

The G(roup)Rumbler: Instructions for Users

Release Version: November 2023, with TurboGrumbler option

Professor Malcolm K. Sparrow
John F. Kennedy School of Government, Harvard University

(Last Revised: 3rd November 2023)

1.0: Introduction: About The GRumbler:¹

1.1: Purpose: The GRumbler is designed as a tool to help Course Administrators divide classes of students into small groups for the purposes of group discussions, exercises, or project work. The GRumbler assigns students to groups in a way that maximizes the mixing across the class and provides the greatest possible degree of diversity within each group.

The GRumbler can also generate *sequences of group assignments*—for use on different days, or for different exercises—that have minimal overlap with one another (i.e. that keep students apart who have been together in any group before).

Using a GRumbler-generated sequence of group assignments for a succession of different exercises should therefore increase social cohesion within the class, allowing students to meet a greater number and a more diverse set of their classmates through a sequence of small-group engagements.

The people who will appreciate the GRumbler most of all will be those who previously spent many frustrating hours shuffling color-coded index cards around on large tables, in order to try and generate balanced and non-duplicative sequences of group assignments!

1.2: Fundamental Framework: Defining and Avoiding Conflicts. The GRumbler allows users to define *conflicts* between students in a number of dimensions simultaneously, and then searches for group assignments that minimize the total number of within-group conflicts.

The following are examples of relationships between students that could be defined as *conflicts* for this purpose:

- students who work for the same agency (generally applicable for executive program audiences).
- students who come from the same country, state, or city.

¹ The logic, algorithmic approach, search strategies and layout of the GRumbler were designed by the author of these instructions, Malcolm Sparrow. The algorithms are controlled and the overall process managed largely by Excel Macros written in Visual Basic. The TurboGrumbler (for faster search capability) was written in C++ using Microsoft Visual Express 2010.

- students who share the same ethnicity (assuming we prefer heterogeneous groups).
- students assigned to the same dormitory suites, or sharing accommodation (who therefore spend time together anyway).
- students who are known to be particularly talented, or particularly weak (and where any clustering within groups along lines of ability might upset the overall balance, making some groups potentially dominant, or potentially vulnerable).
- “ad-hoc conflicts,” e.g. two or more students who just don’t get along, or for whatever reason need to be kept apart!!

1.3: Important Note: Sensitive Data. Student Data is, of course, confidential in any case. By defining conflicts as described above, administrators may well be creating new data of a highly sensitive nature, incorporating judgments about ability or character. Even the underlying choices about *which variables* to use or not to use as the basis for a group assignment process could be contentious. Defining *ad-hoc conflicts* may be particularly sensitive. All spreadsheets containing such data, including the GRumbler itself, should obviously be treated as highly confidential. Such data will exist within the GRumbler (unless you clear it out), and within any Workbooks that the GRumbler creates (i.e., results files). The detailed instructions below will remind you when and how to safeguard this information.

1.4: Capabilities and Technical Constraints. This version of the GRumbler:

- allows users to define up to 20 different types of conflict, to be used simultaneously.
- allows class sizes up to 5000 students.
- permits division of the class into between 2 and 2500 groups.
- can create either gender-balanced or gender-clustered groups, assuming that data on students’ gender is available.
- allows for different weights to be applied to different types of conflict.²
- permits use of the TurboGrumbler for faster search times, as an optional extra. See section 9 below for details and installation instructions.

1.5: Multiple, Sequential, Group-Assignments. Students can be assigned to one set of groups for exercises (say) on Monday. Then, for Tuesday’s class, new groups can be formed that have minimal overlap with Monday’s groupings. And so on, for Wednesday, Thursday, etc., or for different projects or assignments. If it is mathematically possible, the GRumbler will avoid assigning any two students to the same group today if they have been together before, in any group, on any previous day.³ This version of the GRumbler can generate up to 50 sequential Group Assignments.

² By default, all conflicts will be scored equally (value 1), but users may choose to apply greater weights to particular types of conflict, in which case the algorithm will give these priority attention.

³ When searching for a new grouping, the GRumbler algorithm treats all earlier group assignments as a source of additional *conflicts*. It also continues to recognize and accommodate the set of conflicts originally defined by the user.

1.6: About the Users! This tool is designed for users who

- are at least moderately experienced users of Excel, familiar with all the basic data-entry, copying, pasting, and sorting tools.
- understand the nature of the problem they are trying to solve, and will be able to recognize acceptable solutions when presented with them.
- have the relevant student data in the form of Excel spreadsheets
- want the group assignment lists (i.e., the products of the group assignment process) delivered in the form of Excel Spreadsheets.

To operate the GRumbler, users do NOT need to be familiar with Excel's advanced functions, Macros, Visual Basic programming, the use of mathematical formulas in spreadsheet cells, or with any of Excel's mathematical and statistical analysis capabilities.

2.0: Overall Structure of the GRumbler Workbook:

You need a copy of the GRumbler Workbook: [GRumbler.xlsm](#)

Download a zipped version of the Excel file "GRumbler.xlsm" from the GRumbler webpage. Use the "save" option, rather than "open", and then unzip. Copy the unzipped GRumbler.xlsm file to a directory location where you want to use it, rather than trying to run it from inside the zip directory.

Make sure that the "GRumbler.xlsm" file is downloaded to your own physical hard drive, rather than to a shared network location or the Cloud (as rights controls for networked locations may limit the ability of the GRumbler to create and write files).

An old version (January 2014) [GRumbler.xls](#) is also retained and available on the Grumbler webpage for the sake of anyone using older versions of Excel (either in a PC or MAC environment) who might not be able to use the more modern (GRumbler.xlsm) file.

To familiarize yourself with the GRumbler, I recommend you print these instructions out, rather than trying to follow them on the screen at the same time as navigating the GRumbler Workbook.

Unsafe Macros Warning: When you first open the GRumbler file, depending on your Excel security settings, you may confront pop-up windows warning about viruses and potentially unsafe macros.

- If you see the "Protected View" warning near the top of the screen, click "Enable Editing". You may encounter a "Runtime error 91" notice. If so, click end. Save the file, and re-open it. Better still, before opening the file, make sure that the "Protected View" options in your Excel "Trust Center" settings are turned off, so Excel does not try to open the GRumbler file in "Protected View" mode. You can access these security options with *any* Excel file opened. Use the menu sequence:

File/Options/Trust Center/Trust Center Settings/Protected View. Disable the Protected View options.

- If you see **Security Warning: Macros have been disabled**, click “Enable Content.”
- If you see **Security Risk: Microsoft has blocked macros...** you can unblock the macros as follows. Close the file. In “Explorer,” right click on the name of the closed file. On the menu that appears, click “Properties.” On the “General” tab, at the bottom, there is a “Security” option to “Unblock” the file. Check the “Unblock” box. Click “Apply”. Close the “Properties” pop-up box, and then reopen the GRumbler.xlsm file.

Username: The first time you open the file a dialog box will appear asking you to enter your name, which will thereafter appear in the “Username” box on the Main Menu. The GRumbler does not use your name for any other purpose. You can enter any name you like. This just shows you that this copy of the GRumbler file is yours, rather than somebody else’s. It also provides a gentle nudge for new users to download the latest (and therefore fastest and most bug-free) version of the file from the GRumbler website, rather than copying older versions from their friends and colleagues.

With the GRumbler file open, you should see the three worksheet tabs, as follows:

- 1) **Main Menu:** This serves as the GRumbler’s Control Panel. Most of it is protected, so you cannot mess things up by mistake. There are seven cells (Search Parameters) whose values you are permitted to change:
 - a) “Gender Balancing/Clustering” at cell C6. (Where you can choose whether you want the GRumbler to produce *gender-balanced* groups or *gender-clustered* groups, or ignore gender altogether.)
 - b) “Number of Groups” at cell C7. (Where you will specify how many groups you want to split the class into.)
 - c) Search Loop Structure Parameters, cells C13 to C16, which control the length and strategy for a search. Most users will never need to change these, and they can remain set at their default values, {50,4,8,2} which provide for a fairly extensive search.
 - d) Cell B90 contains the Directory Path for the file “TurboGrumbler.exe.” This will remain blank if the TurboGrumbler is not installed. Cell J14 also reports whether or not the TurboGrumbler is installed and available for use. The default value is “Currently Not Available,” as the majority of users do not need the TurboGrumbler.

Also, at rows 19 through 21, you will see four Buttons used to launch various macros. Once student data has been loaded into the GRumbler, users can simply press various of these buttons and watch the GRumbler do all the work. By scrolling down to rows 84-86 you will find three more “Occasional Use” buttons

(normally kept out of sight to avoid inadvertent use!). In between is a table which acts as a record of the sequence of Group Assignments generated.

- 2) **Search Engine**: This tab provides the basis for the GRumbler's calculations and displays group assignment 'solutions' for inspection by the user. This entire worksheet is protected to avoid any inadvertent changes.
- 3) **Students Detail**: This is where users can enter their own student data. The GRumbler as originally provided has a fictional dataset for a class of 49 students, just for demonstration purposes. Once you've enjoyed a demonstration (see below for instructions) you can clean out the fictional class data by using "clear contents" on cell range B14:K62, or by pressing Erase All Data from Students Detail Worksheet, which is one of the "Occasional Use" Buttons on the Main Menu. Once the fictional class data is erased, the GRumbler is ready for your own class data to be pasted in. The major parts of the *Students Detail* worksheet are as follows:
 - a) Columns B through F, where users will paste the students' names and gender, and two other descriptors (such as *organization* and *country*) that you want to appear on the Group Lists the GRumbler produces.
 - b) Columns G through Z, where users will paste whatever information they want to be used to define *conflicts* between students. Two students will be regarded as having a conflict whenever they have *identical non-blank values* in any one of these 20 conflict columns. (Cells left blank in these columns are simply ignored.)
 - c) Twenty cells [G4:Z4], above the conflict columns, where users can adjust the *weights* they want applied to the conflicts defined by the various columns. By default, these are all set to 1 (so every type of conflict has equal weight), which works fine for the majority of applications.
 - d) Two cells [E13:F13] which act as *headers* or *labels* for the two columns of data to appear on the Group Lists. By default, these cells contain the headers "Organization" and "Country," as these are generally applicable in the executive training context and are probably useful for international conferences too.
 - e) The remainder of this worksheet is protected, to avoid inadvertent changes.
 - f) Two Macro buttons are provided for convenience at the top of the worksheet. One re-adjusts all the column widths to fit the data entered, and the other re-sorts the students by lastname/firstname.

3.0: Enjoy a Demonstration using the Fictional Class: Before you load your own class data, I recommend you first run a quick demonstration involving the fictional class data provided. This will show you how to use the Control Panel (Main Menu) and roughly what to expect from the GRumbler when you start working on your own class.

Demo Step 1: Understanding the problem: For the purposes of the demonstration, let us imagine that this class of 49 students needs to be divided into groups for 4 successive exercises. For Monday and Tuesday's exercises, we want 7 groups. For Wednesday's, Thursday's and Friday's exercises, we want 12 groups. We want all the groups to be gender-balanced.

Click on the "Students Detail" tab. Notice that Columns E and F contain *Organization* and *Country* data for each student, and the Headers in cells E13 and F13 reflect these choices.

Now look to see how the conflicts have been defined in columns G through Z.

- Column J is an exact copy of column E, showing all the students' organizations. This means that the GRumbler will (if possible) separate any students who come from the same organization.
- Column K contains the countries from which students come. The effect of the column K entries here is to ensure, for instance, that the 3 students from the Netherlands will be separated from each other, and similarly for the two students from Nigeria. The 11 Canadians and the Americans (who form the majority of the class) will be spread out as much as possible.
- Column I shows the dormitory suite numbers to which some of the students were assigned. Using these dormitory numbers to define conflicts will keep apart students who share suites. The remainder of the class, let's imagine, are accommodated in a hotel, so their "room numbers" are left blank here.
- Columns G, H, and L through Z are left blank, but could readily be used for any other user-defined conflicts. For example, if William Annett, Theodore Beaumont and Tayo Besley all need to be kept apart (e.g. because they passionately support rival soccer teams) then you could simply enter "soccer conflict" in cells G14, G15 and G16. Entering any non-blank text (such as "A") in these three cells would accomplish the same objective. It doesn't matter *what* the text says; it only matters whether two students have the *same* text in the *same* column.

Demo Step 2: Import Student Data into the Search Engine: Click on the "Main Menu" tab, and look at the Control Panel. By pressing the down arrow you can toggle through the search parameters that users are allowed to adjust. Find the four buttons at row 20, and click on the first one, Load Student Data into Search Engine. The Status Bar (purple) will show "Student Data Imported," and the "Class Size" will now show 49. The Status Bar also reminds the user to adjust the number of groups required before launching a search (see Demo Step 3).

Demo Step 3: Perform Search for first set of group assignments: Move the cursor to cell C7, and set the number of groups (for Monday's exercise) to 7. Take care to press ENTER after you've typed 7. Leave the "Gender Balancing/Clustering" parameter (cell C6) set to "Balancing." Now press the second button, Conduct New Search. The GRumbler will proceed to find a suitable set of 7 groups, and will show you in the (purple) Status Bar when the search is finished. This search should be very quick.

Notice the "Target Score," which the GRumbler calculated to be 50 for this search. This means that, given the number of groups chosen (7), and the set of conflicts defined, there must inevitably be *at least 50 conflicts* in any solution. This does not mean that a perfect solution (with a score equal to the Target Score) necessarily exists, but for sure it is impossible to do any better than this Target score. So if and when the GRumbler finds a solution that equals the Target Score, it stops the search as it could not possibly do better. For this Fictional class, there will be some inevitable conflicts because there are more Canadians (11) and more "USA" students (29) than there are groups (7).

Demo Step 4: Inspect the Solution, and record it if acceptable: Even though the GRumbler has found an optimal solution (equaling the Target Score), you don't have to accept any particular solution. You can inspect it, if you like, by moving to the "Search Engine" Tab, and examining the listing of students. They will be arranged in order of the groups assigned (see column BY), and column BZ indicates which students have remaining within-group conflicts. You can check the reasons for these conflicts by looking back to columns G through Z, which show the conflict-defining data. Notice that all the blank conflict columns (and conflict columns assigned a zero weight) are hidden from view for convenience.

Let us imagine you are NOT satisfied, for some reason, with the group assignment the GRumbler has just found. Go back to the Control Panel (Main Menu Tab). Press Conduct New Search again. The GRumbler will start a new search, which will generate a completely different solution.

Once the GRumbler has found another (optimal) solution, press the third button, Record Best Solution Found. Notice that the "Number of Group Assignments Completed" now shows 1, and the record of searches completed is updated. The Status Bar tells you the solution has been recorded, and reminds you to adjust the "Number of Groups" for the next search, if necessary.

Go back to the "Search Engine" Tab for a moment. Notice that the first set of group assignments has now been recorded in column AA. This column will subsequently act just like another "conflict" column, for all subsequent searches. This is how the GRumbler avoids putting any students together again if they've been together previously.

Demo Step 5: Create the 2nd, 3rd, 4th and 5th Group Assignments: Return to the "Main Menu" tab (Control Panel). Leaving "Number of Groups" at 7, click on the 2nd button, Conduct New Search. Again the target score will be 50, and the GRumbler should find

an optimal solution fairly quickly. Once it does, press the 3rd button, to record the solution.

For the 3rd assignment, we want 12 groups. Move the cursor to cell C7 and select 12 from the drop-down menu; or type 12 and press ENTER. Then click on Button 2, Conduct New Search. The Target score for this search (with more groups) falls to 22, and the GRumbler will most likely reach the Target fairly quickly.

Assuming you're satisfied with the solution, go back to the "Main Menu" and press Record Best Solution Found. "Number of Group Assignments Completed" should now show 3.

For the 4th and 5th searches, leave "Number of Groups" at 12. Launch each search and then record each new set of group assignments generated. The final (5th) search may take a little longer than the others, as the "Density of Conflicts" has now risen (to 33.1%), and the higher the density the harder it becomes to find optimal solutions. The density has risen because all four of the previous group assignments are now being used as conflicts, along with all the original (user-defined) conflicts. As the GRumbler works on this problem, you may have the chance to watch the iterative search counters on the Control Panel, at cells D13: D16, and the Search Progress indicator at G16. The GRumbler will most likely find an optimal solution for this search too. If not, it will terminate after the search strategy specified in cells C13:C16 has been completed.

Press Record Best Solution Found to bring the "Number of Group Assignments Completed" to 5.

Demo Step 6: Generate Group Lists and Output File: Finally, press the 4th button, Prepare Group Lists and Master List. The GRumbler will create a new Excel WorkBook, containing all five sets of group-assignments formatted ready for printing.

[Note for MAC Users who are using MAC Office 2016, or later versions of Excel for the MAC: At this point you may encounter an extra pop-up dialog box, asking you to grant permission to access a particular directory on your machine. Go ahead and grant the permissions, and everything should then work fine.]

The new Workbook name begins "GRumbler Results..." and the file name contains a *date-time* stamp to make it unique. This file is automatically saved to the same directory where your working version of the GRumbler resides. The results file is left OPEN, for your inspection, and a link to it will therefore appear in your TaskBar at the bottom of your screen.

The GRumbler-Results workbook also has two other tabs apart from the group lists. One is a "Master List of Assignments" which may be useful to course administrators. It shows all the successive Group Assignments in one page. If users want to incorporate the group assignments back into their own spreadsheets, then it is easiest to use this list. The students' names appear on this list in the concatenated Fullname format, but have been

sorted alphabetically by Lastname, then by Firstname. If the user's class data is sorted the same way, group assignments can be accurately transferred back to your own spreadsheets. This is easier than trying to do the same from the group lists, which are formatted for printing.

The output workbook also contains an exact copy of the "Students Detail" tab from the GRumbler. This means that if you ever want to re-run the group assignments for this class, maybe adding or dropping a few students, or changing the conflicts that you wish to define, you have a record of the way this sheet appeared last time you ran this class through the GRumbler. This is designed to save you from having to re-do any work you had already done in choosing conflict variables and entering their values.

[To re-run a class, you would simply copy the contents of cells B14:K62 from this copy of "Students Detail" back into the GRumbler's "Students Detail" tab; then make any further changes you wanted (e.g., adding and deleting students), and then press the Load Student Data into Search Engine button once again from the Control Panel.]

Bear in mind that the Students Detail Tab may contain sensitive data, revealing which conflicts you chose to use and any other text, comments, or codes you may have entered. A GRumbler User would not normally want to transmit any Workbook containing this Tab to anyone else.

You have now completed the demonstration, and you should have a clear idea of what the GRumbler does for you, and how easy it is to use.

4.0: Preparing a Master Copy and Working Versions of the GRumbler:

I recommend you now create a new version of the GRumbler worksheet, with the Fictional Class deleted. Otherwise, every time you want to process a new class, you will need to start by wiping the “Students Detail” tab clean. Proceed as follows:

- a) with the original GRumbler workbook open, from the main menu, do FILE/SAVE AS, and then edit the file name to read:

[Master--GRumbler.xlsm](#)

- b) Click on “Main Menu” and scroll down to the “Occasional Use” Buttons. Click on Erase All Data from Students Detail Worksheet.
- c) Return to “Main Menu.” Press Load Student Data into Search Engine. With no student data present, this clears out the Search Engine entirely, minimizing the size of the file. “Class Size” should show zero.
- d) Save and close this file. You now have two versions of the GRumbler: the original, as supplied with the Fictional Class, and your own Master Copy, now empty.
- e) In Windows Explorer, select your Master Copy and, without opening the file, RIGHT CLICK, and COPY. Then, pointing to the same directory, RIGHT CLICK and PASTE. Windows will create a new file:

[Master—Grumbler - copy.xlsm](#)

This is the file you will now use, preserving your Master Copy just in case something corrupts your working version. Move this copy to the directory where you want to work (e.g., where your own class data resides).

5.0: Preparing your own Student Data (in your own Spreadsheet):

Before transferring your class data to the GRumbler, I recommend you perform the following checks within your own database first. These things are easier to do in your own spreadsheets, as many operations (such as adding deleting, or re-arranging columns) are highly restricted once you have moved data to the GRumbler.

- a) Check that you have *Lastname* and *Firstname* for each student. If you have only Fullname (e.g., “Julia Roberts”) you can use that instead.
- b) Check that you have *Gender* for each student. You can mark these Male/Female, or M/F, or M/W. But not Mr. and Mrs. For the GRumbler to understand genders correctly, anything starting with an “M” or an “m” is interpreted as Male, and *any other text* is interpreted as Female.
- c) Inspect carefully any variables within your data that you intend to use as conflict variables. You should check to make sure that text which is *supposed* to be identical really *is* identical. For instance, if you decide to use *Organization* as a conflict variable, the GRumbler won’t know that “FDA” is the same organization as “F.D.A.” or “Food & Drug Administration” or “DHHS: FDA.” I recommend you SORT the student data by reference to each variable (in turn) that you plan to use to define conflicts, and visually check the entries for consistency by scanning down the sorted column.
- d) Of course, you don’t need to use *Organization* or *Country* data if these variables are irrelevant in your context. The Fictional Class demo used *Organization* and *Country* data twice on the “Students Detail” tab: first, in columns E and F (which meant that these data elements appeared on the Group-lists produced in the GRumbler-Results file) and, second, in two of the Conflict Columns (Columns J and K). You don’t necessarily need to use *Organization* or *Country* at all. Nor do you need to use any of the data you plan to place in columns E and F, or any subset of it, as conflict variables. In Columns E and F you might choose to use “Degree Program” instead, or “School,” or “Grade,” or even the students’ telephone numbers or e-mail addresses (if that is what you need to appear on the printed Group lists).
- e) Gather and enter into your own spreadsheet any other information you plan to use to define additional conflicts (e.g., room assignments, age, rank, ability levels, etc.)

6.0: Copying your Student Data to the GRumbler (Student Data Tab):

- a) To use the GRumbler on your own class data, open your file called:

[Master—Grumbler – copy.xlsm](#)

Then FILE/SAVE AS, and substitute [GRumbler—“name of your class”](#) for the filename. Save the file to whatever working directory you use for administration of this particular class.

- b) Make sure there is no student data in the “Students Detail” Tab. If there is any, use the “Occasional Use” Button on the Main Menu that deletes all the student data.
- c) Without closing the GRumbler workbook, open your own student data spreadsheet. Use COPY and PASTE-SPECIAL/VALUES⁴ to transfer the students’ names to the “Students Detail” Tab of the GRumbler: Lastnames (to column B), Firstnames (to column C). Make sure you get these the right way around. If you already have concatenated Fullnames instead, then you can paste these into either column B or column C, but leave the other column blank.
- d) Again using COPY and PASTE-SPECIAL/VALUES, transfer students’ *Genders* to column D. Check that Gender indications obey the rule that Males begin with an “M” or “m”, and Females begin with some other letter (like “W” or “F”).
- e) Again using COPY and PASTE-SPECIAL/VALUES, transfer into Columns E and F the data for the students’ *Organizations* and *Country*, or whatever other data you have decided to substitute for these variables. If you do decide to use variables other than *Organization* and *Country* in columns E and F, you should amend the *headers* in cells E13 and F13. The contents of these two cells will also be reproduced as *headers* on the Group-lists.

⁴ For those not familiar with this function...first Select the whole column of Lastnames in your own data. RIGHT CLICK, COPY. Now select cell B14 on the GRumbler’s Students Detail Tab. RIGHT CLICK, PASTE SPECIAL, VALUES, OK.

7.0: Defining Conflicts, and Adjusting the Conflict Weights:

- a) If you want to use the contents of columns E and F to define conflicts (as in the Fictional Class Example), copy the contents of these columns (E and/or F) to any two of the twenty conflict columns: G through Z.
- b) Again using COPY and PASTE-SPECIAL/VALUES, transfer any other data from your own spreadsheet that you want to use to define additional conflicts. Use any of the columns G through Z for such data. If you do NOT want to use the data in columns E and F as conflict variables, you don't need to. All the conflict columns work exactly the same way, and you can place anything you like into them (either as text or numbers). The contents of these 20 conflict-defining columns will NOT appear on the group lists produced for printing.
- c) If you value one conflict more highly than another (i.e., you are more eager to avoid it), you can adjust the *weights* in cells G4 to Z4. For simple problems, weighting them all equally, at 1.00, works fine most of the time, and means that the Scores produced by the GRumbler represent the *total number of within-group conflicts*.

If you choose to alter the weights, you can enter any integer values or real numbers instead. The search algorithm will recognize and use up to 2 decimal places for conflict weights. In order to give a particular type of conflict precedence over all the others, raising its weight to 2 is normally sufficient.

8.0: Running the Search Algorithm: You've now done your share of the work (preparing the data and defining the conflicts). The GRumbler does the bulk of the work from this point forward. Following this sequence of steps, just as you did when running the Fictional Class Demo, will generate the Group Assignments you need:

- a) Move to the "Main Menu" Tab.
- b) Click the Button Load Student Data into search Engine. Check that "Class size" shows the right number of students. If it does not, check the completeness of the data in "Students Detail." Check *names* in particular. Then press Load Student Data into search Engine again.
- c) Check that "Gender Balancing/Clustering" (cell C6) is set to "Balancing", assuming you want gender-balanced groups. (This is the norm for most applications). If you select "Clustering" the GRumbler will create single-gender groups as far as possible. Select "Off" if you want the GRumbler to ignore gender completely.
- d) Enter the "Number of Groups" desired at cell C7, (if you enter a number manually, remember to press ENTER). Then press Conduct New Search.
- e) When the GRumbler either finds an optimal solution (matching the TARGET SCORE) or finishes its search pattern, you can inspect the best solution found (if you wish) by moving to the "SEARCH ENGINE" tab. Or you can just trust the GRumbler to have done a good job (normally a perfectly sound assumption if it has reached the TARGET score). If you are NOT satisfied with the solution found, press Conduct New Search again, and the GRumbler will search again and find a different solution.
- f) When satisfied, press Record Best Solution Found. The "Number of Group Assignments Completed" (cell D9) will increase by 1.
- g) If you want to produce more than one set of group assignments, repeat steps (c) through (f) for each day, or exercise, for which you need groups. Don't forget to adjust the "Number of Groups" and "Gender Balancing/Clustering" each time, before you launch the new search. [Note: You should NOT reload the data each time, as this wipes everything clean and erases all the previous Group Assignments generated.]
- h) Conducting more or less exhaustive searches: When the "Conflict Density" increases much above 40% optimal solutions which hit the target score become much harder to find, and the GRumbler has to work harder (i.e., longer) to find them. If it does not reach the TARGET SCORE during a search, and you want it to search more extensively for an optimal solution, then you can increase the Search Loop Structure Parameters (cells C13:C16) to define a more thorough search. If you want the GRumbler to try harder (and therefore search longer), try changing the parameter in cell C13 from 50 to 200. If your searches seem to be taking too long, you can reduce the search length, or you may choose to install the TurboGrumbler for faster

searching (it runs roughly 20 times as fast). Section 9 below provides details of how to install the TurboGrumbler. It is extremely rare for the GRumbler to find a solution in a subsequent search that scores better than the best one found in a [50:4:8:2] search. So there is generally no point in running such a search twice.

- i) Once the GRumbler has generated the desired sequence of Group Assignments, press the Button Prepare Group Lists and Master List. A results file called [GRumbler Results—date-time.xlsx](#) is created, saved, and left open. The link to it appears at the bottom of your screen. It is saved automatically to the same directory as the one from which you opened the working GRumbler file (i.e., your course directory).
- j) If, by chance, you are not satisfied with the overall set of group assignments and want to try again, simply begin again at step (b). Reloading the student data resets everything, and erases all previous Group Assignments.

CONFIDENTIAL DATA: Now is the time to remember that both the working version of the GRumbler file (still open) AND the newly created results file contain a copy of the STUDENTS DETAIL worksheet. Once the class is over, you would normally DELETE the working GRumbler file (you would still retain the “Master—GRumbler.xlsm” file, with no student data within it). And be careful to strip the “STUDENTS DETAIL” worksheet out of the GRumbler Results file if you plan to distribute that file electronically.

9.0: Installing and Using the TurboGrumbler:

TurboGrumbler.exe is an Optional Add-On for the GRumbler.xlsm, and provides faster search times. The TurboGrumbler will only be required by users working with very large class sizes (i.e., over 1000 students) and/or wanting to conduct extensive and exhaustive searches for optimal solutions.

To use the TurboGrumbler, you should first download and install the latest version of the GRumbler file. Then follow the instructions below for downloading and installing the TurboGrumbler.exe file.

The explanations in this section are designed for users already familiar with the purpose and operations of the GRumbler.

9.1: Who needs the Turbo Version?

The TurboGrumbler takes a few extra steps to install but offers greater computational power than the GRumbler. If you want the simplest possible installation, and deal with class-sizes smaller than 300, then you should probably just use the GRumbler.

The GRumbler relies solely on Excel macros to perform all its computations. With the TurboGrumbler installed, when users launch a new search, the <GRumbler.xlsm> file will call up a precompiled executable program <TurboGrumbler.exe>, which does all the hard mathematical work, rather than relying on Excel macros for this computationally-intensive task.

The search algorithm for <TurboGrumbler.exe>, coded in C++ and precompiled, runs roughly 20 times as fast as searches performed by <GRumbler.xlsm> itself... which means, unless you are dealing with extremely large classes, searches last only a few seconds.

The GRumbler looks and behaves basically the same way, whether or not the <TurboGrumbler.exe> file is installed. If the GRumbler detects that the TurboGrumbler has been installed, when it comes time to conduct a search the GRumbler automatically passes the TurboGrumbler all the data needed to conduct the search, waits for it to complete the search, and then imports the results and shows them on the "Search Engine" tab in the normal way. The User only deals directly with the worksheet file <GRumbler.xlsm>.

9.2 To try the TurboGRumbler for yourself: First, make sure you have downloaded the latest version of the GRumbler.xlsm file. It might be wise to delete other, older versions of the GRumbler from your machine, to avoid confusion. The TurboGrumbler only works with GRumbler versions dated August 2011 or later.

Download the latest version of the Executable file "TurboGrumbler.exe" from the TurboGrumbler webpage.

Where you then place the GRumbler and TurboGrumbler files, and how they get linked together, is important. Please follow these steps:

(1) Choose some out-of-the-way location, and place both files [GRumbler.xlsm & TurboGrumbler.exe] there: You might create a new directory, called "Downloaded GRumbler Files", under "My Documents," and place both files there. Or put them into some other existing directory that you're unlikely to delete. However, you should avoid using any directory location that requires "Administrator Rights" to access, or is on a network, because under some Operating Systems <TurboGrumbler.exe> won't work from such locations. For instance, in Windows 7, you should avoid placing them under "Program Files" or "ProgramData", or anywhere close to system files. Also, under Windows XP and some other operating systems, if you have a partitioned drive, it seems to be important to choose a location on the C: drive, and not under "Program Files". For example, the following path works fine:

C:\Users\YourName\Downloaded Grumbler files\

(2) Launch the <exe> file, by double-clicking its icon or filename. The "Console Screen" will appear (it should be bright green), and it will tell you that TurboGrumbler has created a textfile, called <Where-am-I.txt> containing the Directory Path. It also explains how to adjust the Console Screen color-balance in order to get the right display appearance when TurboGrumbler runs a search. Follow those instructions, and then press the space bar and the Console Screen will close.

(3) Open the worksheet file <GRumbler.xlsm>. When you first open this worksheet, you may need to adjust Excel security settings to permit the operation of macros, and you will also be asked to enter a username the first time you open the file (see section 2.0 above for details).

(4) The GRumbler should automatically detect the presence of <TurboGrumbler.exe>, as long as both files are in the same directory. The directory path will appear in cell B90 of the GRumbler's "Main Menu" tab, down below the "Occasional Use Buttons." The "Main Menu" will also indicate, at cell J14, that the TurboGrumbler is "Installed and Available." Make sure that the directory path displayed in cell B90 is the same as the directory path shown in the "Where-am-I" text file.

You won't have to open <TurboGrumbler.exe> yourself ever again, unless you decide to move it to a different directory location. The <GRumbler.xlsm> worksheet will be doing the calling, and you'll be dealing only with the <GRumbler.xlsm> file.

The set-up procedure is now complete, and you can check the functioning of the TurboGrumbler by running the normal demonstration using a fictional class, as described in Section 3.0 above.

(5) Erase the fictional class data: Once you're done with the demonstration, erase all data from the "Students Detail" tab. Scroll down to the "Occasional Use" Buttons of the

"Main Menu" tab, and press Erase all data from Students Detail worksheet. Return to the "Main Menu" tab (from the "Students Detail" tab), and press Load Student Data into Search Engine. There isn't any student data, so this effectively cleans out the "Search Engine" tab for you.

(6) Save as Master Copy: Now save and close <GRumbler.xlsm>. This file will now act as your Master Copy of the GRumbler. You can make as many copies of it as you like, and paste them into your working directories (probably with your course files), wherever you want to use the GRumbler. It is important that you leave the file <TurboGrumbler.exe> in its out-of-the-way location, because that's where the <GRumbler.xlsm> file expects to find it. Once the set-up procedure has been followed, and <GRumbler.xlsm> knows where the <TurboGrumbler.exe> is located, the two files no longer need to be in the same directory. <TurboGrumbler.exe> stays where you originally placed it, and you can move or copy <GRumbler.xlsm> to any directory location you like.

9.3 Relocating the <TurboGrumbler.exe> file: If at some point in the future you should choose to move the <TurboGrumbler.exe> file to a new location, you will need to repeat the set-up procedure by double-clicking on its icon or filename *in its new location*. Follow the instructions on the Console Screen to reset the color mix for comfortable viewing (as these settings are directory specific).

Then paste the new directory path from the file <Where-am-I> to cell B90 of <Grumbler.xlsm>'s "Main Menu" tab. Alternatively, if you move <GRumbler.xlsm> into the same directory as <TurboGrumbler.exe> and then open <GRumbler.xlsm>, it will detect the presence of TurboGrumbler and automatically update the directory path information in Cell B90.

Whenever you want to use TurboGrumbler, it is worth checking that the directory address in Cell B90 properly reflects the location of the TurboGrumbler.exe file.

10.0: Frequently Asked Questions:

Question: Why do users have to download a Zip-file version of the GRumbler, rather than downloading the xlsx file directly?

Answer: Because the Harvard Scholar website does not permit the uploading or downloading of the xlsx filetype.

Question: Sometimes the conflict scores in the search engine appear to be in the wrong places. When and why would that happen?

Answer: **[Majority Descriptors]**: Invariably the explanation is that the conflict scores that appear to be misplaced relate to one conflict column that has a single predominant value. For instance, imagine that more than two thirds of the students in a class are from the USA, and the rest are from a variety of other countries. In this situation the label “USA” is deemed a “majority descriptor.” Whenever any of the conflict columns has a “majority descriptor” present, the search algorithm replaces that label with its complement, which means effectively marking all the non-USA students as “non-USA” and disregarding the “USA” labels. Evenly distributing the “non-USA” students across the groups is exactly the same, in effect, as evenly distributing the “USA” students. So using the complement of any majority descriptors produces exactly the same result in the end, but generates lower target scores and make the algorithm run more efficiently. This increases the odds it will find optimal solutions quicker.

Question: Can you terminate a search that is underway, without messing everything up?

Answer: Yes. If you are *not* using the TurboGrumbler, then press and hold the ESCAPE key on your keyboard. Click END on the Visual Basic pop-up window that appears. [You may find the screen displays the “Search Engine” instead of the “Main Menu.” If so, return to the “Main Menu.”] Then adjust the search parameters (e.g. for a shorter search, or for a different number of groups) and launch a new search by pressing the Button Conduct New Search.

If you *are* using the TurboGrumbler, you will see a notice appear on the console display screen as soon as the search time reaches 15 seconds. At any point thereafter you can terminate the search by closing the console window (press x in top right corner), and the best-solution-found-so-far will be transferred back to the GRumbler.

Question: If a student drops out, or needs to be added to the class, is it easy to re-run the group assignment process?

Answer: Yes it is. Re-open the previous working version of the GRumbler for the right class, assuming you have not deleted it. [I recommend you retain this file until all the Group Exercises have actually taken place...because last minute changes (and therefore re-runs of the Group Assignment process) seem all too common.] Alternatively create a new working version of the GRumbler for this class, and re-load the STUDENTS DETAIL data you used before. Then make the necessary changes, as follows:

To REMOVE students: On the STUDENTS DETAIL tab, remove the data for any students who have dropped out by selecting the appropriate ROW (which stretches from column B to column Z), RIGHT CLICK, CLEAR CONTENTS. This will leave a blank row in the middle of the student data. That's perfectly fine.

To ADD students: you can either type new students' data into any vacant rows, or add them as new rows at the bottom of the class data. If copying data from another spreadsheet, remember to use PASTE-SPECIAL/VALUES, rather than PASTE. Don't forget to fill in any conflict columns for added students, so they are treated consistently with the rest of the class. There is no need to sort the data alphabetically OR to remove any blank rows. (The GRumbler takes care of both issues when it imports the Student Data into the Search Engine).

Question: I pressed the Load Student Data into Search Engine Button and the "Class Size" still shows zero, even though the Status Bar says the data has been loaded. I have checked the "Students Detail" tab and the student data is all there. What's wrong?

Answer: Most likely, the Macros for this file are DISABLED due to your security settings. (See note on MACRO SECURITY at section 2.0 above).

Question: What if I don't know the gender of the students?

Answer: Mark them all Male, or all Female, or anything else you wish (like "N/A"). But you must put *something* in the gender column! If the gender data is unreliable or meaningless, turn "Gender Balancing/Clustering" to "Off" before you run any searches.

Question: I pressed the Conduct New Search Button and nothing happens. Why?

Answer: Most likely you manually changed (i.e., without using the drop-down menus) the "Number of Groups," or one of the other Search Parameters on the Control Panel ("Main Menu"), and forgot to press ENTER. Press ENTER and then try again.

Question: Do I actually need to define any conflicts? What happens if I leave all the conflict columns completely blank?

Answer: Then the GRumbler will produce for you groups that are gender-balanced (assuming gender-balancing is turned on), and otherwise random. If you generate a *sequence* of Group Assignments, then the GRumbler will still continue to use all the earlier group assignments as conflicts when forming later group assignments.

Question: I notice in the Search Engine that the weights used for prior group assignments are 1.01, 1.02, 1.03 etc. rather than just 1. Why is that?

Answer: If the GRumbler cannot completely avoid duplicating prior pairings, users would generally prefer that it duplicated older ones rather than duplicating more recent ones. Using a progressive weighting scheme for prior assignments reflects that preference.

Question: If I made a mistake in a prior search (e.g., I forgot to reset the number of groups), can I “go backwards” and redo a previous group assignment?

Answer: You can “rewind” the GRumbler, one set of Group Assignments at a time, by using one of the “Occasional Use” Buttons on the Main Menu: Delete
Most Recent Set of Group Assignments

Question: I’d like to keep a copy of the table that shows the record of “Group Assignments Completed & Recorded So Far.” It doesn’t appear in a GRumbler Results file. How can I get a copy of it?

Answer: With the GRumbler “Main Menu” tab open, use the Excel menus and choose FILE/PRINT. It is set up to print the table of Group Assignments. [If you have chosen to use an older .xls version of the GRumbler, you may need to use the sequence FILE/PRINT PREVIEW/PRINT.]

11.0: Occasionally Asked (Advanced and Technical) Questions:

Question: Is there other software in existence that solves these kinds of problems?

Answer: Yes. The repeated Group Assignment problem most closely resembles a problem known (in mathematics and operations research) as the Graph-Partitioning problem, although virtually all other applications of Graph-Partitioning aim to minimize the *between-group* costs rather than the *within-group* costs. Several specialist software packages exist for solving large and complex Graph-Partitioning problems, and many of these packages deploy algorithms designed to solve “large” problems (with thousands of nodes) efficiently, and to run on parallel processors or supercomputers. Such packages *could* be used for this problem.

However, the GRumbler, by contrast, is designed:

- to work on problems of modest size (with merely hundreds or a few thousand students)
- where processing speed and optimization are less critical issues,
- to save users from having to learn how to use any specialized or sophisticated software, and
- to save users from having to prepare “adjacency matrices” (which are the normal inputs into the mathematical software). The GRumbler generates these matrices itself from the user defined conflicts, and from all the previous group assignments.

Question: Exactly what mathematical techniques are involved in the GRumbler search algorithm?

Answer: The search strategy picks initial group-assignments which are gender-balanced (or gender-clustered, if required) but otherwise random; and then combines three different types of “step” as it tries to improve the solution. It only ever shuffles men with men, and women with women, (unless Gender-Balancing is “Off”) thus preserving gender-balanced or gender-clustered groups at all stages of the search. The three types of step are:

- SCRAMBLING: This involves randomized shuffling either all the men, or all the women, that have within-group conflicts remaining. The SCRAMBLING routine works on men and women alternately.
- SCRAMBLING with RANDOM INCLUSION FACTOR: Likewise, this technique scrambles a group of men, or a group of women, but the group in this case includes a small number (4% of the class) of “non-conflicted” individuals as well as all the conflicted students of that gender. Periodic use of RANDOM INCLUSION FACTORS prevents the search from getting stuck in local minima search loops.
- SWAPS: this technique involves pair-wise swapping of students, between groups, in a way that decreases the total class-wide conflict

score. Local HILL-CLIMBING involves making a series of *available SWAPS*, until there are no more SWAPS available. HILL-CLIMBING is performed after every SCRAMBLER iteration.

Question: I'd like to change the "Conflict Weights" and to define some additional conflicts for some of the later exercises in my class. Can I do this?

Answer: Yes. You can alter the Conflict Data and Conflict Weights before any search, assuming no change in class composition. Go to the "Students Detail" tab and make whatever changes you want in columns G through Z. Be careful not to alter the data in any *other* columns. Then return to the "Main Menu" tab and press Update Conflict Data & Weights (one of the "Occasional Use" buttons). You can check, if you like, that the updated conflict data and weights are now reflected in the "Search Engine" tab. Then launch your new search(es).

Question: Can I enter *negative* numbers as conflict weights, so that the defined conflicts actually act as *attractors*? In some circumstances I'd like similar students grouped together rather than keeping them apart.

Answer: Yes, you can. But don't be surprised if the calculated TARGET SCORE turns out to be a negative number. Also, bear in mind that if you are asking the GRumbler to cluster students along one dimension, but separate the same students along some other dimension at the same time, then these two purposes may turn out to be at odds with one another. In such cases the best scores actually achievable may miss the Target Score by a much higher margin than normal, as the GRumbler makes trade-offs between the two conflicting purposes. The GRumbler may nevertheless find perfectly acceptable solutions, and probably the best ones possible.

Question: I want to assign participants to tables for a dinner, and there are some particular people that I need to group together on "high table," even while everyone else gets thoroughly mixed up. Can I make the GRumbler keep these particular individuals together?

Answer: Yes. Use one of the spare "Conflict Columns" on the "Students Detail" Tab to mark your "top table" members. Enter "top table" (or other suitable label) against their names, and leave the rest of that column blank. Change the "Conflict Weight" for this column (from its default value of 1) to -5. The GRumbler will then treat this column as an *attractor* (see question above), and the -5 is significant enough to make sure this consideration trumps everything else. Be sure to set all your other requirements (such as gender-balancing and other conflicts) as you would normally. [Note that the "top table" group won't necessarily come out as group number 1. But you can renumber the groups later, if you need top table to be table 1, within the GRumbler Results file.]

Question: *Gender Balancing* or *Clustering* seems to take precedence over any of the defined conflicts, regardless of how high I set the Conflict Weights. Can I apply a *weight* to the gender aspects of group assignment, so the GRumbler balances gender considerations with other considerations rather than giving them absolute priority?

Answer: You are right: the goal of Gender Balancing or Clustering is given absolute precedence over any user-defined conflicts. If you want gender considerations to be *weighted* against other conflicts you should *use the gender data as a conflict variable*, as follows:

- (a) turn the Gender Balancing parameter in “Main Menu” to “Off,”
- (b) on the “Students Detail” worksheet, copy the students’ gender data into a spare Conflict Column, and
- (c) choose a suitable Conflict Weight for this (gender) Conflict Column. Use a positive value (e.g., 1) if you want gender *balancing*, and a negative value (e.g., -1) if you want gender *clustering*.

Then the GRumbler will treat the gender data just like any other user-defined conflict.

Question: I used the GRumbler to generate four sets of group assignments, but after the first two exercises, we lost a couple of students and I want to re-run the 3rd and 4th sets of group assignments, making some changes to the class composition, *and* still avoid duplication with the 1st and 2nd sets of groupings. Can I load the first two sets of group assignments into the GRumbler, edit the class membership, and then re-run the searches for the 3rd and 4th sets of assignments?

Answer: If you merely want to “rewind” the GRumbler and re-run some searches, but without changing the class composition, you can use Delete Most Recent Set of Group Assignments, which is one of the “Occasional Use” Buttons on the Main Menu.

If you need to change the class composition as well as “rewind”, you should conduct a new search, treating any set of previously used group assignments as *Defined Conflicts* for the new search. A suitable procedure is described in the box below:

- (1) Open a new copy of the GRumbler, and also open your original GRumbler Results file (showing all the previously generated group assignments).
- (2) Copy the “Students Detail” information from your Results file into the “Students Detail” tab of the GRumbler.
- (3) Now go to the “Master List of Assignments” tab in the Results file, SELECT and COPY the data (without the header) from the columns for Group Assignments labeled 1st and 2nd. PASTE this group assignment data into any two of the twenty CONFLICT columns in the “Students Detail” Tab of the GRUMBLER. [Pick Conflict columns that are otherwise empty, or “spare”. Make sure this data lines up correctly with the student names. It *should*, because students are listed in alphabetical order by LASTNAME, both on the “Master List of Assignments” and on the “Student Detail tab.”] Make sure the conflict weights for these conflict columns are set to the default value of 1.
- (4) Now make any changes to the class composition by adding new rows or clearing the contents of existing rows within the GRumbler’s “STUDENTS DETAIL” tab. If you *add* any students, make sure you also add any relevant conflict-defining data so that they are treated consistently with the other students. For added students, leave the conflict columns showing previous group assignments *blank*.
- (5) Move to the “Main Menu” Tab, and proceed as usual. First, press Load Student Data into Search Engine and then launch your new searches. The GRumbler will treat the (already used) group assignment data just like it treats any other conflict-defining data. [Note that the “Number of Group Assignments Completed”, however, will reflect only the *new* group assignments that you generate from this point forward; and any new GRumbler Results File you create will similarly contain Group Lists only for *newly generated* group assignments.]

Question: Could I use the GRumbler as a class scheduling package?

Answer: It is not designed for that, but in fact, yes, you could—because class scheduling is also fundamentally a graph-partitioning problem. In place of students’ names you’d use the names (or numbers) of courses to be scheduled. Then you’d specify conflicts between courses that should not be scheduled at the same time (because, for example, they required the same classroom, faculty, or equipment; or because there were significant numbers of students who wanted or needed to take both courses). You’d then set the “Gender Balancing/Clustering” parameter to “Off” and set the “Number of Groups” equal to the number of separate time slots to which any course could be assigned. With the problem set up that way, the groups subsequently generated by the GRumbler would represent clusters of courses that could be scheduled (with minimal inconvenience) during the same time slots. After generating these clusters of courses, you’d then need to allocate each cluster to one actual time period. Commercial scheduling packages normally offer a lot more choices and sophistication than this. Also note that the course-scheduling application doesn’t need *sequential* group assignments, as constructing a class schedule is generally a one-time thing.

Please Contact me: malcolm_sparrow@harvard.edu

- I would like to hear your experiences with the GRumbler, good or bad.
- I'd value your suggestions as to what enhancements you'd like to see in future versions.
- I'd value your feedback about these instructions, particularly any parts that need further clarification.
- I'm also happy to try to answer any other queries you may have.

Please do NOT send me any spreadsheets containing student or other personal data.

Happy GRumbling...