Demand and capacity models

High complexity model user guidance

August 2018

Published by NHS Improvement and NHS England

Contents

1. What is the demand and capacity high complexity model?

The demand and capacity high complexity model (HCM) is an Excel tool that will assess the demand and capacity requirements for services with complex pathways.

If a service has patients moving between different intensity levels (‘patient groups’), where they can either step up to a more intense level of treatment or step down to a lower level of treatment, the HCM takes this into account when predicting the future demand on the service.

The HCM is suitable for any service that provides treatment and care for patients with a long-term condition, where the patient will receive different modes of treatment in different frequencies based on their particular circumstances and where several groups of patients are seen/treated at any time. Specifically, the HCM is applicable for services such as:

- mental health services
- community services
- chronic conditions.

Detailed guidance on how to group patients will be released in summer 2018.

2. Methodology

The simple model of demand and capacity used by the core model is not applicable to services that experience internal patient transfers.

A more complex methodology has to be used – the 'stock and flow' model.

Consider the bathtub shown below – water flows in at a certain rate (the inflow), and exits at another rate (the outflow):

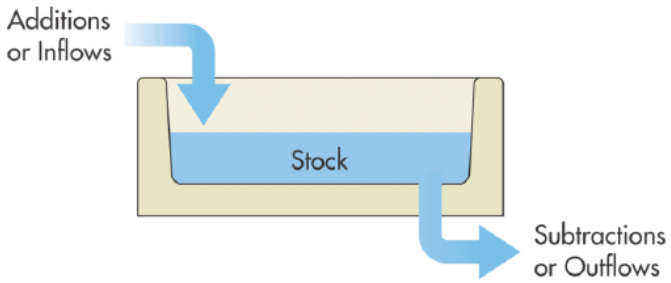

- when the inflow exceeds the outflow, the stock rises
- when the outflow exceeds the inflow, the stock falls
- the peaks and troughs of the stock occur when the net flow crosses zero
- the stock should not show any discontinuous jumps (instant change from full to empty) – it is a continuous process.

We can apply this analogy to any healthcare service – we can consider our 'stocks' to be the patient caseloads, and the inflows and outflows are patients moving in and out of these caseloads. A practical example involving dialysis is shown below.

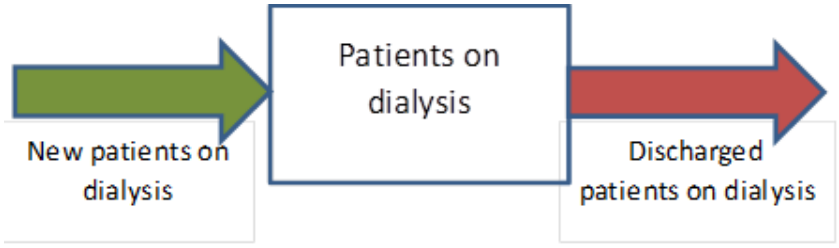

This flow diagram shows new dialysis patients being referred into the service, which adds to the patient caseload. Patients may stay on the caseload for as long as clinically appropriate, and are then discharged.

We will now add an extra level of treatment – patients who undergo a kidney transplant:

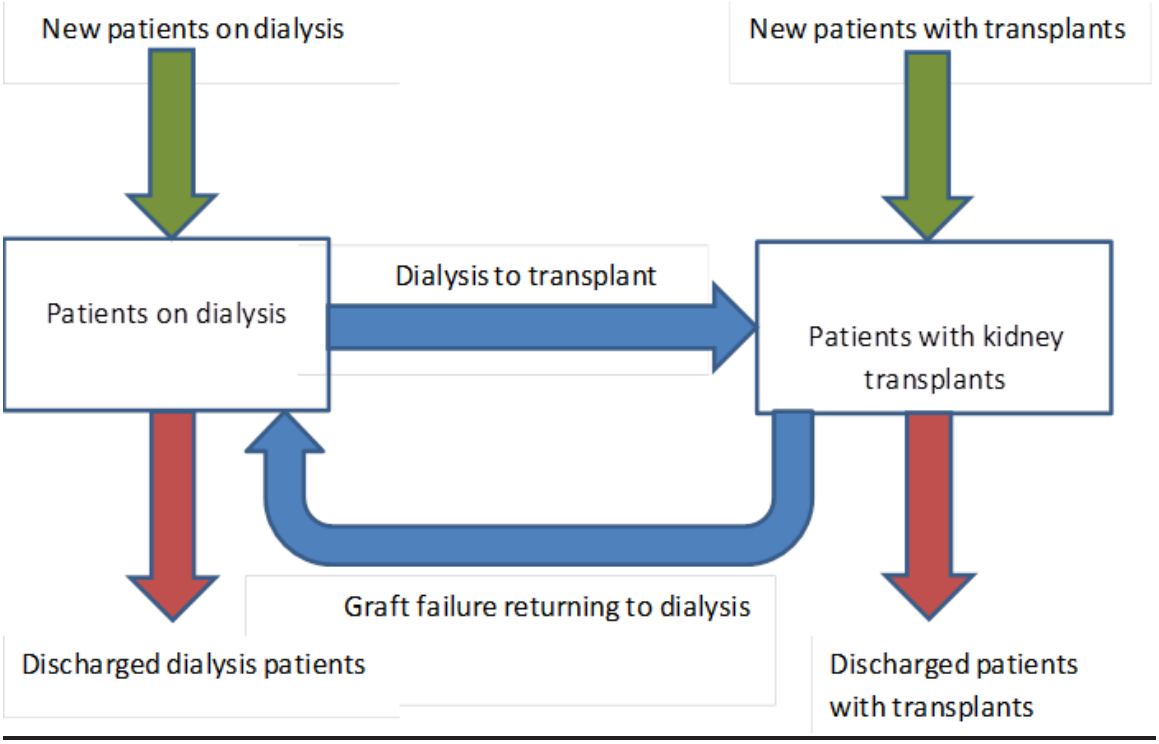

There are now several different possible pathways for patients:

- patients may be referred into the service as either dialysis or transplant patients
- patients on dialysis may be discharged, or become transplant patients
- transplant patients may return to dialysis, or be discharged from the service.

To be able to predict the expected number of patients on dialysis and patients with kidney transplants in the next 52 weeks, we need to understand the flow rates for all the arrows in the previous graph. For that aim, first we will use:

- the historic number of patients on the caseloads (stocks)
- the number of patients transferred between the two caseloads (blue arrows)
- the number of discharged patients from either caseload (red arrows).

We will estimate the rates using the historic data. For instance, if the number of ‘patients on dialysis’ at the end of week one is 100, and during the next week, three patients are transferred to ‘patients with kidney transplants’, the dialysis to transplant rate for that week is 3/100 = 3% per week.

This calculation will be iterated for all the blue and red arrows on a weekly basis to calculate seasonal and yearly rates as well.

External additions to the caseload (green arrows) do not depend on caseload flows – they can be considered external variables, and can be analysed independently using a statistical process control (SPC) chart.

3. How to use this guide

This guide is a quick reference manual that highlights the functionalities of the sheets in the HCM. We assume that you already have some familiarity with basic concepts of demand and capacity management. This includes as a minimum:

- understanding the role of variation, and how to account for it when managing capacity (eg standard deviation, percentiles, SPC charts)
- appreciation of the different levels of intensity/groups of patients
- understanding the methodology used in this model
- differentiation of internal transfers and external additions.

Guidance is also available within the model, and can be accessed in one of two ways:

- If you see the following icon  , you can click on it to bring up an information window. Click on the icon again to hide it.
- If you see a cell with a small red triangle in the upper right corner, hovering your mouse cursor over the cell will make an information window appear.

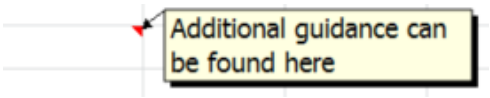

4. Required data items

To populate the HCM, you will first need to group patients into levels of intensity. Patients should be grouped by the level of treatment they require. The HCM can accept a maximum of five levels of intensity on which to predict future demand.

We are also working on a full guide on methods to group patients.

The following data items are required for each level of intensity:

- patients on caseload
- external additions to the caseload
- discharges from the caseload
- internal removals from the caseload to other levels of intensity (transfers from the caseload)
- unit of contact with the service (contacts, minutes, etc)
- frequency of contact with the service
- capacity information for the service
- non-attendance slot lost (NASL) rate (aggregated for the whole service)
- non-attendance slot lost (NASL) rebooking rate (aggregated for the whole service)
- patients past their due date.

We recommend that at least 104 weeks of data (two years) is used for the demand forecasting; however, the model will accept 52 (one year) to 156 (three years) data points.

5. Workbook structure

The HCM, and this guide, are split into the following sections:

- setup
- demand 1 – external transfers
- demand 2 – transfers between levels of intensity
- data validation
- service
- SPC chart – caseload
- SPC chart – external additions
- capacity
- critical resources
- capacity setup
- capacity summary
- parameters
- results.

6. Setup

The ‘Setup’ sheet lets you set up the model to reflect some key characteristics of your service.

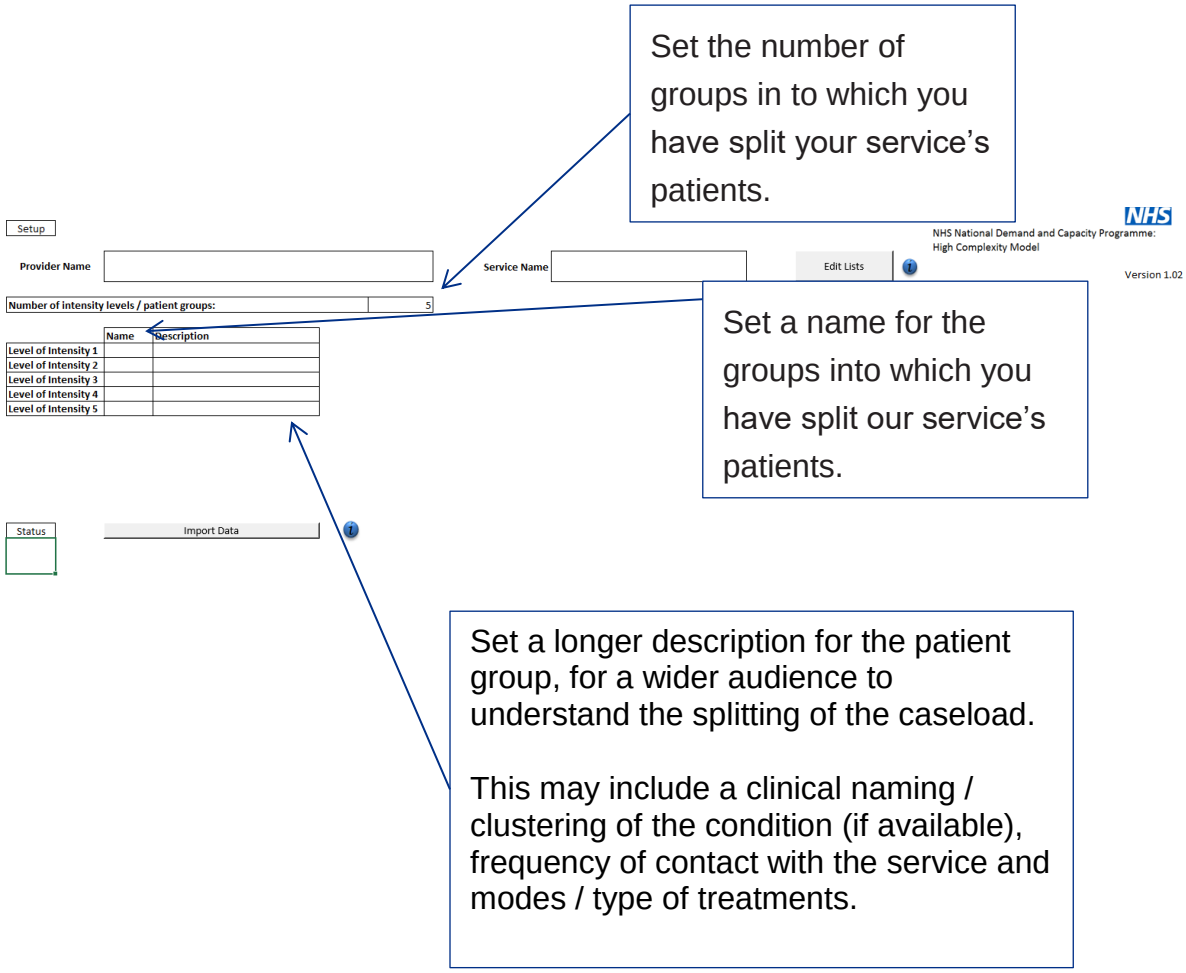

7. Demand – external transfers

The ‘Demand – external transfers’ sheet is used to enter the patient caseload, external additions to the service and discharges from the service, per week per level of intensity.

If you are copying and pasting data from another sheet, please ensure you use the ‘Paste Special (Values)’ or ‘Paste Special… (Text)’ method of pasting data. If you are not sure how to do this, please get in touch with your informatics department for help.

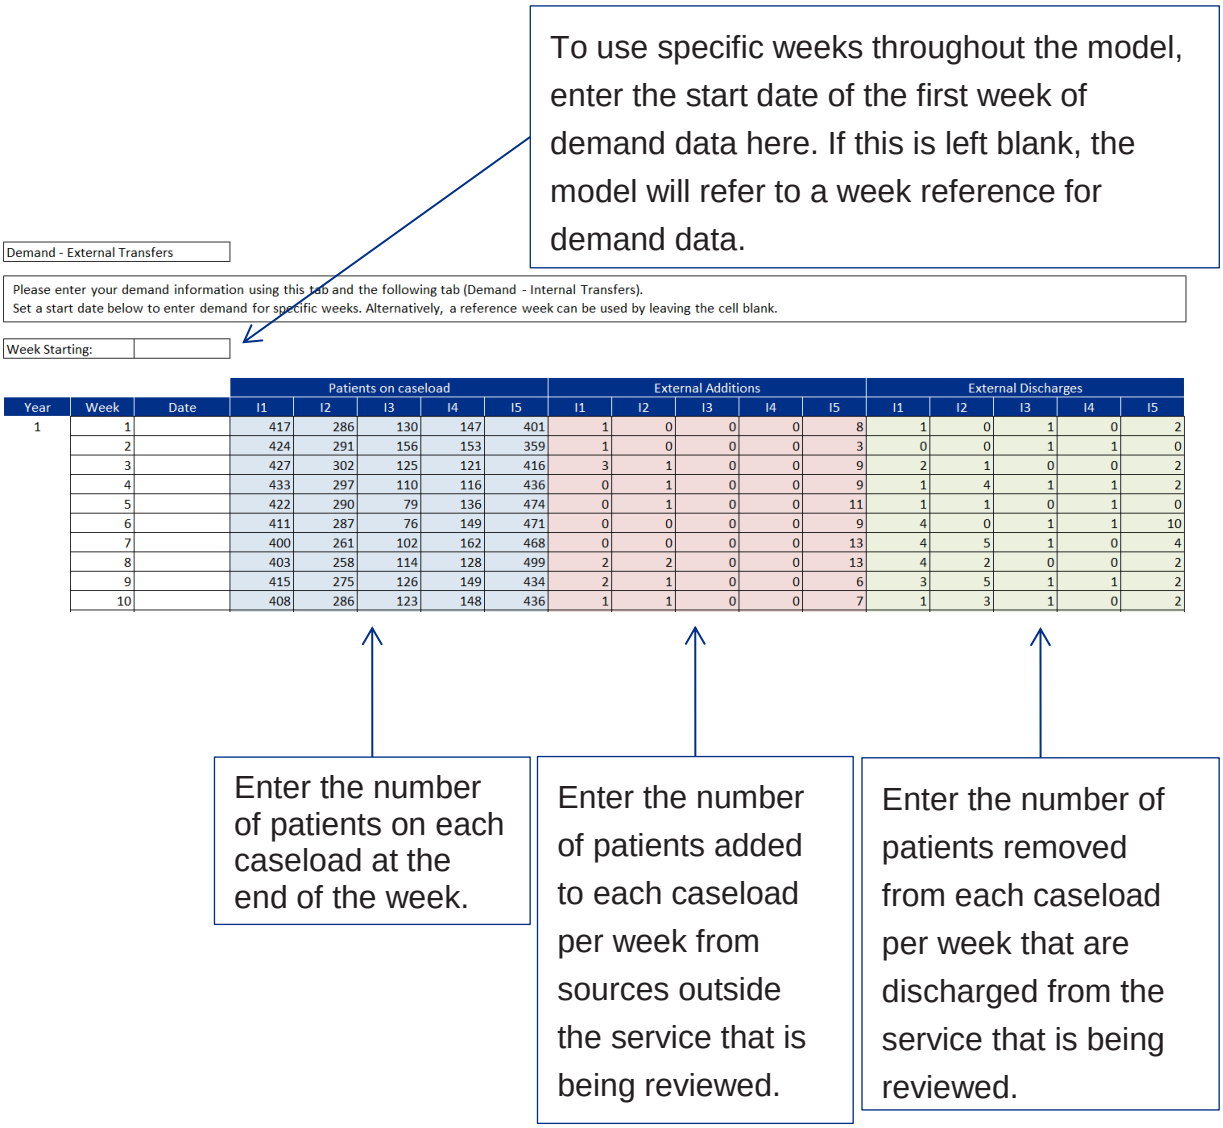

			Patients on caseload					External Additions					External Discharges				
Year	Week	Date	I1	I2	I3	I4	I5	I1	I2	I3	I4	I5	I1	I2	I3	I4	I5
1	1		417	286	130	147	401	1	0	0	0	8	1	0	1	0	2
	2		424	291	156	153	359	1	0	0	0	3	0	0	1	1	0
	3		427	302	125	121	416	3	1	0	0	9	2	1	0	0	2
	4		433	297	110	116	436	0	1	0	0	9	1	4	1	1	2
	5		422	290	79	136	474	0	1	0	0	11	1	1	0	1	0
	6		411	287	76	149	471	0	0	0	0	9	4	0	1	1	10
	7		400	261	102	162	468	0	0	0	0	13	4	5	1	0	4
	8		403	258	114	128	499	2	2	0	0	13	4	2	0	0	2
	9		415	275	126	149	434	2	1	0	0	6	3	5	1	1	2
	10		408	286	123	148	436	1	1	0	0	7	1	3	1	0	2

8. Demand: transfers between levels of intensity

The ‘Demand: transfers between levels of intensity’ sheet is used to track the movement of patients between the different levels of intensity within the service. The patients are tracked from their originating level of intensity to the level of intensity which they are joining. For instance, a patient whose condition has deteriorated may be transferred from a moderate level of intensity to a high level of intensity.

Hovering over the receiving level of intensity will highlight the direction of transfer of patients on the transfer diagram (see below).

Demand - Internal Transfers

Use the matrix below to track patients being transferred between different treatment Intensity Levels.

Transfers between levels of intensity

			I1 to				I2 to				I3 to				I4 to				I5 to			
Year	Week	Date	I2	I3	I4	I5	I1	I3	I4	I5	I1	I2	I4	I5	I1	I2	I3	I5	I1	I2	I3	I4
1	1		0	0	0	10	1	0	0	11	2	25	15	13	0	0	86	24	9	0	0	101
	2		0	0	0	6	2	0	0	18	1	25	13	16	0	0	82	33	9	0	0	109
	3		0	0	0	7	2	0	0	27	1	40	18	37	0	0	65	54	6	0	0	69
	4		0	0	0	7	1	0	0	25	3	24	19	21	0	0	53	23	10	0	0	53
	5		0	0	0	16	4	0	0	27	2	24	18	28	0	0	41	26	0	0	0	70
	6		0	0	0	11	2	0	0	13	2	12	19	20	0	0	51	31	0	0	0	77
	7		0	0	0	13	1	0	0	32	0	12	10	11	0	0	60	34	5	0	0	97
	8		0	0	0	12	4	0	0	20	3	21	12	21	0	0	69	51	10	0	0	74
	9		0	0	0	10	2	0	0	7	6	30	8	5	0	0	62	29	15	0	0	105
	10		0	0	0	15	0	0	0	24	2	37	17	20	0	0	74	49	6	0	0	105

Enter the number of patients moving from one level of intensity to an alternative level of intensity within the same service.

The transfer diagram will highlight the direction of the transfer of patient from the originating level of intensity to the level of intensity they are joining.

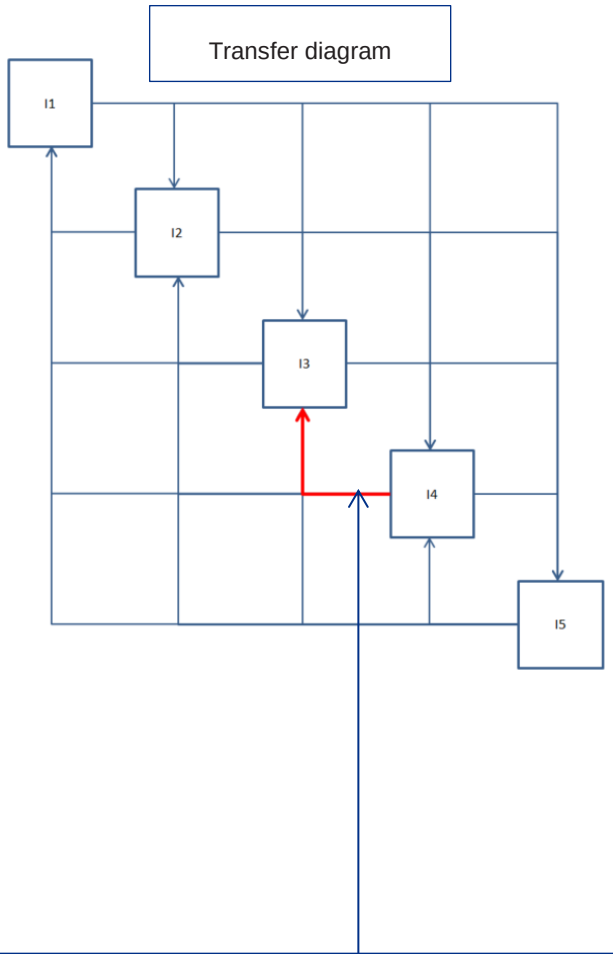

A red arrow will highlight the direction of the transfer when hovering over the receiving level of intensity on the transfer between levels of intensity's table.

9. Data validation

The ‘Data validation’ sheet measures the consistency of the provided dataset. The model estimates the expected number of patients per level of intensity based on the information provided in the first week of data and the subsequent transfers, additions and discharges of patients. The outcome is called validated data and it is compared against the entered number of patients per level of intensity on a weekly basis. The comparison is carried out using the mean absolute percentage error (MAPE) for the pool sizes entered in the ‘Demand’ sections of the model against the validated data. If the MAPE for either level of intensity is larger than 5%, you will need to decide what to do with your data. Your options are to:

- improve the consistency of the data and repopulate the model later with an improved dataset
- use the validated data (use this if you are more confident about the quality of the data for the transfers between levels of intensity rather than the patients per level of intensity). Please note that the model will estimate all the flow rates based on the validated data rather than the patients on the caseloads that you entered
- use the data you have already entered in the ‘Demand’ sections. Please note that due to the lack of consistency in the datasets, the flow rates calculated within the model may not represent your service. Hence, any prediction will not be as realistic as desired.

If your data is consistent, ie MAPE for the different levels of intensity is lower than 5%, you can use either validated data or the data you have entered and the model results will be similar.

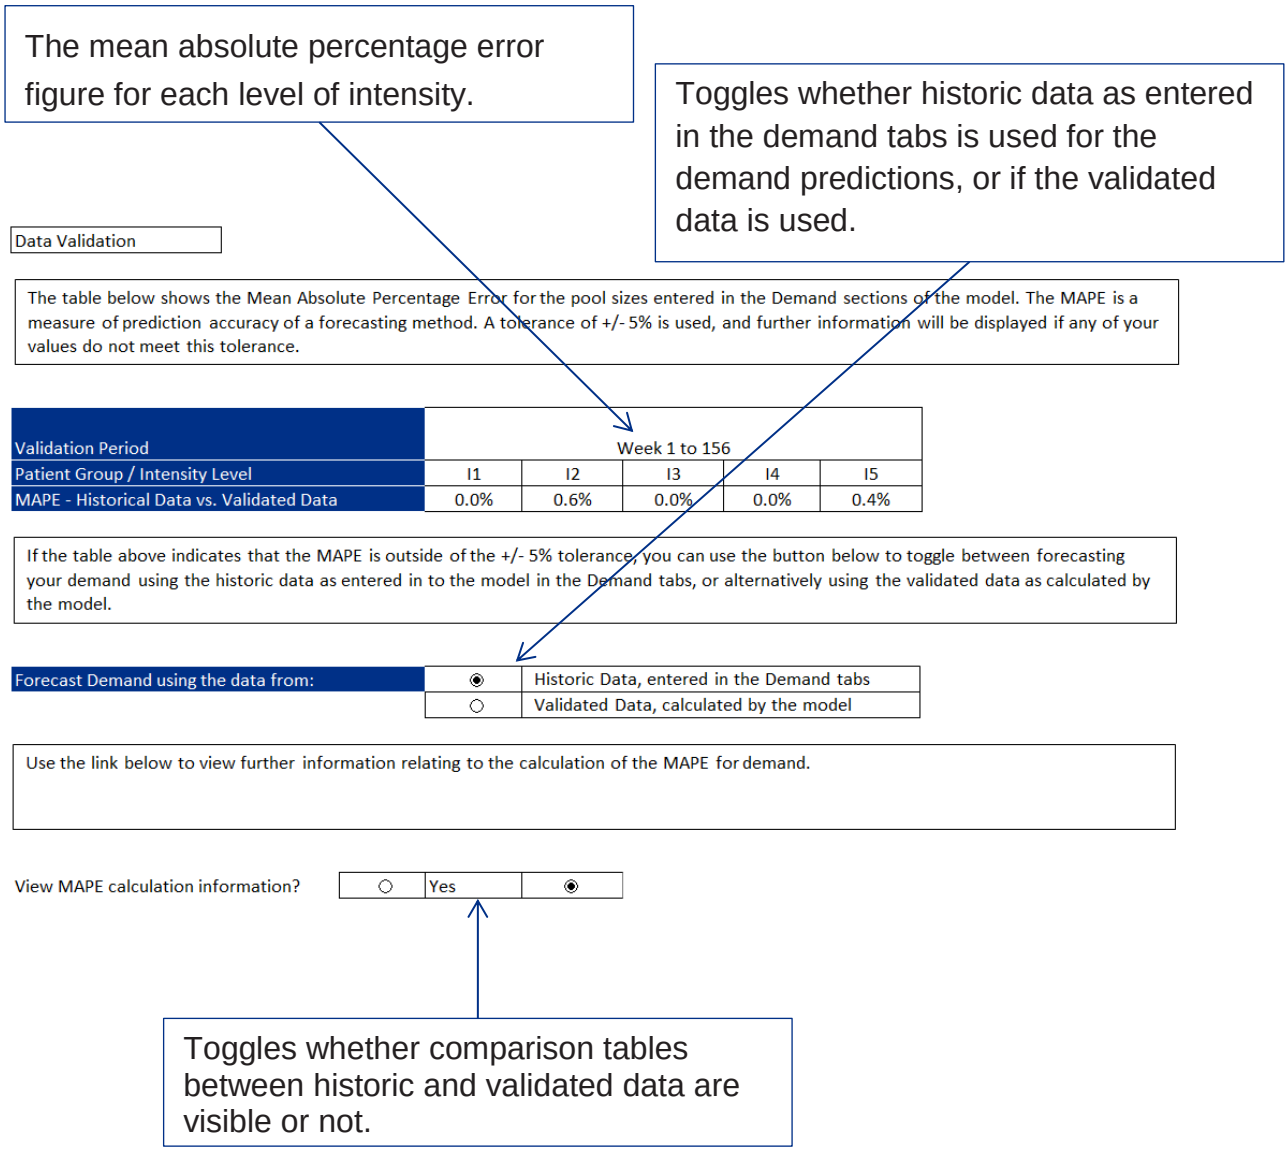
The mean absolute percentage error figure for each level of intensity.
Toggles whether historic data as entered in the demand tabs is used for the demand predictions, or if the validated data is used.
Data Validation
The table below shows the Mean Absolute Percentage Error for the pool sizes entered in the Demand sections of the model. The MAPE is a measure of prediction accuracy of a forecasting method. A tolerance of +/- 5% is used, and further information will be displayed if any of your values do not meet this tolerance.
Validation Period
Week 1 to 156
Patient Group / Intensity Level
I1
I2
I3
I4
I5
MAPE - Historical Data vs. Validated Data
0.0%
0.6%
0.0%
0.0%
0.4%
If the table above indicates that the MAPE is outside of the +/- 5% tolerance, you can use the button below to toggle between forecasting your demand using the historic data as entered in to the model in the Demand tabs, or alternatively using the validated data as calculated by the model.
Forecast Demand using the data from:
Historic Data, entered in the Demand tabs
Validated Data, calculated by the model
Use the link below to view further information relating to the calculation of the MAPE for demand.
View MAPE calculation information?
Yes
Toggles whether comparison tables between historic and validated data are visible or not.

10. Service

The ‘Service’ tab is used to create a homogenous unit of capacity to help predict the service’s required capacity to meet the predicted future demand. The configuration of this sheet depends on the number of levels of intensity that the service has configured, and the units used per type of contact.

Set the number of different modes of treatment used for patients in each level of intensity.

Set the units per type of contact from ‘contacts’, ‘minutes’ or ‘other’.

Service

As the service / modes of treatment and frequency of contact differs among levels of intensity, it is necessary to calculate the demand in a homogenous and comparable unit.

How many services / modes of treatment are used for I1? :	3
How many services / modes of treatment are used for I2? :	3
How many services / modes of treatment are used for I3? :	3
How many services / modes of treatment are used for I4? :	3
How many services / modes of treatment are used for I5? :	3
What are the "units" used per type of contact?	Minutes

Intensity	Mode of Service / Treatment	Group Session?	Length of appointment / session	Number of patients seen / treated	Frequency of Contact	Minutes per week required per patient per type of treatment	Minutes required per patient per level of intensity - Non-Group Sessions	Minutes required per level of intensity - Group Sessions
I1	Group Session	yes	60	5	0.02	0.24	0.9	60
	Face to Face	No	15	1	0.02	0.3		
	Session	No	30	1	0.02	0.6		
I2	Group Session	Yes	120	5	0.2	4.8	9	120
	Face to Face	No	15	1	0.2	3		
	Session	No	30	1	0.2	6		
I3	Group Session	Yes	180	5	0.5	18	22.5	180
	Face to Face	No	15	1	0.5	7.5		
	Session	No	30	1	0.5	15		
I4	Group Session	Yes	240	5	0.8	38.4	36	240
	Face to Face	No	15	1	0.8	12		
	Session	No	30	1	0.8	24		
I5	Group Session	Yes	300	5	1	60	90	300
	Face to Face	No	60	1	1	60		
	Session	No	30	1	1	30		

The service required per level of intensity is calculated for non-group and group sessions

The data points required to complete the Service tab are detailed in the table below.

Mode of service/treatment	Enter a description of the treatment type.
Group session?	Define if the session is a group session or not. If it is a group session, the capacity of the session is shared amongst the number of patients seen/treated.
Length of appointment / session	This option is only available if the units per type of contact is set to ‘Minutes’ or ‘Other’. If the units per type of contact is ‘Other’ and length of appointment/session is not applicable, set it to 1.
Number of patients seen/ treated	If group session is selected, enter the number of patients seen per group session.
Frequency of contact	Select from the dropdown list the frequency of contact per mode of service, or enter a value. To work out the frequency, divide the number of contacts per year by 52, eg if the frequency is once every three months, there will be four contacts in a year. The entered value would be 4/52, or 0.077.

11. SPC chart: caseload

The ‘SPC chart: caseload’ sheet has two functions:

- checks the historic demand data for anomalies that may affect our assumptions around future demand
- sets seasonality periods, which can be used as individual baselines when predicting future demand.

Please note that the seasonal periods configured in the ‘SPC chart: caseload’ sheet will affect the seasonal periods throughout the model.

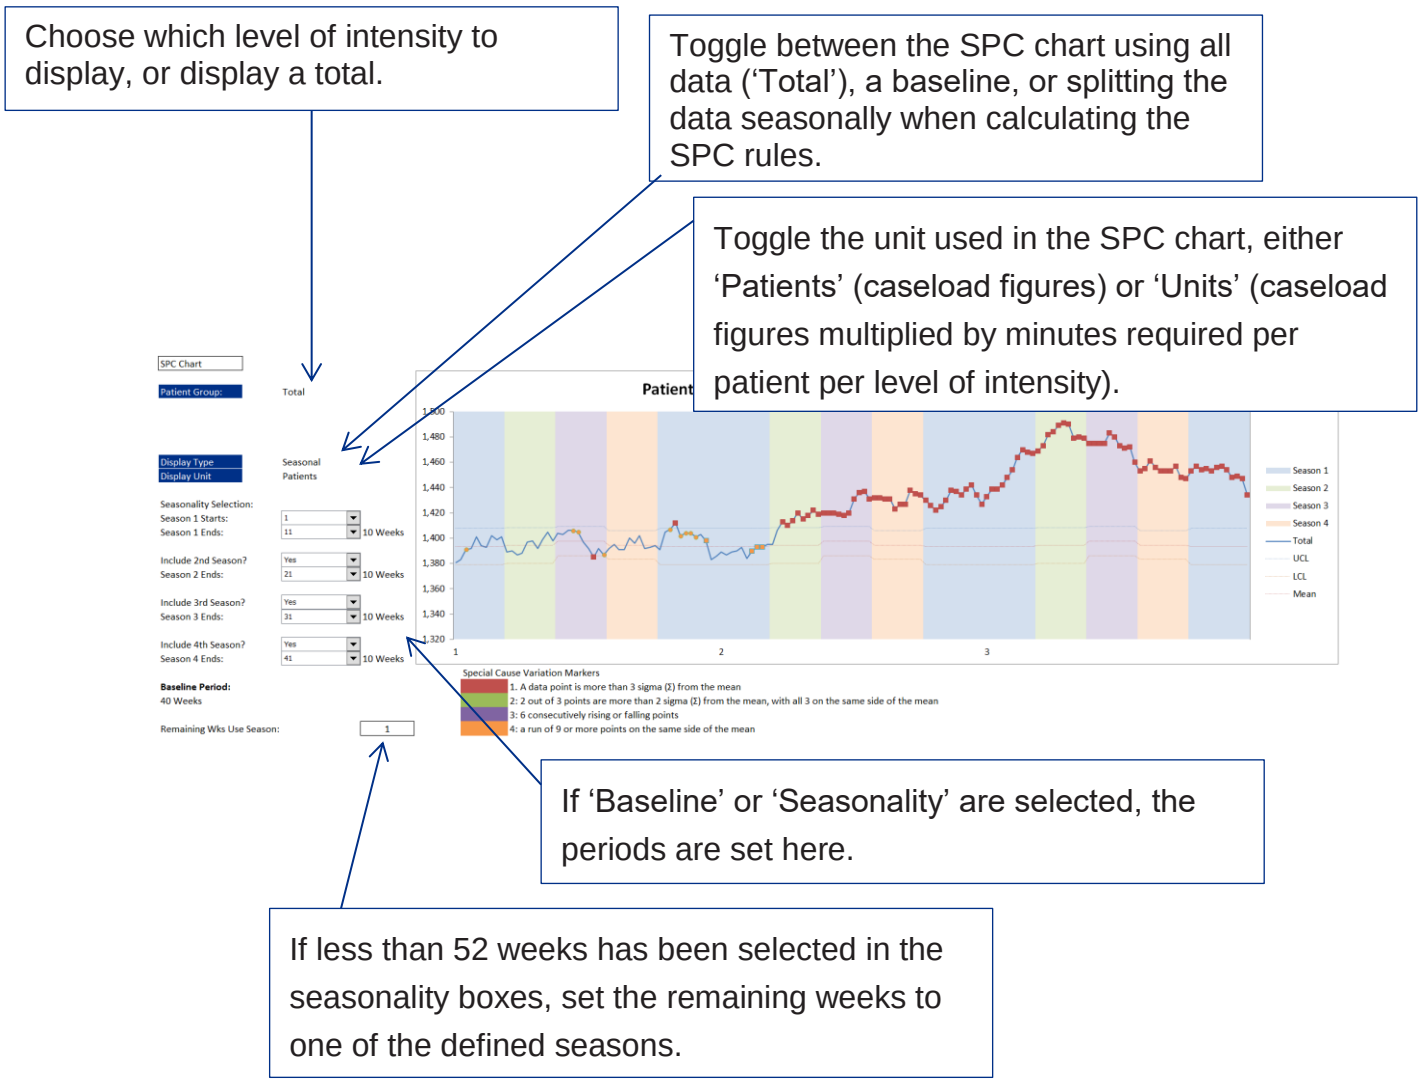

12. SPC chart – external additions

The ‘SPC chart – external additions’ displays the historic demand data for external additions to the service, ie new patients added to any level of intensity that were not in the system already. We recommend that you analyse it independently to the caseloads as there is no direct correlation between the current number of patients in an intensity level and the external additions. Based on your analysis, you will be able to define how you expect it to behave in the ‘Parameters’ sheet.

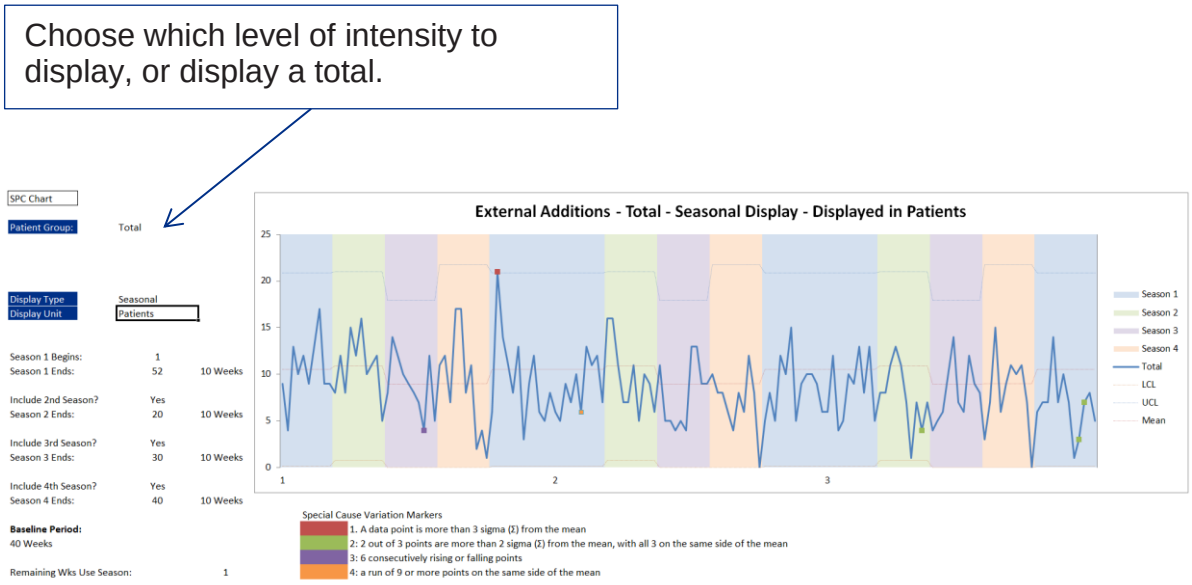

The remaining options are dictated by the seasonal selections made in the caseload SPC chart sheet, and are not editable in this sheet.

13. Capacity

The ‘Capacity’ sheet is used to set up the blocks of capacity that the service uses on a weekly basis. It allows you to enter capacity for ‘Core’ capacity and ‘Ad hoc’, or plus cost, capacity. Although only the core capacity table is shown below, the same data points and configuration apply to the ad hoc set up.

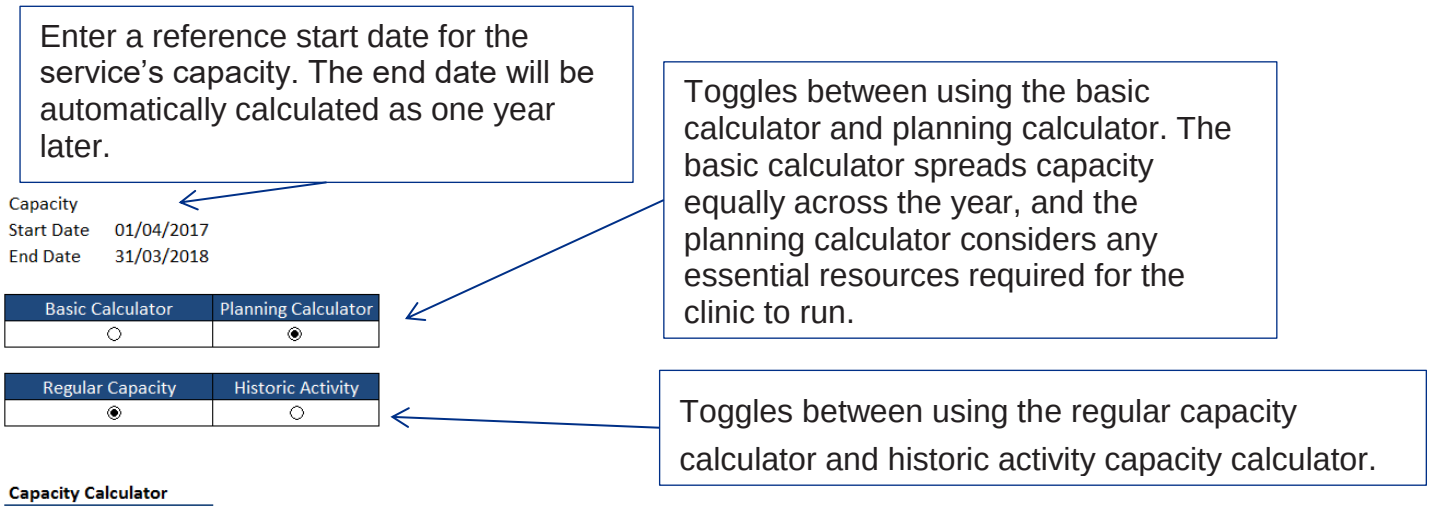

Capacity Calculator

Core Capacity

										Totals	65800	1265.385	1265.384615
Clinic Name	Clinic Description 1	Clinic Description 2	Suitable Intensity Levels					Group Session?	Weeks Per Year	Minutes Per Week	Total (Minutes)	Average per Week	Average per Intensity Level
			I1	I2	I3	I4	I5						
Test	Test	Test	X					Yes	52	700	36400	700	700
Test2	test	test		X				No	42	700	29400	565.3846	565.3846154

The data points to complete the ‘Capacity’ sheet are detailed in the table below.

Clinic name	Set a name for the unit of capacity.
Clinic description 1 / 2	Set a description for the unit of capacity.
Suitable intensity levels	Select the levels of intensity that the capacity can be used for. Multiple selections can be made here, and the capacity will be evenly spread across the different intensity levels.
Group session	Specify if the clinic is a group session.
Weeks per year	Set the number of weeks per year the clinic should run.
[Minutes/contacts/other] per week	Set the capacity of the clinic in the unit defined in the Service tab.

Total ([Minutes/contacts/other]) is the capacity multiplied by the number of weeks per year that the clinic should run. The average per week then spreads this total over 52 weeks. If the clinic can service multiple intensity levels, then the ‘Average per Intensity Level’ field will split the average per week evenly over the levels of intensity.

14. Critical resources

The ‘Critical resources’ tab displays if ‘Planning Calculator’ is selected in the ‘Capacity’ sheet. This sheet allows you to define resources which dictate whether or not a clinic is able to run.

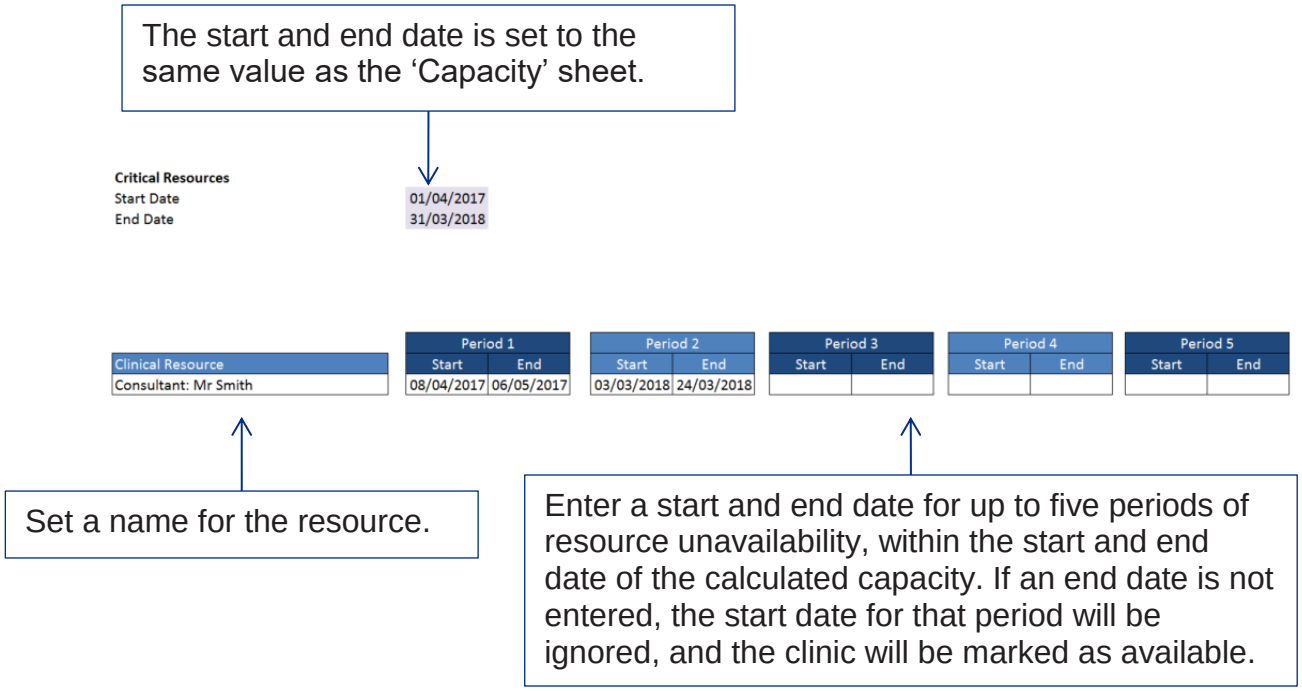

15. Capacity setup

The ‘Capacity setup’ sheet links the resources defined in the ‘Critical resources’ sheet to the clinics configured in the ‘Capacity’ sheet. Although only the ‘Core’ clinics are shown below, data entry for the ad hoc clinics mirrors this.

Moves between configuring resourcing allocation for core capacity and for ad hoc capacity.

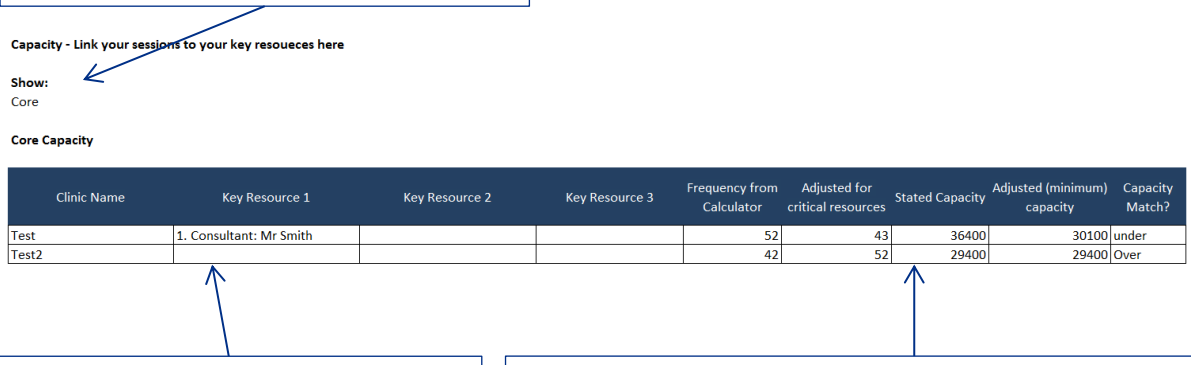

Capacity - Link your sessions to your key resoueces here

Show:
Core

Core Capacity

Clinic Name	Key Resource 1	Key Resource 2	Key Resource 3	Frequency from Calculator	Adjusted for critical resources	Stated Capacity	Adjusted (minimum) capacity	Capacity Match?
Test	1. Consultant: Mr Smith			52	43	36400	30100	under
Test2				42	52	29400	29400	Over

Configure up to three key resources for each clinic. If any of these resources have been marked as unavailable in the ‘Resource configuration’ sheet, the clinic will be marked as unavailable.

The ‘Capacity setup’ will determine the frequency for the clinic as set up in the ‘Capacity’ sheet, the frequency available with the essential resources now configured, the stated capacity from the ‘Capacity’ sheet and the now adjusted total capacity.

16. Capacity summary

The ‘Capacity summary’ sheet displays the entered capacity configuration across 52 weeks, using the session data entered in the ‘Capacity’ sheet, and any extra options selected by the user.

If you are using the basic calculator, the model assumes an even spread of capacity across the year.

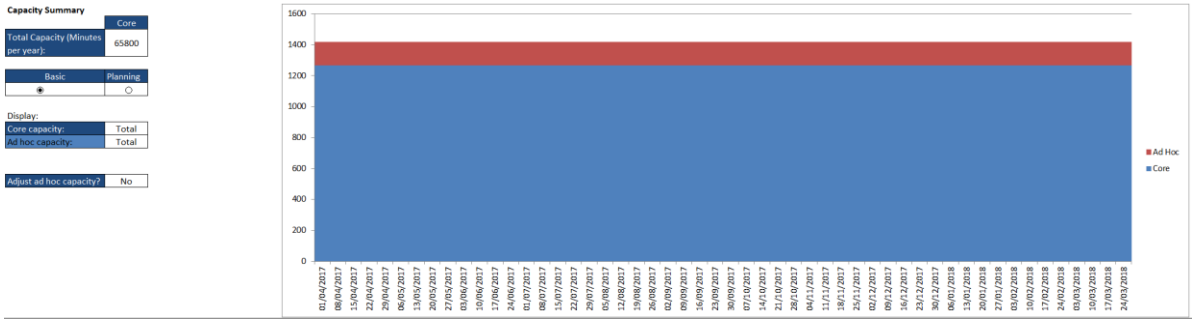

Selecting the ‘Planning’ option will instruct the model to allocate capacity depending on the availability of key resources for clinical sessions, which will generate a more detailed map of your capacity.

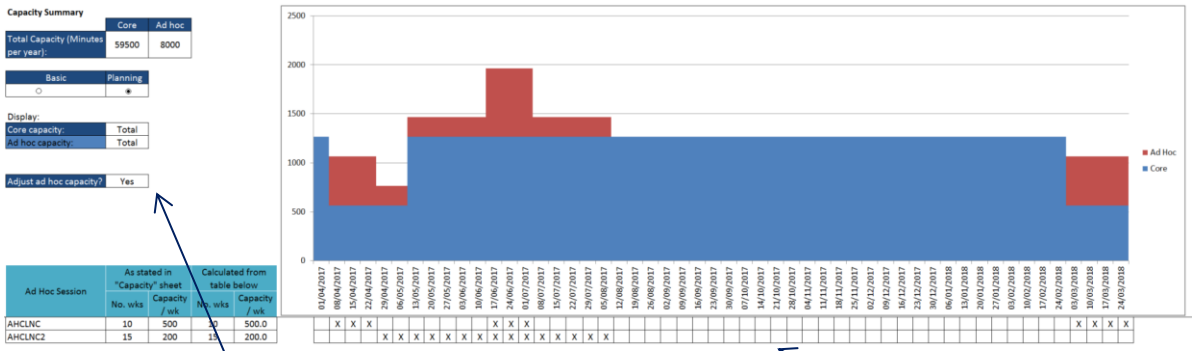

If you enable ‘Adjust ad-hoc capacity’, you can select specific weeks your ad hoc capacity will run.

If you have linked ad hoc sessions to a key resource, they will automatically block weeks in which your key resource is unavailable.

17. Parameters

The ‘Parameters’ sheet is used to establish the behaviour of the service once a request for service has been received.

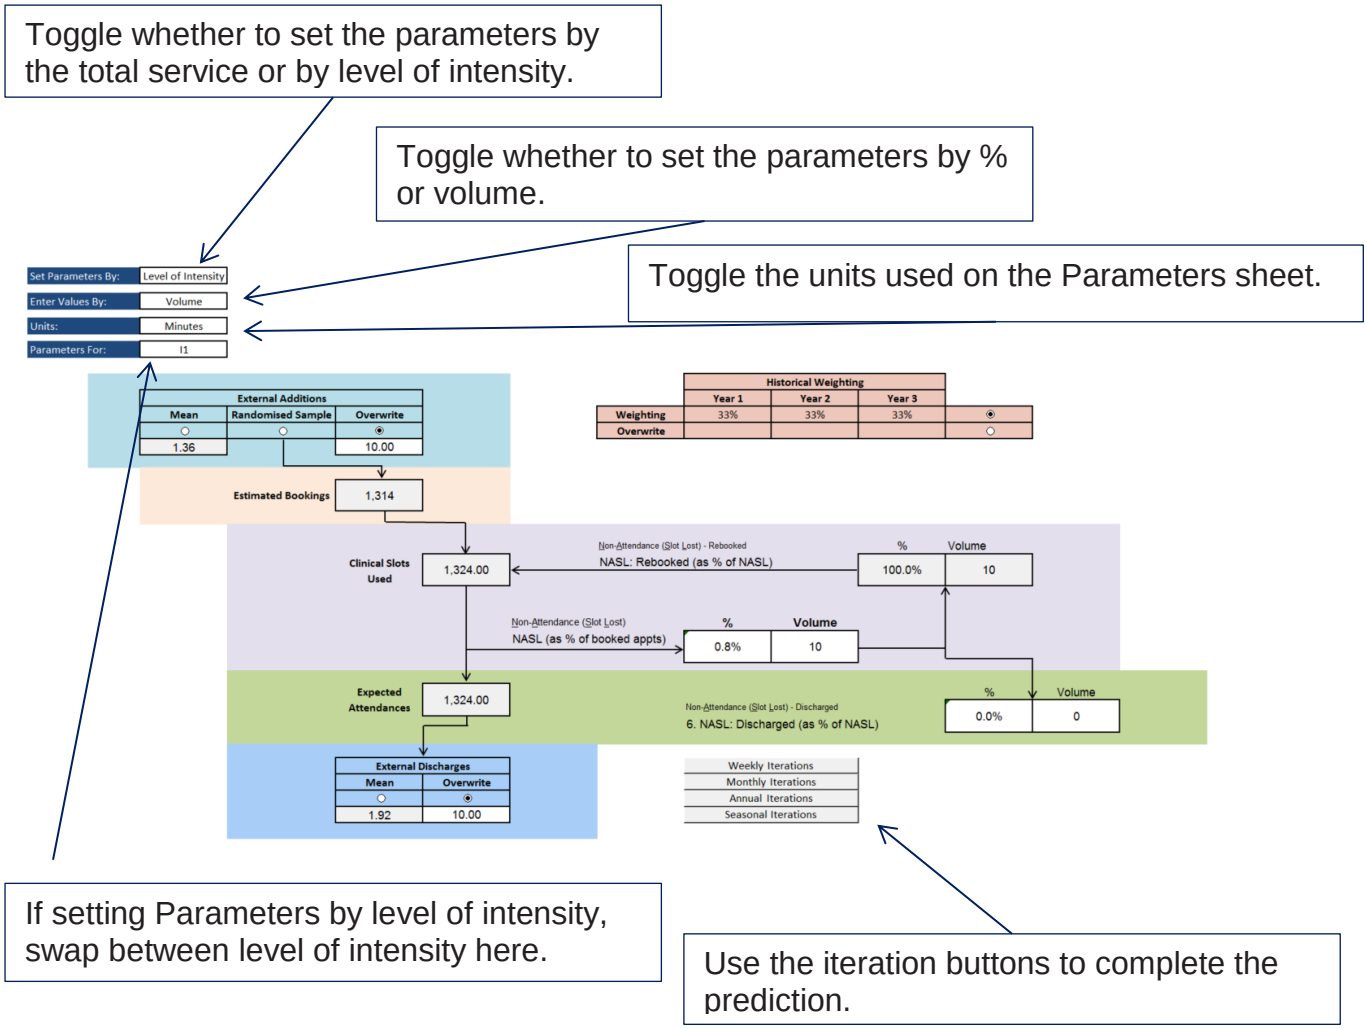

The parameters entered are detailed in the table below.

External additions	Choose the mean of historic data, a bootstrapped sample of historic data, or enter a user defined value, for the external additions per week.
Non-attendance (slot lost) rate	Enter the % or volume of patients who did not attend an appointment, where the clinical slot was lost.
Non-attendance (slot lost) discharge	Enter the % or volume of NASL patients who were subsequently discharged from the service.
External discharges	Choose the mean of historic data or enter a user-defined value for the external discharges per week.
Historic weighting	Edit the relative historic weighting for external additions and external discharges if using the calculated values.

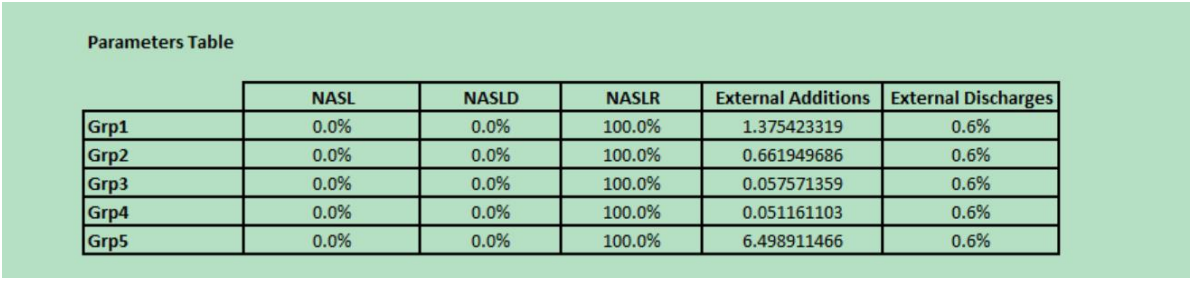

Parameters Table

	NASL	NASLD	NASLR	External Additions	External Discharges
Grp1	0.0%	0.0%	100.0%	1.375423319	0.6%
Grp2	0.0%	0.0%	100.0%	0.661949686	0.6%
Grp3	0.0%	0.0%	100.0%	0.057571359	0.6%
Grp4	0.0%	0.0%	100.0%	0.051161103	0.6%
Grp5	0.0%	0.0%	100.0%	6.498911466	0.6%

The Parameters table displays the parameters used per intensity level. This is particularly useful if parameters are being set per intensity level.

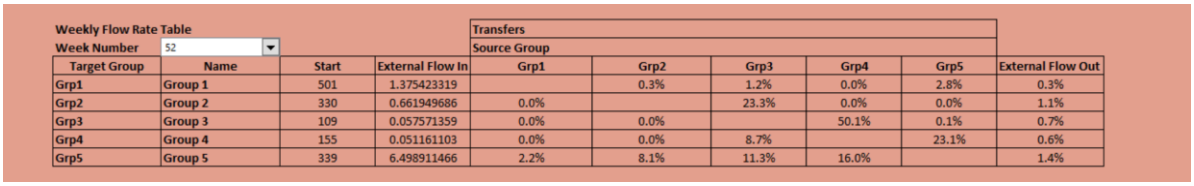

Weekly Flow Rate Table

Week Number 52

				Transfers					
				Source Group					
Target Group	Name	Start	External Flow In	Grp1	Grp2	Grp3	Grp4	Grp5	External Flow Out
Grp1	Group 1	501	1.375423319		0.3%	1.2%	0.0%	2.8%	0.3%
Grp2	Group 2	330	0.661949686	0.0%		23.3%	0.0%	0.0%	1.1%
Grp3	Group 3	109	0.057571359	0.0%	0.0%		50.1%	0.1%	0.7%
Grp4	Group 4	155	0.051161103	0.0%	0.0%	8.7%		23.1%	0.6%
Grp5	Group 5	339	6.498911466	2.2%	8.1%	11.3%	16.0%		1.4%

The flow rate table displays the flow rates from one group to another. The table will update to the flow rate of the last run iteration. For iterations that run over several periods (weekly, monthly and seasonally), the dropdown list can be used to view the flow rates at different stages.

Once the Parameters sheet has been completed, use the iteration buttons to run the prediction. The ‘weekly iterations’ button uses a change in the flow rate (and randomised external additions sample, if necessary) to make the prediction for each of the 52 weeks. If ‘monthly iterations’ is selected, the factors will change every four weeks, for 52 weeks. If ‘annual iterations’ is selected, the factors will be constant for every week, for 52 weeks. If seasons were set in the model, then ‘seasonal iterations’ will be available. The factors will change in synchrony with the seasonal periods selected in the ‘Caseload SPC chart’ sheet.

18. Results (summary)

The ‘Results’ sheet displays the predicted Caseload, Demand and Required Capacity as a total for the service, or for each level of intensity.

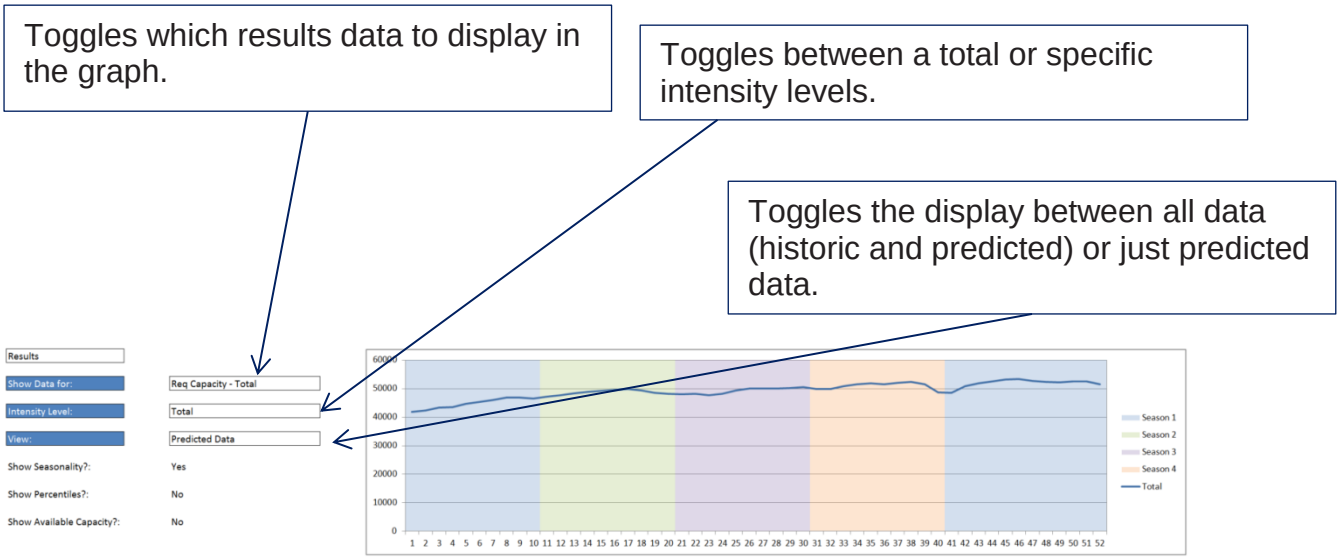

If ‘Predicted data’ is chosen, the following option becomes available:

1. Show seasonality?

    i) Apply the same seasonality profile to the newly predicted data

If any of the required capacity options are selected in the ‘Show Data for:’ field, and predicted data is being displayed, the following options become available:

2. Show percentiles?

    i) Show the 65th and 85th percentile of required capacity

3. Show available capacity?

    i) Show the available capacity for the intensity level selected

Required vs available capacity

The required vs available capacity chart displays the balance of required capacity against available capacity. The coloured bar represents the range of required capacity, and the small diamond represents the available capacity (core and ad hoc).

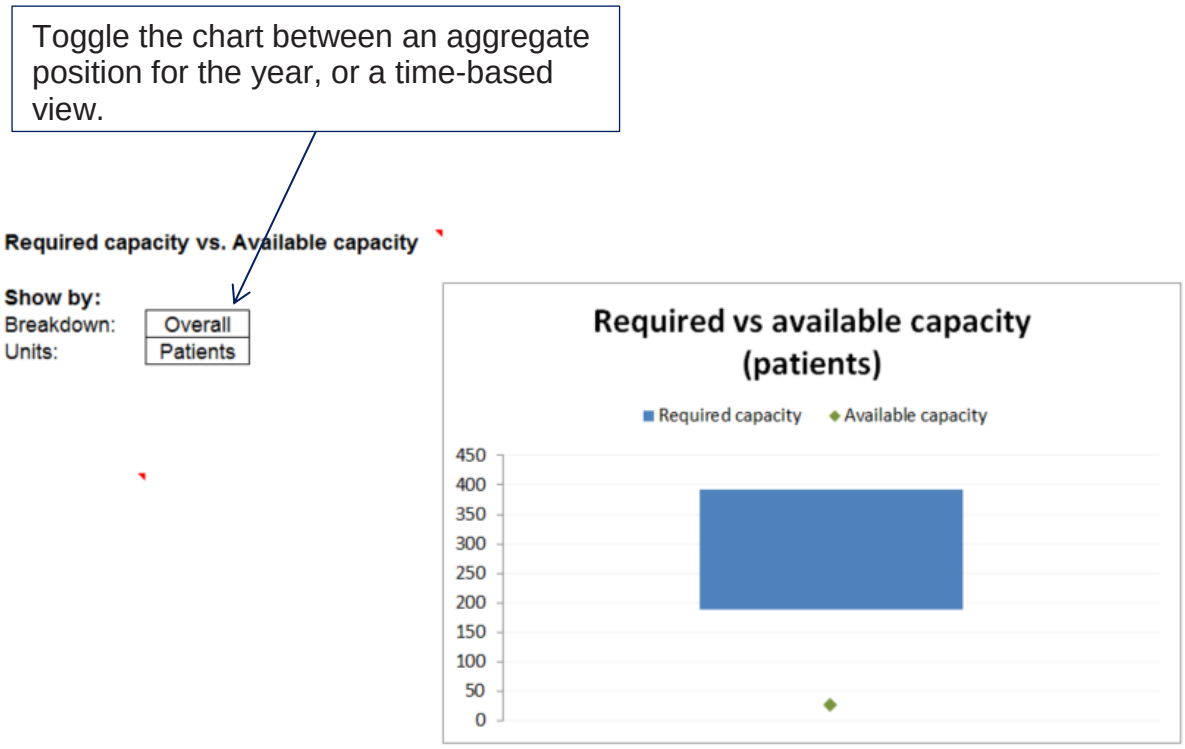

Appendix: Glossary of key terms

Activity

Clinical contact that has taken place. As activity is simply a reflection of what the service is capable of delivering, this is not the same as demand.

Activity can be biased by changes in capacity, or by additional ad hoc capacity such as waiting list initiatives, so we discourage using historic activity as a basis for planning your service.

Capacity

Available capacity is the resource you can deploy to provide a service for your patients. This needs to be operationally verified and compared against the required capacity.

Required capacity is what needs to be provided by your service so that your waiting list does not increase over time. Required capacity is a combination of your demand (adjusted for variation), and the removals and additions to your waiting list resulting from non-attendances and discharges.

Demand

Requests for service – this can be in the form of a referral, a decision to admit (DTA), or even simply an appointment in an earlier part of your service as part of the continuation of a treatment pathway.

Level of Intensity

A unit into which patients can be grouped based on different methods of treatment at different frequencies.

Non-attendances

There are two types:

- non-attendances, where the slot was reused (NASR)

- non-attendances, where the slot was lost (NASL).

Any patients who were discharged from the service subsequent to their non-attendance are described here as NASR (discharge) or NASL (discharge).

Contact us:

NHS Improvement
Wellington House
133-155 Waterloo Road
London
SE1 8UG

0300 123 2257
enquiries@improvement.nhs.uk
improvement.nhs.uk

@NHSImprovement

This publication can be made available in a number of other formats on request.

 Publication code: IG 17/18