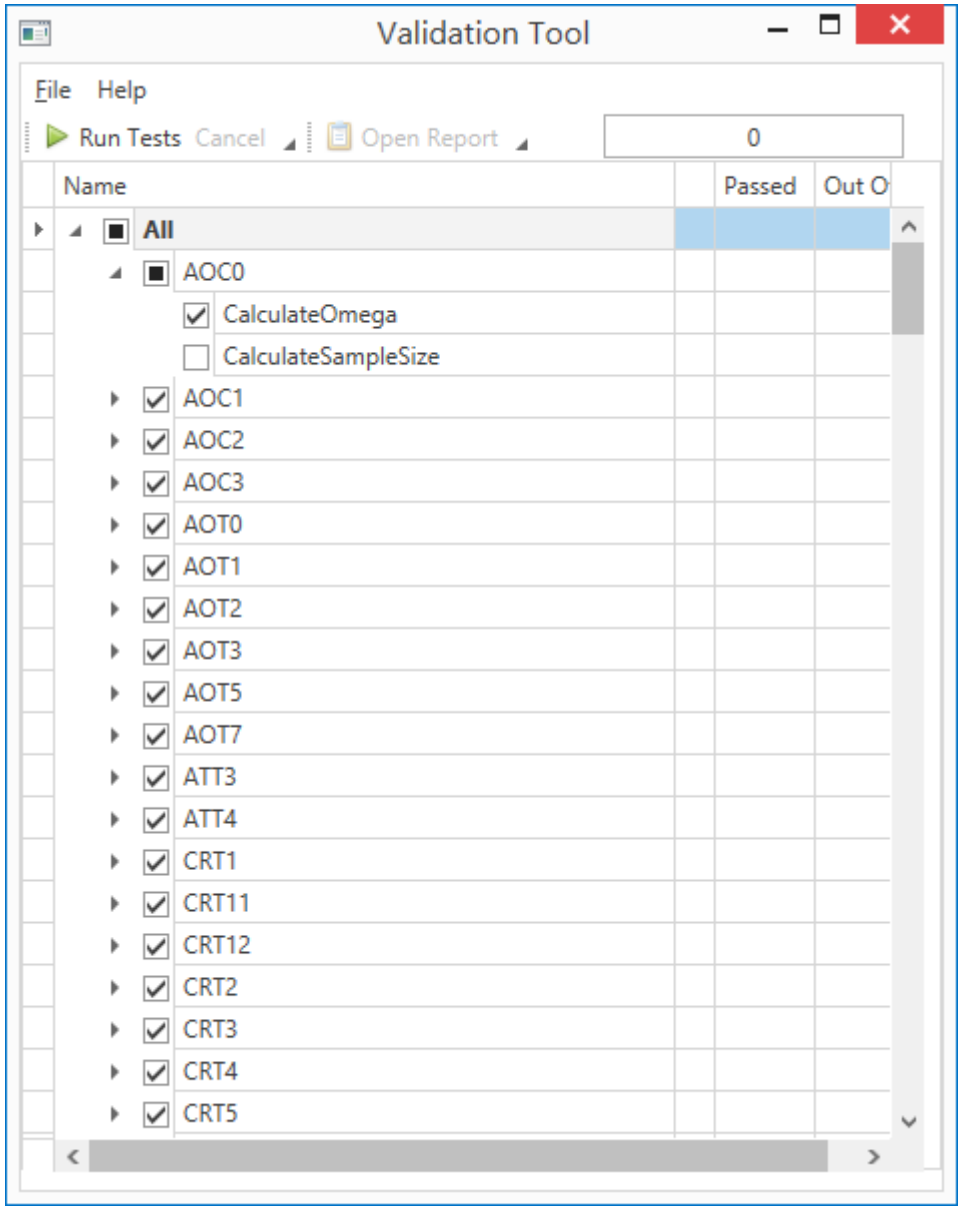

Figure 6.2: Operational Qualification Tool

The nQuery OQ window provides a list of all of the nQuery design tables by code (AOT1, MTE1 etc.) and has a three level menu hierarchy: All Tables, Per Table, Per Solver. Selecting the arrow ▶ to the left of the "All" option at the top of the menu will open and collapse the per-table options. Selecting the arrow ▶ to the left of a table code option will open and collapse the solver options for that table.

For all options, the user can include or exclude a table or solver by selecting or de-selecting the check-box to the left of that option. Selection is indicated with a tick icon.

Note that the tables selected by default will depend on when the OQ tool is opened. If the OQ tool is opened on the Getting Started tab or when no tab is open then all tables will be selected by default. If the OQ tool is opened on a specific design table then only that table will be selected by default.

Note that a small number of shortcuts are available for the OQ tool. Select the Help file menu item to see these.

Once the user has selected the tables which will be qualified, select the "Run Tests" button in the window menu bar. Note that since the nQuery testing data is specifically designed to include extreme values to test robustness that some tests may take some time to run. If the user wishes to stop qualification before it is finished, select the "Cancel" button in the menu bar.

As the OQ tool progresses, each item in the menu will show the Qualification results per-feature as it is completed in the columns to the right of the code-name.

The first column will show a tick symbol if the feature has passed and an "x" symbol if it has failed. The next "Passed" column will show the number of passed individual table/solver tests. The final "Out of" column will display the total number of tests ran for that individual table/solver tests. If a table/solver passes then the tests passed should equal the total number of tests.

A progress bar in the menu bar will display the total number of tables tests and indicate overall progress by the amount of the bar filled.

An example of the OQ tool in progress is shown in Figure 6.3.

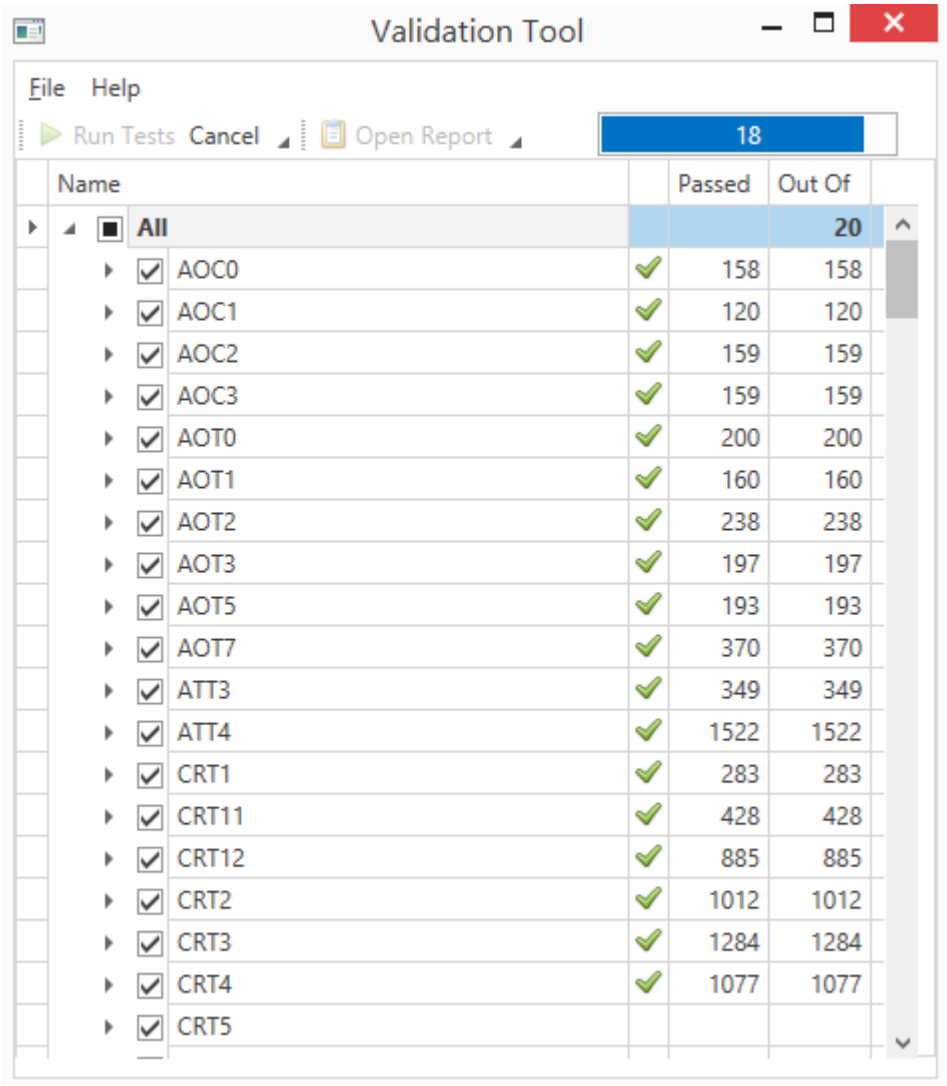

Figure 6.3: Operational Qualification Tool - In Progress

6.2.2 nQuery OQ Results

To view the full report of the OQ results, select the "Open Report" option in the menu bar. This will automatically output a HTML report which will display the results automatically in the machine's default HTML viewer. This will often be the users default internet browser.

To OQ report file will automatically be saved on the machine when a "Open Report" is used. It will be saved to the C:\Users\<User Account>\AppData\Local\nQuery\Reports folder where <User Account> is the named User which currently logged into on the machine. Alternatively, you can save a report to an alternative folder using the "Save as" option in the File menu of the OQ window.

If OQ report is opened in a supported application, the structure of the report will be as follows:

At the start, the report will contain the time and date the tests were run, the number of passed design table tests and the time taken for testing.

Within the report, there will be an individual report for each design table test run. These will show the table code, table name, an indication whether the table passed overall, the time taken to test that table and the number of passed tests for that table.

Within each design table report, there will be a report for each solver in that design table. This will show the solver name, an indication whether the table solver passed overall, the time taken to test that table and the number of passed table solver tests. When shown within a browser, it allows a user to see the individual design parameters used for each test applied to a solver. To open these, select the "Details" option to the right of the solver name. An example an OQ report with a Details menu open is shown in Figure 6.4.

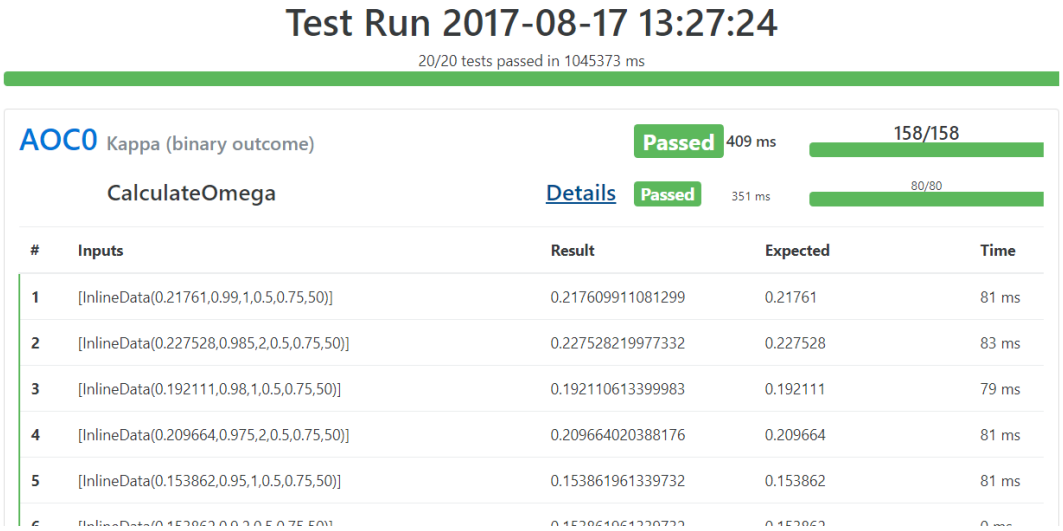

Figure 6.4: Operational Qualification Tool Report

7 nQuery Updates and Licensing

7.1 Renewing a License

Important: Internet Access is required to renew nQuery. If you need to renew offline, please contact support@statsols.com or login to your online account

To renew an nQuery License, select the Activate/Renew License option from the Help menu. This will open the Product Activation dialog. Enter your provided nQuery Activation Key in the Activation Key field. This is shown in Figure 7.1.

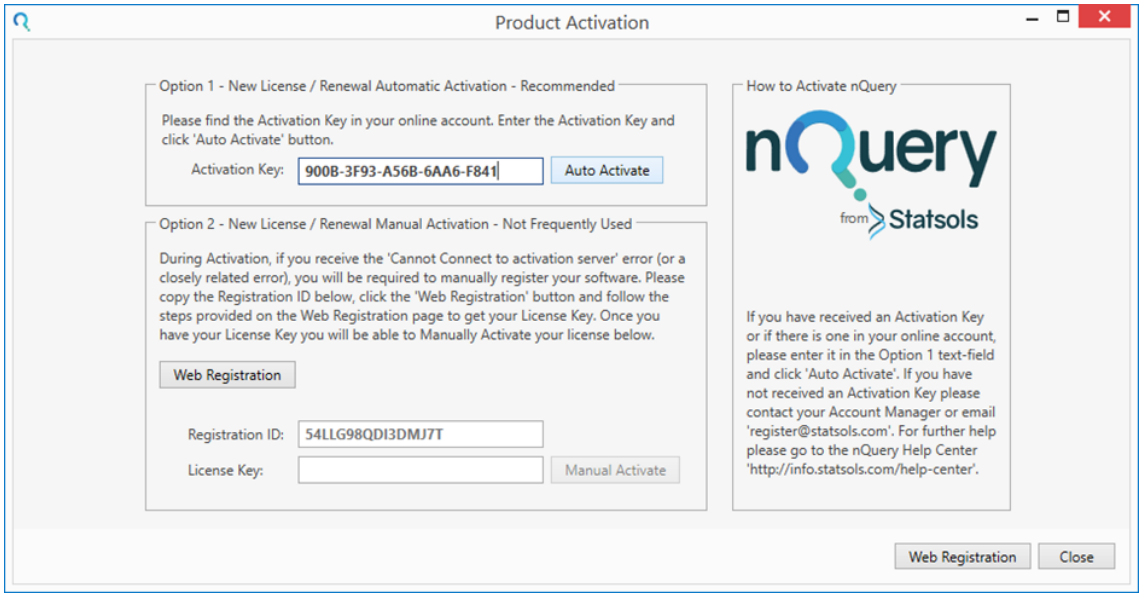

Figure 7.1: Renewal Example

If you have issues you can contact support@statsols.com or login to your online account and create a support ticket for fastest response.

7.2 Activating an Add-on Module

Important: Internet Access is required to add an Add-on module nQuery. If you need to add a module offline, please contact support@statsols.com or login to your online account

To purchase an add-on module, either use the link provided in the Getting Started screen or visit the Statsols website at www.statsols.com

To activate an add-on module, select the Enable Modules option in the Help file menu. The Enable Modules dialog box should open. This is shown in Figure 7.2.

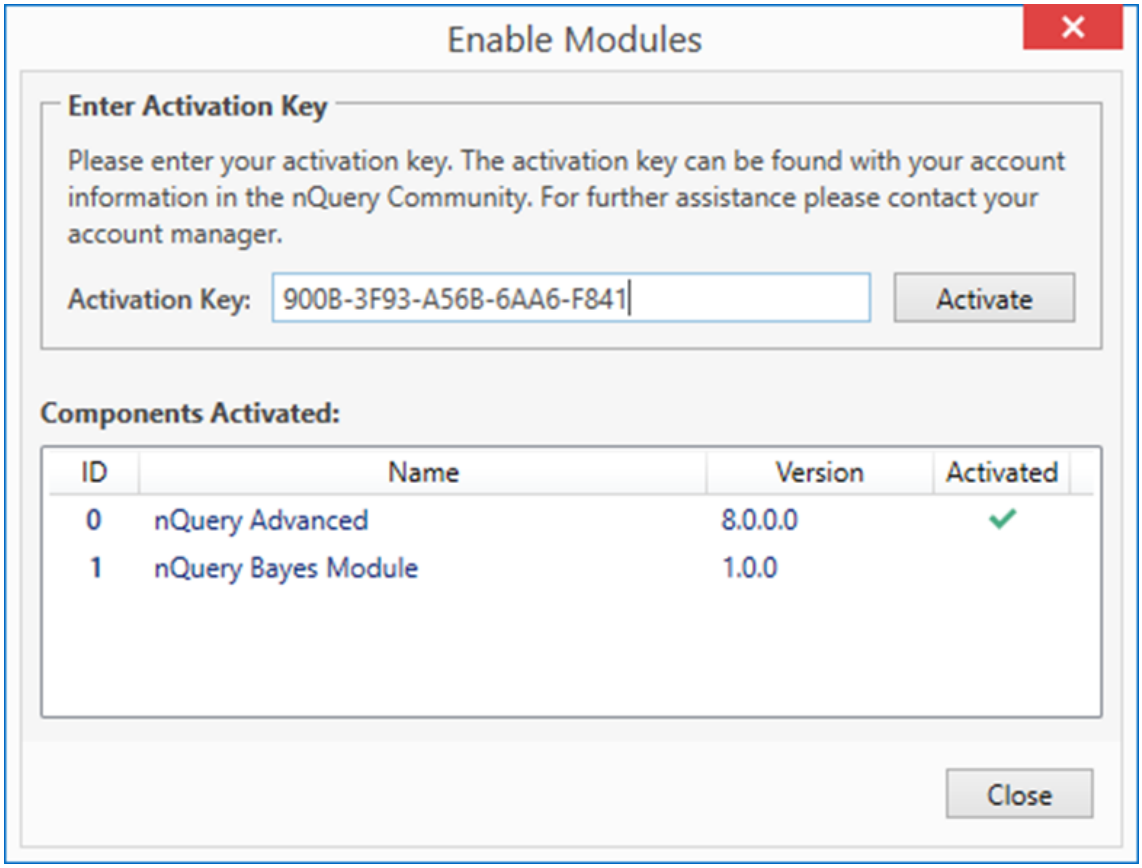

Figure 7.2: Module Activation Screen

If you have used a valid nQuery Activation Key in nQuery before, this key will appear in the Activation Key area. Otherwise, you will need to enter the key provided by Statsols.

To activate the module, click "Activate". The Components Activated area will refresh after a couple of seconds and the nQuery Module(s) should have a tick in the Activated column for the selected module(s) .

All tables in the module(s) will be enabled in nQuery Advanced

If you have issues you can contact support@statsols.com or login to your online account and create a support ticket for fastest response.

7.3 Checking for Software Updates

Important: Internet Access is required to update nQuery. If you want to manually update, please contact support@statsols.com or login to your online account

7.3.1 Checking for Updates

nQuery provides automated tools to allow the user to upgrade their nQuery application to the latest version. There are two methods to update your nQuery application: using the Help menu option or using the system tray option.

To select the Help menu option, open the Help file menu and select "Check for Updates". To use the system tray option, find the Update Statsols - nQuery system tray item, right-click and select "Open Updater".

If there are no updates, these will open the Update Statsols - nQuery dialog which will say that no upgrades are available.

If updates are available, a prompt will open asking the user if they wish to upgrade. An example of this prompt is shown in Figure 7.3.

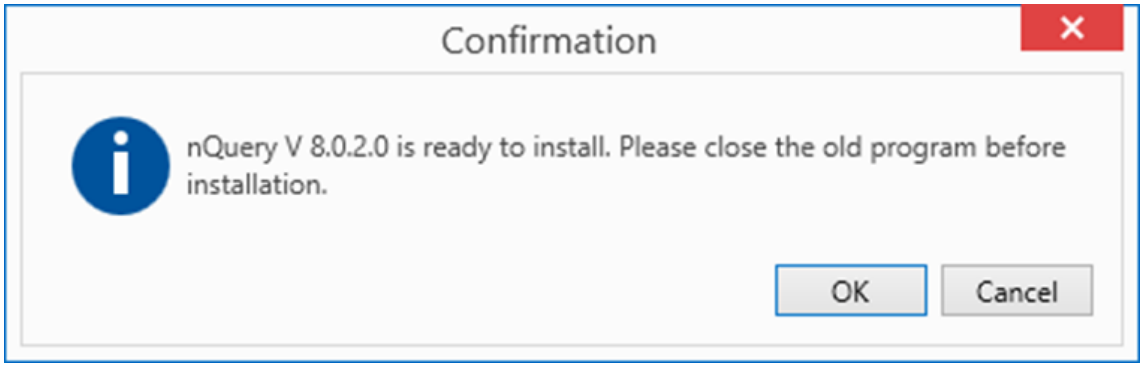

Figure 7.3: Update Prompt Example

Select OK to install the update or Cancel to close this dialog. If you select OK, the Update Statsols - nQuery dialog will appear and show the download progress. When downloading is completed, the Figure 7.3 dialog will appear again. Select OK and this will close any currently open versions of nQuery. An install screen will automatically appear and update your nQuery application. On completion, your nQuery application will be updated and can be used again.

Bibliography

[Bretz et al., 2005] Bretz, F., Pinheiro, J. C., and Branson, M. (2005). Combining multiple comparisons and modeling techniques in dose-response studies. *Biometrics*, 61(3):738–748.

[Chen et al., 2004] Chen, Y. J., DeMets, D. L., and Gordon Lan, K. (2004). Increasing the sample size when the unblinded interim result is promising. *Statistics in medicine*, 23(7):1023–1038.

[Cohen, 1988] Cohen, J. (1988). Statistical power analysis for the behavioral sciences. *NJ: Lawrence Earlbaum Associates*, 2.

[Cui et al., 1999] Cui, L., Hung, H. J., and Wang, S.-J. (1999). Modification of sample size in group sequential clinical trials. *Biometrics*, 55(3):853–857.

[Demets and Lan, 1984] Demets, D. L. and Lan, G. K. (1984). An overview of sequential methods and their application in clinical trials. *Communications in Statistics-Theory and Methods*, 13(19):2315–2338.

[Demets and Lan, 1994] Demets, D. L. and Lan, K. G. (1994). Interim analysis: the alpha spending function approach. *Statistics in medicine*, 13(13-14):1341–1352.

[Gao et al., 2008] Gao, P., Ware, J. H., and Mehta, C. (2008). Sample size re-estimation for adaptive sequential design in clinical trials. *Journal of Biopharmaceutical Statistics*, 18(6):1184–1196.

[Hwang et al., 1990] Hwang, I. K., Shih, W. J., and De Cani, J. S. (1990). Group sequential designs using a family of type i error probability spending functions. *Statistics in medicine*, 9(12):1439–1445.

[Jennison and Turnbull, 1999] Jennison, C. and Turnbull, B. W. (1999). *Group sequential methods with applications to clinical trials*. Chapman and Hall/CRC.

[Mehta and Pocock, 2011] Mehta, C. R. and Pocock, S. J. (2011). Adaptive increase in sample size when interim results are promising: a practical guide with examples. *Statistics in medicine*, 30(28):3267–3284.

[O'Brien and Fleming, 1979] O'Brien, P. C. and Fleming, T. R. (1979). A multiple testing procedure for clinical trials. *Biometrics*, pages 549–556.

[Pampallona et al., 1995] Pampallona, S., Tsiatis, A., and Kim, K. (1995). Spending functions for the type i and type ii error probabilities of group sequential tests. *J Statist Plan Inference*, 42(19):1994–35.

[Pampallona et al., 2001] Pampallona, S., Tsiatis, A. A., and Kim, K. (2001). Interim monitoring of group sequential trials using spending functions for the type i and type ii error probabilities. *Drug Information Journal*, 35(4):1113–1121.

[Pinheiro et al., 2006] Pinheiro, J., Bornkamp, B., and Bretz, F. (2006). Design and analysis of dose-finding studies combining multiple comparisons and modeling procedures. *Journal of biopharmaceutical statistics*, 16(5):639–656.

[Pocock, 1977] Pocock, S. J. (1977). Group sequential methods in the design and analysis of clinical trials. *Biometrika*, 64(2):191–199.

[Tukey, 1985] Tukey, J. (1985). The problem of multiple comparisons. *The Collected Works of John W. Tukey*, 2:1965–1984.