
CASAL
(C++ algorithmic stock assessment laboratory)

CASAL User Manual v2.01-2003/08/01

B. Bull
R.I.C.C. Francis
A. Dunn
A. McKenzie
D.J. Gilbert
M.H. Smith

NIWA Technical Report 124
ISSN 1174-2631
2003

CASAL
(C++ algorithmic stock assessment laboratory)

CASAL User Manual v2.01-2003/08/01

B. Bull
R.I.C.C. Francis
A. Dunn
A. McKenzie
D.J. Gilbert
M.H. Smith

NIWA Technical Report 124
2003

**Published by NIWA
Wellington
2003**

Enquiries to:
Science Communication, NIWA,
Private Bag 14901, Wellington, New Zealand

ISSN 1174-2631
ISBN 0-478-23268-3

© NIWA 2003

Citation:
Bull, B.; Francis, R.I.C.C.; Dunn, A.; McKenzie, A.; Gilbert, D.J.; Smith, M.H. (2003).
CASAL (C++ algorithmic stock assessment laboratory): CASAL User Manual v2.01-
2003/08/01. *NIWA Technical Report 124*. 223 p.

*The National Institute of Water and Atmospheric Research
is New Zealand's leading provider
of atmospheric, marine, and
freshwater science*

Visit NIWA's website at <http://www.niwa.co.nz>

CONTENTS

Contents.....	3
Introduction	7
1. Getting started.....	9
1.1 CASAL end user licence	9
1.2 Version	9
1.3 Citing CASAL.....	10
1.4 System requirements	10
1.5 Necessary files	10
1.6 Useful add-ons	10
1.7 Getting help.....	10
1.8 Technical specifications	11
2. Running CASAL	13
2.1 Command line arguments	13
2.2 Running a Bayesian analysis in CASAL.....	15
2.3 Free parameter file format used by CASAL	17
2.4 Constructing a CASAL data file	18
3. Overview of the CASAL model	21
3.1 Model components	21
3.2 Parameters	21
3.3 Observations.....	21
4. The population section.....	23
4.1 Overview	23
4.2 The state object and the partition	24
4.3 The time sequence.....	25
4.4 Transitions between states.....	27
4.4.1 Ageing (in an age-based model)	27
4.4.2 Recruitment.....	28
4.4.3 Maturation.....	31
4.4.4 Migration	32
4.4.5 Growth (in a size-based model).....	33
4.4.6 Mortality (natural and fishing).....	34
4.4.7 Disease mortality	37
4.5 Setting the initial state	37
4.6 Applying ogives	39
4.7 Ogives descriptions	41
4.8 Calculation of size-at-age (in an age-based model)	45
4.9 Calculation of mean weight	46
4.10 Weightless model (running CASAL as a numbers only model).....	47
4.11 Maturity, in models without maturity in the partition.....	47
5. The estimation section.....	49
5.1 Role of the estimation section	49
5.2 Specifying the free parameters.....	49
5.3 Point estimation.....	50
5.4 Likelihood or posterior profiles.....	51
5.5 Bayesian estimation	52
5.6 Observations.....	56
5.6.1 Types of observations	56
5.6.2 Age/size observations	58

5.6.3	Age-at-maturation observations.....	60
5.7	The objective function.....	62
5.7.1	Weighted least-squares	63
5.7.2	Likelihoods	63
5.7.3	Process error	68
5.7.4	Calculating nuisance q 's	68
5.7.5	Priors.....	71
5.7.6	Penalties.....	73
5.7.7	Ageing error.....	75
5.8	Residuals	76
5.9	Generate simulated observations.....	76
6.	The output section	81
6.1	Printouts from CASAL	81
6.2	Output quantities	82
6.3	Projections.....	84
6.3.1	Carrying out projections	84
6.3.2	Calculating projected fishery performance estimators (FPIs).....	86
6.4	Deterministic yield calculations	87
6.4.1	Yield per recruit analyses	87
6.4.2	Deterministic MSY	88
6.5	Stochastic yield estimates	89
6.5.1	MCY/CAY.....	89
6.5.2	Current surplus production (CSP).....	91
7.	The population.csl file	93
7.1	Defining the partition	93
7.2	Defining the annual cycle and the time sequence	94
7.3	Defining recruitment	98
7.4	Defining recruitment variability.....	101
7.5	Defining recruitment in yield simulations.....	102
7.6	Defining growth (in a size based model)	102
7.7	Defining maturation (when maturity is in the partition)	104
7.8	Defining maturity (when maturity is not in the partition).....	104
7.9	Defining migrations	105
7.10	Defining natural mortality.....	106
7.11	Defining fishing mortality.....	108
7.12	Defining disease mortality	109
7.13	Defining selectivities.....	110
7.14	Setting the initial state.....	111
7.15	Defining ogive preferences	113
7.16	Defining size-at-age	113
7.17	Defining the size-weight relationship	116
8.	The estimation.csl file.....	119
8.1	Defining the estimation method	119
8.2	Defining point estimation.....	119
8.3	Defining likelihood or posterior profiling.....	120
8.4	Defining MCMC	120
8.5	Defining the free parameters and priors.....	123
8.6	Defining the relativity constants q	126
8.7	Defining the observations	126
8.8	Defining the objective function associated with the observations	139
8.9	Defining the penalties	143
8.10	Defining the ageing error	148

8.11	CASAL extensions.....	149
9.	The output.csl file.....	151
9.1	Defining the printouts	151
9.2	Defining the output quantities.....	152
9.3	Defining projections.....	156
9.4	Defining yield calculations	156
9.5	Defining deterministic yields	158
9.6	Defining stochastic yields	159
10.	Troubleshooting.....	163
10.1	Typical errors	163
10.2	Other errors	164
10.3	Reporting errors	164
11.	CASAL extensions	167
11.1	New parameterisations (via <code>user.parameterisation.C</code>)	167
11.2	New priors and penalties (via <code>user.prior_penalty.C</code>).....	170
11.3	New likelihoods (via <code>user.likelihood.C</code>).....	172
12.	Post-processing of CASAL output	175
12.1	Introduction.....	175
12.2	<code>extract.header</code>	175
12.3	<code>extract.objective.function</code>	176
12.4	<code>extract.free.parameters</code>	177
12.5	<code>extract.fits</code>	177
12.6	<code>extract.quantities</code>	178
12.7	<code>extract.free.parameters.from.table</code>	179
12.8	<code>extract.quantities.from.table</code>	180
12.9	<code>read.files.for.BOA</code>	181
13.	An example of an age-based model.....	183
13.1	Introduction.....	183
13.2	The input parameter files	183
13.3	CASAL output	191
13.4	Generating simulated output	194
14.	Enhancements and other changes from CASAL v1.02-2002/10/21	197
14.1	Issues fixed.....	197
14.2	Changes and enhancements.....	198
	Acknowledgments	201
	References	203
	Index	205
	Quick reference	210
	Command line arguments.....	210
	Optional command line arguments.....	210
	Order of transitions within a time step	210
	Available ogives	211
	Available penalties	212
	List of commands and sub-commands in the <code>population.csl</code> data input file.....	212
	List of commands and sub-commands in the <code>estimation.csl</code> data input file.....	216
	List of commands and sub-commands in the <code>output.csl</code> data input file	221

INTRODUCTION

CASAL (C++ algorithmic stock assessment laboratory) is a generalised age- or length-structured fish stock assessment model that allows a great deal of flexibility in specifying the population dynamics, parameter estimation, and model outputs.

This manual provides information on how to use CASAL, including how to run CASAL, how to set up the input data files, descriptions of the population dynamics and estimation methods, and how to generate outputs. It also contains a brief overview of the technical specifications of the software, and an example of an age-based model using CASAL.

CASAL is designed for flexibility. It can implement either an age- or length-structured model, optionally also structuring the population by sex, maturity, and/or growth-path. It can be used for a single stock for a single fishery, or for multiple stocks, areas, and/or fishing methods. The user can choose the sequence of events in a model year. The data used can be from many different sources of information, for example catch-at-age or catch-at-size data from commercial fishing, survey and other biomass indices, and survey catch-at-age or catch-at-size data. Estimation can be by least-squares, maximum likelihood, or Bayes.

As well as generating point estimates of the parameters of interest, CASAL can calculate likelihood or posterior profiles and can generate Bayesian posterior distributions using Monte Carlo Markov Chain methods. CASAL can project stock status into the future using stochastic recruitment and can generate a number of yield measures commonly used in New Zealand stock assessment, including MCY, CAY, F_{max} , $F_{0.1}$, deterministic MSY, and CSP.

1. GETTING STARTED

1.1 CASAL end user licence

CASAL (including the software, documentation, examples and other ancillary files) is not free software. You may not distribute CASAL or modify CASAL under any circumstances without the written authorisation of the National Institute of Water and Atmospheric Research Limited of 269 Khyber Pass Road, Newmarket, Auckland (NIWA).

You may use CASAL for non-commercial evaluation purposes only. You may make copies of this software only as reasonably required for backup purposes. You must not distribute, sell or otherwise make CASAL available for use by a third party. You must not use all or any part of CASAL in conjunction with any product or service (including training or consulting) for commercial gain.

NIWA grants you a non-exclusive and non-transferable right to use CASAL only in accordance with the terms of this licence. NIWA reserves the right to refuse to license CASAL to any person without giving reasons thereof. Requests for the use of CASAL outside the terms of this licence should be directed to NIWA (see <http://www.niwa.co.nz/> or email CASAL@niwa.co.nz).

This licence agreement is formed when you accept the terms of this licence by using or running CASAL.

NIWA reserves the copyright and all other intellectual property rights in CASAL.

NIWA makes no representations or warranties regarding the accuracy of CASAL, the use to which CASAL may be put or the results to be obtained from the use of CASAL. Accordingly NIWA accepts no liability for any loss or damage (whether direct or indirect) incurred by any person through the use of or reliance on CASAL.

NIWA is to be acknowledged in publications relating to the use of or from conclusions drawn from CASAL. However, you must not, without written permission, use the name or any trademark or logo of NIWA to claim any sponsorship, endorsement, approval or affiliation or other association with NIWA by virtue of this licence.

The CASAL software, documentation, example and other ancillary files are distributed in the hope that they will be useful for non-commercial evaluation purposes only, but without any warranty; without even the implied warranty of merchantability or fitness for a particular purpose.

This licence is governed by and construed in accordance with the laws of New Zealand.

1.2 Version

This document details the usage of CASAL version v2.01-*yyyy/mm/dd*. The version number printed by CASAL is suffixed with a date in format *yyyy/mm/dd*. This is the last UTC date on which its source files were officially modified. User manual updates will usually be issued for each minor version or date release of CASAL, and can be obtained, on request, from CASAL@niwa.co.nz.

1.3 Citing CASAL

A suitable reference for CASAL and this document is:

Bull, B.; Francis, R.I.C.C.; Dunn, A.; McKenzie, A.; Gilbert, D.J.; Smith, M.H. (2003). CASAL (C++ algorithmic stock assessment laboratory): CASAL User Manual v2.01-2003/08/01. *NIWA Technical Report 124*. 223 p.

1.4 System requirements

CASAL is available for Redhat Linux 7.3 and from the command prompt under most Microsoft Windows operating systems.

Several of CASALs tasks are highly computer intensive and a powerful processor is recommended. We recommend a minimum of 64 megabytes of free RAM for running CASAL (although, depending on the scope of the problem, you may need much more). The program itself requires less than 10 megabytes of hard-disk space but output files can consume large amounts of disk space. Depending on number and type of user output requests, the output could range from a few hundred kilobytes to several hundred megabytes.

1.5 Necessary files

In Linux, only the executable file `casal` is required to run CASAL. In Windows, you need the executable file `casal.exe`.

1.6 Useful add-ons

No software other than the appropriate operating system or emulation package is required to run CASAL. However, as CASAL offers little in the way of post-processing of the output, most users will wish to have a package available that allows tabulation and graphing of model outputs. We recommend the use of software packages such as Microsoft Excel (<http://www.microsoft.com>), S-Plus (<http://www.insightful.com>), or R (<http://www.r-project.org>) (Ihaka & Gentleman 1996).

You may also wish to use the “extract CASAL output” S/S-Plus/R functions for post-processing CASAL output (see Section 12). This is distributed as the script file `extract_CASAL.v2.01.s`.

A useful package for post-processing and analysis of Monte-Carlo Markov Chain (MCMC) Bayesian output is the S-Plus/R library “Bayesian Output Analysis Program (BOA)” — see Smith (2001). Information about this package can be found at <http://www.public-health.uiowa.edu/boa>. A function to read the MCMC output from CASAL for use with BOA is included with `extract_CASAL.v2.01.s`.

1.7 Getting help

CASAL is distributed as unsupported software. NIWA does not provide help for users of CASAL outside of NIWA. While we would appreciate being notified of any problems or errors in CASAL, updates may or may not correct these problems or errors — see Section 10.3 for how to report errors. Information about CASAL can be found at

<http://www.niwa.co.nz/ncfa/tools/casal/>. The maintainer of this software, documentation, and associated files can be contacted at CASAL@niwa.co.nz.

1.8 Technical specifications

CASAL is compiled using `gcc`, a freeware C/C++ compiler developed by the GNU Project (<http://gcc.gnu.org>). The current version has been compiled on Linux using `gcc` version 3.2.3 (20030425) and on Microsoft Windows using MinGW `gcc` version 3.2.3 (mingw special 20030504-1). Note that the output from CASAL may differ slightly on the different platforms due to different precision arithmetic or other platform dependent implementation issues.

CASAL uses a quasi-Newton optimiser and scalar, vector, and matrix types from the `Betadiff` automatic differentiation software package. `Betadiff` emulates most of the functionality of an early version of AUTODIF (Fournier 1994), and is based on a modified version of the program ADOL-C v1.8.4 “A package for automatic differentiation of algorithms written in C/C++” (<http://www.math.tu-dresden.de/~adol-c>) developed by a team including Andreas Griewank (Technical University of Dresden, griewank@math.tu-dresden.de). A suitable reference for ADOL-C is Griewank et al. (1996).

The optimiser used by `Betadiff` is based on the main algorithm of Dennis Jr. & Schnabel (1996).

The random number generator used by CASAL is the `newran` random number generation package (Davies 1998), and uses the Lewis-Goodman-Miller algorithm with Marsaglia mixing.

2. RUNNING CASAL

CASAL is controlled by command line arguments, which are used to tell it what task you want to do, for example, run the model, estimate the parameters, or do a MCMC run. Section 2.1 lists these command line arguments.

CASAL gets most of its information from input data files. The program looks for three files, `population.csl`, `estimation.csl`, and `output.csl`, which contain population, estimation, and output parameters respectively (although the names of these files can be modified, see later). Section 2.4 describes how to construct a CASAL data file — the population, estimation, and output file parameters are listed in Sections 7, 8, and 9 respectively.

Note that the information is read in from the three data files at the start of each CASAL run. As a result, you can change the data files and start another run in the same directory before the first run is finished — providing that you make sure that the outputs of the two runs are sent to different destinations.

CASAL uses both the standard output and standard error; we suggest redirecting both into files. With the bash shell, you can do this using the command structure,

```
(casal [arguments] > out) >& err
```

For `casal -r`, `-e`, or `-E`, the standard output dump can be processed using the extract CASAL output S/S-Plus/R functions (Section 12).

2.1 Command line arguments

The call to CASAL is of the following form.:

```
casal [-l] [-r] [-e] [-E] [-p] [-m] [-a number] [-C filelist
  -S outfile] [-s number prefix] [-v outfile] [-P outfile]
  [-Y] [-f prefix] [-F suffix] [-q] [-i infile]
  [-O outfile] [-o outfile] [-g RNG_seed] [-n name]
```

The call should include exactly one of the following “task” arguments.

- l Display the CASAL end user *licence*.
- r *Run* the population section once only and calculate the objective function (see Section 0). Print out the free parameters, the objective function and its components, the fits and residuals if requested (Section 5.8), and the output quantities (Section 6.2).
- e Calculate the point *estimate* of the parameters (Section 5.3). Print outputs as per -r.
- E As per -e but using finite differences instead of automatic differentiation (Section 5.3).
- p Calculate likelihood or posterior *profiles* (Section 5.4).
- m Use *MCMC* (Section 5.5) to sample the posterior distribution of the parameters. See Section 2.2 for MCMC procedure.
- a [number] Recover the specified MCMC run from its results files; continue the run and *append* further results to the results files. See Section 2.2 for MCMC procedure (see also -n).

- C [*filelist*] *Concatenate* the MCMC results files for the specified files into a single set of samples from the posterior. Reduce the sample size by random or systematic sub-sampling if requested. Optionally apply prior reweighting (Section 5.5). The *filelist* argument should be a list of the (full) names of samples files, separated by white space. Use -S to specify the file into which to dump the results. See Section 2.2 for the MCMC procedure.
- s [*number prefix*] Generate *simulated* observations, i.e., use CASAL as a simulator (Section 5.9). You must use -i to provide the name of a file containing free parameters, either one set (e.g., a point estimate) or multiple sets (e.g., a posterior sample). For each set of parameters supplied, *number* simulations are carried out. The results are dumped to files whose names are generated by combining the filename *prefix* specified, the number of the parameter set, and the number of the individual simulation (e.g. if *prefix* = *my_simulate*, then the third set of simulated observations for the second set of true parameters will be dumped into a file *my_simulate.par2.sim3*). If *number* = 1, then the *.sim[n]* part of the filename is omitted.
- v [*outfile*] Output the *values* of the output quantities (Section 6.2). You must use -i to provide the name of a file containing a posterior sample. Results are dumped into *outfile*. Use this option to analyse the results of a MCMC run.
- P [*outfile*] Calculate *projected* outputs (Section 6.3). You must use -i to provide the name of a file containing free parameters, either one set (i.e., a point estimate) or multiple sets (i.e., a posterior sample). Results are dumped into *outfile*.
- Y Calculate *yield* estimates (Sections 6.4, 6.5), e.g., MCY, CAY, deterministic MSY, CSP.

In addition, you can use any of the following arguments:

- f [*prefix*] Use a prefix on the names of the three input parameter files.
- F [*suffix*] Replace the standard *cs1* suffix used on the input parameter filenames with a user defined suffix.
- q Run *quietly*, i.e., suppress printing from within the population section.
- i [*file*] *Input* one or more sets of free parameter values from a text file.
 - With -r, run the model with each.
 - With -e, do a separate point estimate starting at each.
 - With -p, use the first set as the initial minimum. (The actual minimum, not the starting point of the minimiser. See Section 5.4).
 - With -m, use the first set as the starting point for the pre-MCMC point estimate, the second set (if there are two) as the starting point for the chain, and ignore the rest.
 - With -s, produce simulated observations for each.
 - With -v, calculate output quantities for each.See Section 2.3 for the free parameter file format and Section 2.2 for MCMC procedure.
- O [*outfile*] *Output* (no append) a set of free parameter values to a text file.
 - With -r, output the free parameter values used for the run (either those specified with -i, or the base values in the parameter files if -i is not used)
 - With -e or -E, output the estimated parameter values.

- The text file is in an appropriate format for use with `-i` (Section 2.3).
If the file exists already, it is overwritten.
- `-o [outfile]` *Output* (with append) a set of free parameter values to a text file.
Same as `-O` above, except that if the file exists already, it is appended to.
 - `-S [outfile]` With `-C`, dump the posterior *sub-sample* into `outfile`.
 - `-g [RNG_seed]` With `-m`, `-s`, or `-Y`, seed the random number *generator* with this positive (long) integer value. If this is not specified, then the default is defined as a number based on the computer clock time.
 - `-n [name]` Used with `-m` or `-a` when chains are being carried out on several computers. The argument is the *name* of the current machine, which is inserted into the names of the MCMC results files. That way you can copy the files from all the chains onto a single computer and, because they have different names, they will not overwrite each other.

2.2 Running a Bayesian analysis in CASAL

A full Bayesian analysis is more time consuming than the other CASAL tasks, and involves editing input files to achieve different tasks. The process of how to get CASAL to run a Bayesian analysis is described here. Section 5.5 describes the algorithms used by CASAL.

The first step in producing Monte Carlo Markov Chain (MCMC) results is to do a `-m` run. This produces a single Markov Chain. An initial point estimate is produced before the chain starts. This is done in order to calculate an approximate covariance matrix of the free parameters, but may also be used as the starting point of the chain. You can specify the free parameter values used as the starting point of the point estimation (as the first row of the file invoked with `-i`), and also optionally specify the free parameter values from which to start the chain (as the second row of the file invoked with `-i`). If you have specified that a free parameter is fixed in MCMC, you still need to supply a value for it when using `-i`. Once the MCMC commences (as opposed to the initial point estimate), the parameter will be fixed at the supplied value.

The `-m` run produces two results files. The first file is `samples.[run number]`, or if the `-n` option is set it is `samples.[name].[run number]`. It uses the free parameter file format described in Section 2.3, i.e., a header row followed by many rows of parameter values. The second file is `objectives.[run number]` or `objectives.[name].[run number]`. It contains the standard output header produced by CASAL, the covariance matrix used (if the covariance matrix is modified at one or more iterations of the chain, only the initial version of the matrix is shown), and a columnar table (with one row per posterior sample, giving the sample number, the posterior, prior, likelihood, and combined penalties (all on the negative log-scale), the current step size, the acceptance rate so far, and the number of times the covariance matrix has been modified so far.

The run number is the first available positive integer, i.e., if the directory already contains a file `samples.[name].1` but not `samples.[name].2` or `objectives.[name].2`, the next run number will be 2. Incidentally, can we suggest that the *first* thing you do after an `-m` run is *back up the results files*. It may be distressing to inadvertently lose your only copy of a chain that had been running for some time.

If your chain gets interrupted for some reason, you can use `casal -a [run number]` to continue it rather than starting again from scratch (also include the argument `-n [name]` if this was used). Make sure to use the same parameter files in the rerun as in the original run

(CASAL does not check). You will not get the same results as you would have if the original chain had continued, because the random number sequence will be different. If the original run was interrupted or crashed, make sure that the last lines of the samples and objectives files were complete and that they each have the same number of lines once the headers are removed (the printing process might have stopped partway through a line, in which case CASAL would be confused by the partly finished results).

You may want to run multiple Markov Chains simultaneously if you have the hardware resources. If you are using a shared file system, you can run multiple chains in the same working directory. If your chains all use the same data files and you want to combine them later to produce a single posterior sample, use `-n` with each chain with a different 'machine name' argument to send the results to a different file. All the chains will then share the same run number. Give each chain a different random number seed using `-g` (or the results may be identical).

Use `casal -C` to combine a list of MCMC results files into a single posterior sample, decimate it down to a sub-sample of a specified, manageable size, and apply posterior reweighting if requested. You need to provide the samples file names, which are supplied after the `-C` (the objective files with the same suffixes should also be present), and the name of the file into which the sub-sample is dumped (in the free parameter file format described in Section 2.3), which is supplied with `-S`. You should also set the burn-in period at this stage.

Use `casal -v` to calculate output quantities for a posterior sample, either the sub-sample generated by `casal -C` or the original single-chain sample generated by `casal -m`.

An example of running a Bayesian analysis

A typical sequence for a ten-chain MCMC might be as follows. First, generate the chains, specifying random number seeds and machine numbers,

```
(machine 1): casal -m -g 144 -n PC1
(machine 2): casal -m -g 1812 -n PC2
...
(machine 10): casal -m -g 71 -n PC10
```

Files such as `samples.PC4.1` will be generated (assuming this is the first MCMC run in the directory). Back them up. Uh-oh: machine 2 had a power failure. To resume the run from where it stopped,

```
(machine 2): casal -a 1 -g 1812 -n PC2
```

CASAL finds the previous output files, `samples.PC1.1` and `objectives.PC1.1`, resumes the MCMC where they ended, and appends the results from the rest of the chain to these files.

Following this, copy all the output files onto one of the ten computers and run them through an external MCMC diagnostics package. Next pool the ten chains (and sub-sample to reduce the size of the result),

```
casal -C samples.PC1.1 ... samples.PC10.1 -S subsample.dat
```

having first added the following to `estimation.csl`,

```
@MCMC
burn_in 100000
```

A sub-sample file is generated. Summarise the posterior,

```
casal -v quantities.dat -i subsample.dat
```

Suppose that you want to check out the effect of using a different prior. Change the `estimation.csl` file to specify the new prior, add the following to `estimation.csl`,

```
@MCMC
prior_reweighting 1
```

and then repeat the last commands,

```
casal -C samples.PC1.1 ... samples.PC10.1 -S subsample.2.dat
casal -v quantities.2.dat -i subsample.2.dat
```

Finally, note that files in free parameter file format (including the posterior samples output by `casal -m` and `-C`), and tables of output quantities (including the output of `casal -v` and `-P`) can be read into S/S-Plus/R using the functions in Section 12.

2.3 Free parameter file format used by CASAL

In various situations it is useful to either write sets of free parameters to a file or to read sets of free parameters from a file. For example:

- When doing MCMC with `-m`, a long list of sets of parameter values is generated. They are saved to disk (so that they don't consume memory and so that they can be recovered if the program crashes partway through the chain) and can be re-loaded later on.
- When a point estimate has been calculated with `-e`, the user may want to save the parameter values and reload them later. For example, when running the model with `-r` at the 'optimal' point.
- When a point estimate has been calculated by another stock assessment package, the user may want to run CASAL using the parameter values estimated by the other package.

The same *free parameter file format* is used in all cases. There is one header line, consisting of the name of each parameter (in `command[label].subcommand` format), followed by the length if it is a vector parameter, separated by single spaces. The header is followed by 1 or more sets of parameters, each written as a long vector on a single line.

A simple example of this is,

```
initialization.B0 size_at_age.k 1 recruitment.YCS 30
10000 0.1 0.87 0.95 1.12 ... (27 more YCS)
15000 0.2 0.93 0.98 1.14 ...
```

Note that the 1 argument for `size_at_age.k` is compulsory — it is not a scalar but a vector of length 1.

For input, the header must be exactly accurate or the program will reject the file. This is done to check that the right parameters have been provided in the right order. However, there is no check that the right number of parameters are supplied in the rows of the table.

Note that CASAL generates a line of data as output, suitable for use in a file with `-i`, automatically when doing a run or an estimation (i.e., with `-r`, `-e`, or `-E`). To use this as an input for a subsequent run, copy the appropriate lines (i.e., the lines immediately after “In a format suitable for `-i` :”) into a text file, and rerun CASAL with the `-i [file]` option. CASAL can also generate a free parameter file using the `-o` or `-O` options (the first appends to a file if it already exists, the second replaces it) automatically from a `-r`, `-e` or `-E` run (in this case, if `-e` or `-E` is used with a multi-row `-i` input file, then multiple estimations will be done and multiple rows will be written to the `-o` or `-O` file).

Free parameter files can be read into S/S-Plus/R using the functions in Section 12.

2.4 Constructing a CASAL data file

The model is specified to CASAL via the population, estimation, and output parameters. These are specified in the `population.csl`, `estimation.csl`, and `output.csl` input data files (though you can modify these names using the `-f` and `-F` command-line options). All the parameters that can be used are listed in Sections 7, 8, and 9 respectively.

The parameter files use the *command-block format*. A parameter file consists of any number of command-blocks in any order. Each command-block either consists of a single command (starting with the symbol `@`) and its arguments, or a command (starting with `@`) and an optional label and one or more subcommands, i.e.,

```
@command arguments
or
@command [label]
[subcommand arguments]
[subcommand arguments]
[...]
```

Blank lines are ignored, as is extra white space between arguments. Comments beginning with `#` are removed. If you want to remove a group of commands or subcommands using `#`, then comment out the whole block, not just the first line. Alternatively, you can comment out an entire block by placing curly brackets around the text that you want to comment out. Put in a `{` as the first character on the line to start the comment block, then end it with `}`. All lines (including line breaks) between `{` and `}` inclusive are ignored. (These should ideally be the first character on a line, but if not, then the entire line will be treated as part of the comment block.)

Don't put extra white space before a `@` character (which must also be the first character on the line). Make sure the file ends with a carriage return. Commands and subcommands must consist of letters and/or underscores, and must not contain a full-point (`.`).

There is no need to mark the end of a command block. This is automatically recognised by either the end of the file or the start of the next command block, which is marked by the `@` on the first character of a line.

Also note that the commands, sub-commands, and arguments in the parameter files are case sensitive.

Some commands can never have subcommands (such as @initial). If a command has no subcommands, then it has to have arguments, which are placed on the same line as the command.

All other commands have no arguments, but have subcommands instead. Also,

- Some commands can be used multiple times in the same parameter file and must have a different label each time (such as @abundance).
- Some commands can be used only once and may never have a label (such as @annual_cycle).
- Some commands can be used one or more times: if used once they don't need a label, but if used more than once they do need labels (such as @recruitment).
- Some commands can be used one or more times and don't need labels: they are internally labelled 1 the first time they are used, then 2, 3... (such as @growth).

The parameter listings say what kind of label or argument each command and subcommand takes. Arguments can be of the following types:

<i>switch</i>	true/false
<i>integer</i>	an integer
<i>constant</i>	a real number
<i>estimable</i>	a real number which can be estimated
<i>constant vector</i>	a vector of real numbers
<i>estimable vector</i>	a vector of real numbers which can be estimated
<i>ogive</i>	an ogive which can be estimated
<i>string</i>	a string
<i>vector of strings</i>	a list of strings.

Parameters of type *constant vector*, *estimable vector*, or *vector of strings* contain one or more entries separated by white space (tabs or spaces).

Switches are parameters which are either true or false. Enter 'true' as true, t, or 1, and 'false' as false, f, or 0. Note that this is one of the few situations where CASAL is case insensitive.

Ogive parameters (Section 4.6) are the most complex to enter. You need to specify the type of the ogive, then the ogive parameters. For example, a logistic selectivity ogive with the label 'trawl' with parameter values $a_{50} = 5$, $a_{1095} = 2$ might be entered as,

```
@selectivity trawl
all logistic 5 2
```

where 'all' specifies that this ogive applies to all fish, i.e., males and females, mature and immature. If you want a size-based ogive in an age-based model, you need to insert the word 'sizebased' between the subcommand and the ogive type. For example,

```
@selectivity trawl
all sizebased knife_edge 30
```

See Section 4.6 for an explanation of how the size-based ogive is converted to an age-based ogive.

Not all parameters can be estimated — only those of type *estimable*, *estimable vector* or *ogive* can be estimated. You decide which of these CASAL should estimate, the *free parameters* (Section 5.2). Sometimes an ogive has some non-estimable parameters, for example, an `allvalues_bounded` ogive has two non-estimable parameters, the lower and upper bounds — the remaining parameters give the values between these bounds and can be estimated normally (as a single vector parameter).

When CASAL processes these files, it translates each command and each subcommand into a parameter. Each parameter has a name. For commands, the parameter name is simply the command name. For subcommands, the parameter name format is either,

1. `command[label].subcommand` if the command has a label, or
2. `command[i].subcommand` if the command is occurring for the *i*th time and is auto-numbered, or
3. `command.subcommand` if the command has no label and is not auto-numbered.

The user needs to convert commands to parameter names in this way in several situations. For example, if you have constant natural mortality,

```
@natural_mortality
all 0.3
```

and you want to estimate the mortality rate M , you need to tell CASAL that the parameter named `natural_mortality.all` is to be estimated, by putting in commands,

```
@estimate
parameter natural_mortality.all
```

Similarly, if you have,

```
@selectivity trawl
male logistic 5 2
...
```

and you want to apply some kind of penalty to the logistic selectivity ogive, you will need to tell CASAL that the parameter named `selectivity[trawl].male` is to be penalised

3. OVERVIEW OF THE CASAL MODEL

3.1 Model components

A fisheries model in CASAL consists of three parts.

1. The *population section* is the model of the fish population dynamics. It includes processes such as recruitment, migration, natural and fishing mortality.
2. The *estimation section* carries out the estimation of free parameters. The estimation will be based on an *objective function* (weighted sum of squares, negative log likelihood, negative log posterior, etc.). The estimation section is used to find a *point estimate*, which is the set of parameter values that minimises the objective function. It may also be used to characterise the uncertainty in the point estimate, via either profiling or producing a Bayesian posterior.
3. The *output section* produces results for the user. These may include parameter estimates, the objective function and its components, fits and residuals, projections, yield estimates, etc.

3.2 Parameters

Parameters are quantities that describe how things work. There are three types:

1. *population*: both those related to population biology, e.g., size at age, weight at size, maturation, stock-recruitment relationship, natural mortality, migration parameters, and those concerning the fishery, e.g., catches, selectivity ogives, maximum exploitation rates.
2. *estimation*: needed for the estimation procedure, e.g., choice of estimation method, observations and their error structures or weights, which parameters are to be estimated, priors, starting values, minimiser control values.
3. *output*: indicating which outputs the program should produce, e.g., what should be printed as the model runs, which quantities should be written to file, etc.

Some parameters may function as switches, allowing the user to choose between available options (e.g., between Ricker or Beverton & Holt stock-recruitment relationships, or between normal or lognormal distributions).

Each time a model is run the population and estimation parameters will fall into two classes; those which are assumed *known*, and those that are *free* (i.e., to be estimated). It is up to the user to specify which parameters are free. Not all parameters are *estimable*. Some, such as switches, would never be estimated. (Note that while CASAL may allow a parameter to be estimated, this does not mean that the modeller should necessarily allow it to be estimated.)

3.3 Observations

Observations are data which allow us to make inferences about a fishery (i.e., to estimate parameters). Examples include CPUE indices, survey biomass estimates, catch at age, commercial catch length frequencies, etc. The process of estimation in CASAL involves finding values for each of the free parameters so that each observation is as close as possible to a corresponding expected value. Note that catches are treated as population parameters, not observations.

4. THE POPULATION SECTION

4.1 Overview

The basic structure of a CASAL population model is defined in terms of an *annual cycle*, *time steps*, *states*, and *transitions*.

The *annual cycle* defines what processes happen in each model year, and in what sequence. (In line with the New Zealand fisheries management framework, CASAL runs on an annual cycle rather than, for example, a 6-monthly cycle.)

Each year is split up into one or more *time steps*, with at least one process occurring in each time step. You can think of each time step as representing a particular part of the calendar year, or you can just treat them as an abstract sequence of events.

The *state* is the current status of the population, at any given time. The state can change one or more times in every time step of every year. The state object must contain sufficient information to figure out the future course of the fishery (given a model and a complete set of parameters).

There are a number of possible changes in the state, which are called *transitions*. These include processes such as *recruitment*, *natural mortality*, *fishing mortality*, *disease mortality*, *ageing*, and *migration*.

The division of the year into an arbitrary number of time steps allows the user to specify the exact order in which processes and observations occur. The user needs to specify the time step in which each process occurs. If you ask for more than one process to occur in the same time step, there is a default order in which they occur (see Section 4.3). If you don't want things to happen in this default order, just split them into different time steps.

The key element of the state is the *partition*. This is a broadly applicable concept that can be used to describe many different kinds of fish model. The partition is simply a breakdown of the total number of fish in the current population into different kinds of fish. (Note that the partition records numbers of fish, not biomass.) The fish are categorised by various *characters*. The permissible characters are: size class or age class, sex, maturity, area, stock, and growth-path. The user chooses:

- Whether the partition is subdivided by size class or age class (not both).
- Which of the other characters are included in the partition.
- The number of areas, stocks, or growth paths (if any of these characters are included in the partition).

The resulting partition can be conceptualised as a matrix, where the columns are size or age classes and the rows represent combinations of the other characters. Then the number in each cell of the matrix is the number of fish with the corresponding combination of characters.

For an example of these ideas, consider a model of a single stock with a spawning and non-spawning fishery. The non-spawning fishery happens over most of the year (say 10 months) in the home area. The mature fish then migrate to the spawning area, where the spawning fishery operates. At the end of spawning, these fish, along with the recruits from the previous year, migrate back to the home area. The modeller decides that fish will be divided in the partition by age, sex, maturity, and area (spawning and home grounds). So the partition has 8 rows ($2 \text{ sexes} \times (\text{mature or immature}) \times 2 \text{ areas}$) and one column per age class.

The modeller decides to use the annual cycle in Table 1.

Table 1: The annual cycle of a simple model.

Time	Time step	Area		Activity
		Non-spawning	Spawning	
Jan–Oct	1	Mature and immature fish	Empty	Fishing in the home area.
End of Oct	2	Immature fish	Mature fish	Mature fish migrate to the spawning area.
Nov–Dec	3	Immature fish	Mature fish	Fishery in the spawning area.
End of Dec	4	Mature and immature fish	Empty	Recruits from the previous year appear in the spawning area. Along with the mature fish, they migrate to the home area.

So they define four time steps, labelled 1 through 4. Step 1 includes the non-spawning fishery. Step 2 includes the migration to the spawning area. Step 3 includes the spawning fishery. Step 4 includes recruitment and the migration back to the home area. (In fact, they could have used only 3 time steps, by using a single step in place of their steps 2 and 3. Because the default order of processes within a time step places migrations before fisheries, the processes would still have occurred in the right order.) There are other details to be sorted out, such as the proportion of natural mortality occurring in each time step, but this gives the basic idea.

This structure can be used to implement complex models, with intermingling of separate stocks, with complex migration patterns over multiple areas, and multiple fisheries using different fishing methods and covering different areas and times. Note that there is little point in using a complex structure to model a stock when there are no observations to support that structure. In other words, you should use a structure for your model that is compatible with the data you have available.

The model is run from an *initial* year up to the *current* year. It can also be run past the current year to make *projections* — things that happen in the future — up to the *final* year. Alternatively, for yield calculations, it is run over an abstract *simulation period*.

4.2 The state object and the partition

The key component of the state object is the partition, a matrix of *numbers of fish* by combinations of characters. The columns can either be age or size classes, the rows are combinations of the following characters:

- Sex (male or female).
- Area (any number of areas, named by the user).
- Stock (any number of stocks, named by the user).
- Maturity (immature or mature).
- Growth-path (any number of growth-paths).

A stock is defined as a subpopulation of fish which recruits separately. See Section 4.11 for the treatment of maturity when it is not a character in the partition.

Growth-paths are a feature used to implement some persistence of size at age in an age-based model that uses some length or size data. Each growth-path has its own growth curve, and the size-based model features will hence have different effects on different growth-paths. So, you need to tell CASAL the following:

- Whether the model is age- or size-based.
- The number and nature of size classes in a size-based model.

- The minimum and maximum age classes in an age-based model.
- Whether there is a plus group.
- Whether the partition is divided by sex.
- Whether the partition is divided by maturity.
- Whether the partition has growth-paths, and, if so, how many.
- Whether the partition has multiple stocks, and, if so, how many, and their names.
- Whether the partition has multiple areas, and, if so, how many, and their names.

Age classes are always 1 year wide, except that the maximum age group can optionally be a plus group. You need to choose the minimum and maximum age classes. Size classes are defined by the user. You need to specify how many size classes there are, the lower bound of each size class, and whether the last size class is a plus group, or if not, what its upper bound is. The relevant parameters are `class_mins` and `plus_group`. The `class_mins` parameter contains the lower bound of each class, and concludes with the upper bound of the last class if it is not a plus group. If, for example, you wanted size classes of 30–40, 40–50, 50–60, and 60–70+ cm, in which case you would set `class_mins 30 40 50 60` and `plus_group true`. Whereas if you wanted 30–40, 40–50, 50–60, and 60–70 cm, you would set `class_mins 30 40 50 60 70` and `plus_group false`.

The user can specify that some combinations of characters are not possible. For example, immature fish might never occur in the area you have labelled `spawn_ground`. To do this, you use the `exclusions` parameters. In this case, you would set,

```
exclusions_char1 maturity
exclusions_val1 immature
exclusions_char2 area
exclusions_val2 spawn_ground
```

It's a good idea to use the `exclusions` parameter wherever it is appropriate because it reduces the size of the partition (so, with the above example, there will be no rows in the partition corresponding to immature fish in area `spawn_ground`) and can save memory and calculation time.

The other component of the state object in CASAL is a vector of spawning stock biomasses (SSBs, mid-spawning season biomasses of spawning fish) for each stock. CASAL needs to include this in the state object so as to calculate future recruitments, if there is a stock-recruitment relationship.

4.3 The time sequence

The time sequence of the population model includes the years over which it is to run and the annual cycle for each year. The model runs from the start of year `initial` and runs to the end of year `current`. Projections extend up to the end of year `final`. The annual cycle can contain the following transition processes:

- Ageing (in an age-based model).
- Recruitment.
- Maturation (if maturity is a character in the partition).
- Migration (if the model includes more than one area).
- Growth (in a size-based model).
- Natural and fishing mortality.
- Disease mortality.

If two or more processes are specified for the same time step then they will happen in the above order. This ordering is imposed only to simplify the specification of the annual cycle. It does not restrict the user because it applies only to processes within the same time step. If, for example, it is desired that maturation occur before recruitment then this can be done by putting these processes in separate time steps.

The basic unit of fishing mortality is a *fishery*, defined as fishing mortality in a single area and time step. You may need to split a single administrative fishery into multiple CASAL fisheries, in which case you will need to partition the catch. (However this should often be avoidable. If you have an observation partway through a fishery, you can specify that a certain proportion of the mortality occurs before the observation, without needing to split the time step into two.)

If there is more than one stock, recruitment is handled separately for each stock, but all stocks must recruit in the same time step. There can be more than one maturation episode per year, each of which can apply to only one stock, or all stocks equally. Similarly there can be more than one growth episode per year, each of which can apply to only one stock, or all stocks equally. The user can define any number of migrations in a given year.

To specify the time sequence, you need to tell CASAL the following:

- The initial, current, and final years.
- The number of time steps in each year.
- The time step in which recruitment occurs, and the area to which each stock recruits
- How SSB is calculated¹.
- In an age-based model, the time step at which ages are incremented.
- If there are any migrations, the time step at which each migration occurs and the source and destination areas. Note that if there are multiple migrations in a time step and an area is the source of more than one migration, then the migrations will happen in the order that they are defined in the `population.csl` file.
- If maturity is a partition character, the number of maturation episodes per year, and the time step at which each maturation episode occurs.
- In a size-based model, the number of growth episodes per year, and the time step at which each growth episode occurs.
- In an age-based model, the proportion of the year's growth which has occurred by the start of each time step².

¹ The SSB (spawning stock biomass) is a common model output and is also the measure of abundance used in stock-recruitment relationships in CASAL (where applicable). Different models define SSB in quite different ways so we allow several options in CASAL as to how SSB is calculated. By default, SSB is calculated for each stock as the mature biomass (of both sexes), in an area of your choice, halfway through the natural and fishing mortality in a time step of your choice. It can alternatively be calculated after some other specified proportion of the mortality (see Section 4.4.6). A 'proportion spawning' multiplier can be applied to the mature biomass to get the SSB (in multi-area models this would typically not be done, instead the appropriate proportion of fish would be migrated to the spawning area). If maturity is not in the partition, then the modeller may nevertheless know that all fish in the spawning area should be mature (i.e., because only mature fish are meant to migrate) but the model does not 'know' this because maturity is not persistent. In this case the user can specify that the SSB is the total biomass in the area, rather than using the mature biomass.

² Fish growth in an age-based model is handled quite differently from a size-based model. The simplest option is to assume that the mean size of a fish is based on its age, rounded down to the next lowest whole number of years. So, for example, 2-year old fish have the same mean size whether they have just passed their 2nd birthday or whether they are about to turn 3. An alternative is to allow some fish growth between birthdays. You can do this using the `growth_props` parameter. This is a vector with one entry per time step. The mean size of fish of age a years (rounded down) in the i th time step is

- The proportion of the year's natural mortality occurring in each time step.
- If there is a disease mortality event, in which time step this occurs.
- The time step and area in which each fishery occurs.
- Whether fishing mortality is instantaneous or uses the Baranov equation¹.

You then need to provide CASAL with details about how each process works. These processes are described individually in Section 4.4.

When you define your annual cycle, there are a number of errors you can make. Some of the less obvious ones are listed here. It is an error if:

- The sum of the proportions of the year's natural mortality over time steps is not 1.
- In an age-based model, any element of `growth_props` is outside [0,1]; or if `growth_props` is not 0 in the time step in which fish age; or if `growth_props` diminishes between consecutive time steps without age incrementation having taken place.
- In a size-based model, more than one growth episode occurs in the same time step, unless they involve different stocks.
- You want to use the Baranov equation and there is a time step that includes two or more fisheries in the same area.

4.4 Transitions between states

This section describes the various transition processes in CASAL. The transition processes, in their default order, are:

1. Ageing i.e., age incrementation (in an age-based model)
2. Recruitment
3. Maturation (if maturity is a character in the partition)
4. Migration (in a multi-area model)
5. Growth (in a size-based model)
6. Mortality (natural and fishing)
7. Disease Mortality

4.4.1 Ageing (in an age-based model)

The ageing process is straightforward. Every fish increases in age by one year, except those already in the plus age group (if it exists), which are unaffected. (Note that if there is no plus group in the partition, then all fish older than the maximum age are “dropped off” the end of the partition, i.e., die.)

calculated as if their age was $(a + \text{growth_props}[i])$. So, if the first entry of `growth_props` is 0.5, then, in time step 1, the mean size of 2-year-old fish is calculated as if they were age 2.5. The default is `growth_props = 0` (i.e., no growth between birthdays).

¹ Natural mortality and fishing mortality occurring in the same area and time step can be sequenced in two different ways. The first option is to apply half the natural mortality, then to apply the mortalities from all the fisheries instantaneously, then to apply the remaining half of the natural mortality. The second options is to use the Baranov catch equation, which implies that natural and fishing mortalities are simultaneous. We prefer the first option — the calculations are more straightforward and the result typically about the same. However you can use Baranov if you want, except that we have not yet implemented the Baranov equation for multiple fisheries in the same area in the same time step. Whichever option you use is applied to all fisheries. More on this in Section 4.4.6.

4.4.2 Recruitment

A number of fish are added to the partition. In an age-based model, all recruiting fish are of the minimum age. In a size-based model, you need to tell CASAL the mean and c.v. of the size distribution of recruiting fish, which is assumed to be a normal distribution (and can depend on sex and stock). Note that fish below the minimum of the range that defines the first size class appear in that size class; and similarly fish above the maximum of the range that defines the last size class appear in that size class.

For each stock, the number of fish added in year y is

$$R_y = R_0 \times YCS_{y-y_{enter}} \times SR(SSB_{y-y_{enter}}) \times CR(T_{y-y_{enter}})$$

where R_0 is the stock's average recruitment (ignoring the stock-recruitment and climate-recruitment functions); YCS are year class strength multipliers (also known as recruitment multipliers); y_{enter} is the number of years after it is spawned that a year class enters the partition; SR is the stock-recruitment function ($SR \equiv 1$ if there is no stock-recruitment relationship); CR is the climate-recruitment function (T is a single exogenous variable such as sea surface temperature, $CR(T) \equiv 1$ if no climate-recruitment relationship).

R_0 is an important parameter because it defines how large the stock would be, on average, if there were no fishing. From R_0 , CASAL can calculate B_0 , which is defined to be the SSB that would exist if recruitment were equal to R_0 every year and there were no fishing (alternatively, CASAL can calculate R_0 from B_0 , if the latter is specified). B_0 has several special roles in CASAL: in the stock-recruitment function (where, by definition, $SR(B_0) = 1$); and as a reference biomass in stock projections (see Section 6.3.2) and yield calculations (Sections 6.4.2 and 6.5.1).

You should provide YCS s starting from year (`initial - y_{enter}`) and extending up to year (`current - y_{enter}`).

It can be a bit tricky to figure out what y_{enter} should be. In an age-based model, this depends on the order of recruitment, ageing, and spawning processes within a year:

- If recruitment then ageing then spawning, then y_{enter} should equal `min_age + 1`.
- If spawning then ageing then recruitment, then y_{enter} should equal `min_age - 1`.
- If any other order, then y_{enter} should equal `min_age`.

CASAL will output a warning if the value of y_{enter} you supply does not obey the above rule, but will continue running.

The stock-recruitment functions available are Beverton-Holt and Ricker (the alternative is no stock-recruitment relationship, $SR \equiv 1$). These are parameterised by the parameter *steepness*, defined as $h = SR(0.2 B_0)$. The functional forms for these relationships are:

$$\text{Beverton-Holt: } SR(SSB) = \frac{SSB}{B_0} \left/ \left(1 - \frac{5h-1}{4h} \left(1 - \frac{SSB}{B_0} \right) \right) \right.$$

$$\text{Ricker: } SR(SSB) = \frac{SSB}{B_0} \left(\frac{1}{5h} \right)^{\frac{5}{4} \left(\frac{SSB}{B_0} - 1 \right)}$$

The basic climate-recruitment relationships available are *exponential*, *arctan*, *logistic* (the alternative is no relationship). All three are functions of a single exogenous variable T , which should be provided for years (`initial - y_enter`) to (`current - y_enter`) at least, and can also extend further into the future (for use in projections). Two additional climate-recruitment functions have been added for situations where the ‘climate variable’ T is actually a prediction of year class strength, typically from a climate-recruitment regression analysis (e.g., Bull & Livingston 2001). The *identity* climate-recruitment relationship allows the predictions to be used in an unmodified form. The *linear-combination* climate-recruitment relationship allows the model to decide how much credence to give the predictions, when you estimate the parameter p (which must be between 0 and 1, otherwise you can potentially get negative recruitments).

exponential:	$CR(T) = \alpha \exp(\beta T)$
arctan:	$CR(T) = \alpha (0.5 + \tan^{-1}(\beta T) / \pi)$
logistic:	$CR(T) = \alpha / (1 + \beta \exp(\beta_2 T))$
identity:	$CR(T) = T$
linear-combination:	$CR(T) = pT + (1 - p)$

Note that this formulation allows various levels of relationship between recruitment and climate. At one extreme, when $CR(T) \equiv 1$, there is no relationship. At the other extreme, when the YCS are constant and $SR \equiv 1$, recruitment is completely determined by climate (apart from the factor R_0). In between these extremes, climate affects recruitment but does not determine it.

Warning: Since the climate-recruitment relationship was coded into CASAL, it has become apparent that some aspects do not work as intended. The climate-recruitment option has been marked obsolete and will not be usable until it is repaired in a future version.

As an option, you can use the initial recruitment $R_{initial}$ (see Section 4.5) as the recruitment for the first $n_{rinitial}$ years of the model. So, for fish recruiting in years `initial` to `initial + n_rinitial - 1`, the $R_0 \times YCS$ term of the recruitment equation above is replaced by $R_{initial}$. (Or, if $R_{initial}$ is defined as a deviate, by $R_0 \times R_{initial}$, see Section 4.5.) This option is used to avoid estimating year class strengths about which there is little information. If you use it, you don’t need to provide the early year class strengths. Supply YCS starting from year (`initial - y_enter + n_rinitial`).

It will usually be a good idea to provide a penalty function (see Section 5.7.6) to force the YCS s to average 1. This ensures that the average recruitment for the years in which YCS s are estimated is close to R_0 . Unfortunately, this penalty function may need to be large (i.e., have a large weight), which can lead to poor MCMC performance in the calculation of a Bayesian posterior. Because of that, CASAL supports two alternative parameterisations of YCS s.

The first of these alternatives (the Haist parameterisation) was suggested by V. Haist. Here, the model parameter YCS is a vector Y , covering years from `initial - y_enter + n_rinitial` to `current - y_enter`. The year class strengths are calculated by $YCS_i = Y_i / \text{mean}(Y_i)$ where the mean is calculated over the user-specified years `first_free` to `last_free`. Then,

$$YCS_i = \begin{cases} Y_i / \text{mean}_{i \in R} (Y_i) & [i \in R] \\ Y_i & [i \notin R] \end{cases}$$

where R is the set of years from `first_free` to `last_free`. One effect of this parameterisation is that R_0 is now defined to be the mean estimated recruitment over the years `first_free` to `last_free` (because the mean YCS over these years will always be 1). Often, the user will wish to force $Y_i = 1$ for $i \notin R$ (this is equivalent to forcing $R_i = R_0$) by setting the lower and upper bounds to be 1. An exception to this might occur for the most recent YCS s, which the user may want to estimate, but not include in the definition of R_0 (because the estimates are based on too few data).

The advantage of the Haist parameterisation is that the user need no longer use a large penalty to force the mean of the YCS parameter to be 1 (though they should still use a small penalty to stop the mean of Y from drifting). This may improve MCMC performance. Simulated and projected YCS are not affected by this feature, nor are those YCS that are set to $R_{initial}$. A disadvantage with this parameterisation in a Bayesian analysis is that the prior refers to the Y 's, *not* the YCS .

The second alternative is the Francis parameterisation of YCS . This uses two distinct concepts of mean recruitment: R_{mean} is the theoretical mean recruitment over all years (past and future), and, as in the Haist parameterisation, R_0 is the mean over the user-specified years `first_free` to `last_free`. There are two corresponding biomasses: B_{mean} is the biomass that would exist if recruitment was always equal to R_{mean} and there was no fishing, and B_0 is the analogous SSB with constant recruitment R_0 . With this parameterisation, R_{mean} is used in place of R_0 in the calculation of R_y , so

$$R_y = R_{mean} \times YCS_{y-y_enter} \times SR(SSB_{y-y_enter}) \times CR(T_{y-y_enter})$$

The user may also force the recruitment to be equal to R_0 for years at the beginning and end of the period `initial - y_enter + n_rinitial` to `current - y_enter`. This can be achieved by providing YCS s for a subset of this period. CASAL will set the recruitment equal to R_0 for all years from `initial - y_enter + n_rinitial` to the year before `YCS_years`, and also for any years after `YCS_years` and up to and including `current - y_enter`. For these years, CASAL replaces YCS_{y-y_enter} in the above equation by \bar{Y} , which is the mean YCS calculated over the years `first_free` to `last_free`.

With this parameterisation, R_0 (and thus B_0) become derived parameters, which are calculated from the user-specified R_{mean} (or B_{mean}) using the equation $R_0 = R_{mean} \bar{Y}$. Note that the only uses of R_{mean} in CASAL are to calculate R_0 and R_y . Also, the “special roles” of B_0 (see above) are unchanged by the Francis parameterisation.

Two advantages of the Francis parameterisation are that there is no need for a penalty function to constrain the YCS s, and the prior distributions specified for parameter YCS do apply to the YCS s (not true for the Haist parameterisation). A disadvantage is the need for the additional parameter R_{mean} . With likelihood estimation, this parameter is not well determined, because if we double R_{mean} and halve all the YCS s we do not affect either the biomass trajectory or the fit to any observations. This will not be a problem with Bayesian estimation unless the priors on R_{mean} and the YCS s are both uniform (not recommended). Because the estimated value of R_{mean} depends on these priors it seems best to treat this parameter, and the associated B_{mean} , as nuisance parameters with little biological meaning.

Incidentally, the output documentation (Section 6.2) refers to ‘`true_YCS`’, which are defined as the $YCS \times CR \times SR$ part of the recruitment equation. These are more informative than the YCS alone, when there is a climate-recruitment or stock-recruitment relationship — let alone the Y 's.

So, to specify the recruitment for each stock, you need to tell CASAL the following:

1. *YCS*, starting from year ($\text{initial} - y_{\text{enter}} + n_{\text{rinitial}}$) and extending up to year ($\text{current} - y_{\text{enter}}$). With the Francis parameterisation you may provide *YCS* for a consecutive subset of these years.
2. The value of y_{enter} .
3. The stock-recruitment function (if any) and the `steepness` parameter.
4. The climate-recruitment function (if any) and the values of the climate-recruitment parameters.
5. In a sexed model, the proportion of recruits which are male.
6. In a size-based model, the mean and c.v. of the size distribution of recruiting fish (which can depend on fish sex).
7. In a growth-path model, the proportion of recruiting fish on each growth-path.
8. If R_{initial} is to be used as the recruitment for the first n_{rinitial} years of the model, the value of n_{rinitial} .
9. If you want to use the Haist or Francis parameterisations of year class strengths, you need to say so, and specify the range of free *YCS*.

4.4.3 Maturation

Maturation is the process in which immature fish become mature and are moved accordingly in the partition. See Section 4.11 for how to treat maturity when it is not a character in the partition.

You can specify a single maturation episode in each year, or you can have multiple maturation episodes. Each episode can apply to one stock, or all stocks equally, and can be applied in one area, or all areas equally. Maturation rates are expressed as an ogive (and note that this ogive contains the rates of maturation, not the proportions of mature fish).

If you try to mature fish in an area where fish are constrained to be immature, CASAL will issue a warning, and will not mature those fish.

So, to specify each maturation episode, you need to specify the following:

- If it applies to only one stock, which is it?
- If it applies to only one area, which is it?
- The maturation rates, as an ogive, optionally by sex.

4.4.4 Migration

Migration is the process of moving fish from one area to another. It only occurs in multi-area models. You can specify any number of migrations occurring in each year. If two or more migrations are specified in the same time step then they take place in the order in which they are given.

A migration can involve only one stock in an area, or all stocks. You can migrate immature fish only, or mature fish only, or both. You can state that a given proportion of these fish migrate (constant across all age or size classes), or you can provide an ogive of proportions migrating by age or size class.

You cannot migrate fish to an area where their combination of characters is not allowed (CASAL errors out). So, for example, if you are moving fish to an area where only mature fish are allowed, you need to specify that only mature fish migrate.

CASAL currently supports two-wave migrations. These migrations consist of two waves in different time steps. If p_i is the specified proportion of fish migrating from the i th partition element, proportion (pwave p_i) will migrate in wave 1 and proportion $(1 - \text{pwave}) \times p_i / (1 - (\text{pwave} \times p_i))$ will migrate in wave 2. Specify these as two separate migrations, give pwave for each, and specify that the first is a '1st wave' and that the second is a '2nd wave'. (No checking is currently carried out that there are two matching waves with the same parameters. Remember that you should specify pwave for each, not pwave for the first and $(1 - \text{pwave})$ for the second. If you want to estimate pwave, you need to set the estimate.same parameter to make sure that pwave takes the same value for both waves.)

CASAL also supports density-dependent migrations. These allow you to make the migration rate depend on the fish abundance in the source area, the destination area, or both. So, you can encourage fish to move into an under populated area and/or out of an overpopulated area.

In each year y , for each density dependent migration from area a to area b , a 'density dependence factor' is calculated as

$$F_{a,b}(y) = \exp\left(-S\left(\frac{A_{a,y} - A_{a,0}}{A_{a,0}}\right) - D\left(\frac{A_{b,y} - A_{b,0}}{A_{b,0}}\right)\right)$$

where S is a number expressing the dependence on the abundance in the source area (negative values mean that fish are encouraged to leave an overpopulated area. Set $S = 0$ for no dependence); D is a number expressing the dependence on the abundance in the destination area (positive values mean that fish are encouraged to move to an under populated area — set $D = 0$ for no dependence); $A_{j,y}$ is the total abundance of all fish in area j , year y (with $y = 0$ meaning the unfished equilibrium level) just before the migration occurs.

Note that a single density dependence factor is applied to all fish in a given migration in a given year, regardless of age, sex, etc.

Now let $P_{a,b}^i(y)$ be the proportion of fish in element i of the partition which migrate from area a to area b in year y , prior to the application of density dependence. (These values depend on the migration rate, or ogive of migration rates, etc.) And let the corresponding odds be

$$O_{a,b}^i(y) = \frac{P_{a,b}^i(y)}{1 - P_{a,b}^i(y)}.$$

Then the effect of density dependence is to change the odds to

$$\vartheta_{a,b}^i(y) = O_{a,b}^i(y) \times F_{a,b}(y)$$

and hence the proportion of fish migrating to

$$\Pi_{a,b}^i(y) = \frac{\vartheta_{a,b}^i(y)}{1 + \vartheta_{a,b}^i(y)}.$$

Density dependence is never applied during the calculation of the initial state.

The specification of the annual cycle includes the time step, source area, and destination area of each migration. You also need to tell CASAL the following:

- If there are multiple stocks and only one stock migrates, which is it?
- Do only mature fish migrate, or immature fish, or both?
- If a proportion of these fish migrate (constant across age or size classes), what is it? Or, if fish migrate according to an ogive across age or size classes, what is it?
- Is density dependence applied? If so, what are the values of the density dependence parameters S and D ?

Two-wave migrations require more details — see earlier.

4.4.5 Growth (in a size-based model)

In a size-based model, growth is the process by which fish move between size classes in the partition. See Section 4.3 for the treatment of fish growth in an age-based model.

You can specify a single growth episode in each year, or you can have multiple growths. Each episode can apply to one stock, or all stocks equally, and applies to all areas equally.

There are many possible fish growth models. So far, CASAL only implements the Francis (Francis 1988) parameterisation of the von-Bertalanffy growth curve (see Figure 1). This is referred to as the ‘basic’ model.

Here, we assume that the growth of fish in size class i is normally distributed with mean

$$\mu = g_a + (g_\beta - g_a)(l_{ci} - l_a) / (l_\beta - l_a),$$

and standard deviation

$$\sigma = \max(c\mu, s_{min}),$$

where l_i is the lower size bound of this size class and $l_{ci} = 0.5(l_i + l_{i+1})$.

For $i < j$, the $[i,j]$ th element of the transition matrix (which defines what proportion of the i th size class move to the j th size class) is simply the integral of this distribution between the bounds $(l_j - l_{ci})$ and $(l_{j+1} - l_{ci})$.

If there is a plus group, the corresponding integrals extend to ∞ . The $[i,i]$ th element is the integral between the bounds $-\infty$ and $(l_{i+1} - l_{ci})$.




Figure 1: Graphical representation of the parameterisation of the von-Bertalanffy growth curve.

So, you need to tell CASAL the following, for each growth episode:

1. If there are multiple stocks and only one stock grows, which is it?
2. The growth model to be used (so far there is only the “basic” model above).
3. The parameters of the growth model, with reference sizes l_α and l_β , the corresponding mean growth rates g_α and g_β , a c.v. c and a minimum standard deviation s_{min} . All of these may depend on sex and maturity.

4.4.6 Mortality (natural and fishing)

Mortality includes natural and fishing mortality — the processes by which fish are removed from the partition. CASAL combines the two processes when they occur in the same time step, hence they are discussed in a single section here.

Each time step can include a proportion of the year’s natural mortality and/or one or more fisheries. Natural mortality is applied to all areas and can depend on sex, maturity, stock, and age or size class. A *fishery* is defined as fishing mortality in a specified area and time step. You need to supply a catch for each fishery in each year.

Natural mortality and fishing mortality occurring in the same area and time step can be sequenced in two different ways. The first option, *instantaneous mortality*, is to apply half the natural mortality, then to apply the mortalities from all the fisheries instantaneously, then to apply the remaining half of the natural mortality. The second option is to use the *Baranov* catch equation, which implies that natural and fishing mortalities are simultaneous. In general, the first option is recommended because it requires much less computation. Note that the use

of the Baranov equation for multiple fisheries in the same area in the same time step has not yet been implemented. Whichever option you use is applied to all fisheries.

With *instantaneous mortality*, the following equations are used.

1. An exploitation rate (actually a proportion) is calculated for each fishery, as the catch over the selected biomass,

$$U_f = \frac{C_f}{\sum_{\substack{\text{rows of the partition } i, \\ \text{age/size classes } j}} \bar{w}_{ij} S_{fij} n_{ij} \exp(-0.5tM_{ij})}$$

where, for element $[i,j]$ of the partition, i indexes the rows of the partition for the area in which fishery f operates, S_{fij} is the selectivity for fishery f , \bar{w}_{ij} is the mean weight, n_{ij} is the pre-mortality number of fish, M_{ij} is the natural mortality, and t is the proportion of the year's natural mortality in the time step.

2. The fishing pressure associated with fishery f is defined as the maximum proportion of fish taken from any element of the partition in the area affected by fishery f ,

$$U_{obs}(f) = \max_{i,j} \left(\sum_{\substack{\text{fisheries } k \\ \text{in the same area} \\ \text{and time step} \\ \text{as fishery } f}} S_{kij} U_k \right)$$

(Not, as in some other models, as the catch over the vulnerable biomass.)

There is a maximum fishing pressure limit of $U_{max}(f)$ for each fishery f . So, no more than proportion $U_{max}(f)$ can be taken from any element of the partition affected by fishery f by fisheries in that time step. Clearly $0 \leq U_{max} \leq 1$. It is an error if two fisheries sharing the same area and time step do not have the same U_{max} .

For each f , if $U_{obs}(f) > U_{max}(f)$, then U_f is multiplied by $U_{max}(f) / U_{obs}(f)$. The fishing pressures are recalculated, and stored if requested.

3. The partition is updated using

$$n'_{ij} = n_{ij} \exp(-tM_{ij}) \left(1 - \sum_{\text{fisheries } f} S_{fij} U_f \right)$$

With *Baranov mortality*, the following equations are used:

1. For each fishery, calculate the fishing mortality rate by solving the following Baranov equation for F_f :

$$C_f = \sum_{\substack{\text{rows } i, \\ \text{age/size classes } j}} \left(\frac{F_f S_{fij}}{tM_{ij} + F_f S_{fij}} \bar{w}_{ij} n_{ij} \left(1 - \exp(- (tM_{ij} + F_f S_{fij})) \right) \right)$$

where C_f is the catch weight and F_f the instantaneous fishing mortality rate for fishery f . There is no closed form solution for F_f given the other parameters, so this equation must be solved iteratively for F_f .

2. The fishing pressure for fishery f is defined as the maximum instantaneous fishing mortality rate for any element of the partition in the area affected by fishery f . Since there can be no more than one fishery per area per time step, the fishing pressure is

$$F_{obs}(f) = F_f \max_{i,j} (S_{fij})$$

There is a maximum fishing pressure limit of $F_{max}(f)$ for each fishery. So, $F_{max}(f)$ is the maximum instantaneous fishing mortality rate on fish affected by the fishery.

For each f , if $F_{obs}(f) > F_{max}(f)$, then F_f is reduced to $F_{max}(f) / \max_{i,j}(S_{fij})$. The fishing pressures are recalculated, and stored if requested. (Note that $F_{max}(f)$ is the maximum permissible value of $F_f \times S_{fij}$ not of F_f . This is confusing, but is allowed so-as to maintain compatibility with previous NIWA software. This is another reason why we do not recommend the Baranov option for use in new models.)

3. The partition is updated using

$$n'_{ij} = n_{ij} \exp\left(-\left(tM_{ij} + F_f S_{fij}\right)\right)$$

where f is the fishery affecting row i in the time period (if none, then $F_f = 0$).

Your `population.csl` data file should contain a list of selectivities. Each fishery should use one of these selectivities. More than one fishery can share the same selectivity. Also fisheries can share selectivities with observations (for example, a CPUE index could use the same selectivity as the corresponding fishery).

Note that if there are not enough fish to take the catch, CASAL simply reduces the actual catch below the specified catch. If you are estimating parameters, a parameter set which leads to fishing pressure limits being exceeded is not automatically disallowed. So your point estimate may break fishing pressure limits. If you want to prevent this (as is generally the case), you will need to add catch limit penalties in the estimation section (Section 5.7.6).

You can specify that observations occur partway through a mortality episode, or that SSBs are calculated partway through mortality. Either way, CASAL needs a method of determining the contents of the partition “after a given proportion p of the mortality”. There are two options:

1. *Weighted sum*: after proportion p of the episode, the partition elements are given by $n_{ij}^p = (1-p)n_{ij} + pn'_{ij}$. Arguably this is the most natural approach if Baranov is not used, although unless $p = 0, 0.5, \text{ or } 1$ it's not logically consistent with the half- M , fishing, half- M sequence used in the instantaneous mortality option.
2. *Weighted product*: after proportion p of the episode, the partition elements are given by $n_{ij}^p = n_{ij}^{1-p} n'_{ij}^p$. This is the most natural approach if Baranov is used, although it might be desirable to use ‘weighted sum’ instead for consistency with analyses not using Baranov.

When the Baranov equation is used, CASAL gives the user the option of specifying an F for some years rather than a catch in tonnes. This is intended for modelling the early history of a

fishery, if catches were not recorded but the modeller has a vague idea about the level of historical fishing pressure. Be clear that this F is an instantaneous mortality at a selectivity of 1, and that individual partition elements may suffer more or less mortality, depending on the selectivity.

Annual selectivity shifts are also provided for. These allow selectivities to shift to the left or right with changes in an exogenous variable. (In the 2002 hoki assessment, this exogenous variable is either related to the depth being fished or the time of the fishing season, see Francis et al. 2003.) The ogive is shifted by $a_f(E_f - \bar{E}_f)$, where a_f is a shift factor (presumably estimated) and E_f is the exogenous variable. This is accomplished by changing the parameters of the ogive, for example, in a logistic ogive, the a_{50} parameter is shifted. Not all ogives support this feature (see Section 4.6 for a complete list). For size-based ogives in an age-based model, the shift is applied before the ogive is converted to age-based.

So, to specify the mortality processes, you need to tell CASAL the following:

- The value of M , which may depend on sex, maturity, age and/or size.
- The total catch for each fishery in each year.
- Which selectivity is used by each fishery.
- The maximum fishing pressure limit for each fishery, as U_{max} for instantaneous mortality or F_{max} for Baranov mortality.
- Whether you want to use the ‘weighted sum’ or ‘weighted product’ approach to calculate the contents of the partition partway through a mortality episode.
- Optionally if Baranov is used, the instantaneous mortality F to be applied, by year, for a range of years that does not overlap with the range of years for which catches are provided.
- The details of each selectivity, which may include a exogenous shift variable E and a shift parameter a .

4.4.7 Disease mortality

Disease mortality is a special, additional, mortality that is implemented to occur after natural and fishing mortality during a time step. This process removes fish from the partition, is applied to all areas, and can depend on sex/age/size class. It can only occur during one time step in the annual cycle.

The partition is updated using

$$n'_{ij} = n_{ij} \exp\left(-\left(t_{year} M_d S_{ij}\right)\right),$$

where M_d is an the disease mortality rate to apply, t is an annual multiplicative scalar (and can be used to index the years in which disease mortality is applied), and S_{ij} is a selectivity to apply to the disease mortality over the sex/age/size classes. As earlier, your `population.csl` data file should contain a list of selectivities. The disease mortality should use one of these selectivities.

4.5 Setting the initial state

Before setting the initial state of the population you need to supply the equilibrium abundance for each stock. Usually, this is done by specifying either B_0 (equilibrium SSB) or R_0 (equilibrium constant recruitment level, as a number of fish). If you specify B_0 it is used to

calculate R_0 , and conversely. Alternatively, if the Francis parameterisation of year-class strengths is used you must supply either B_{mean} or R_{mean} , rather than B_0 or R_0 . If you specify B_{mean} it is used to calculate R_{mean} , and conversely; in either case, CASAL calculates R_0 from R_{mean} (see Section 4.4.2).

CASAL has an alternative parameterisation of equilibrium abundance for use in two-stock models only. You can specify R_0 or B_0 as the sum over stocks (optionally on the log-scale) and the proportion in each stock. An analogous option based on R_{mean} and B_{mean} is available for when the Francis parameterisation of year-class strengths is used.

CASAL offers the following three methods for setting the initial state of the population.

1. Use the equilibrium state based on constant recruitment R_0 .
2. Allow the initial abundance to be different from the equilibrium abundance. You need to supply an initial abundance $B_{initial}$ or $R_{initial}$ for each stock as well as B_0 or R_0 (or, if the Francis parameterisation is used, B_{mean} or R_{mean}). The equilibrium state is calculated, then the numbers of fish of each stock s are multiplied by $B_{initial}(s) / B_0(s)$ if you supplied $B_{initial}$, or by $R_{initial}(s) / R_0(s)$ if you supplied $R_{initial}$. There is also an option for you to express $R_{initial}$ as a deviate, i.e., supply $R_{initial}$ relative to R_0 , in which case the numbers of fish of each stock s are multiplied by $R_{initial}$.
3. Allow the initial age or size distribution to be different from the equilibrium distribution. (Using this option has approximately the same effect as starting the model some years earlier and estimating the earliest year class strengths.) You need to supply an initial number of fish $C_{initial, i}$ for each age or size class i of each stock. The equilibrium state is calculated, then the numbers in each age or size class i are multiplied by a factor such that they sum to the relevant $C_{initial, i}$. (Alternatively, you can specify $C_{initial}$ separately for males and females.) Then $B_{initial}$ is calculated for each stock by running the model forwards for one year, with constant recruitment at equilibrium levels and no fishing, and recording the SSB. (The model is put back to the initial state after doing this.) $R_{initial}$ is calculated as $(B_{initial} / B_0) \times R_0$. The SSBs for all years before the initial year are set to $B_{initial}$ (perhaps not ideal, but CASAL needs to fill them in with something in case they are needed for the stock-recruitment relationship or if you ask for them to be printed out).

CASALs current algorithm for determining the equilibrium state in a size-based model involves running the model over a number of simulated years with constant recruitment. You need to tell it how many years to use; this would usually be the approximate maximum age of the fish.

So, to specify the initial state of the population, you need to supply:

1. R_0 for each stock, or B_0 for each stock (or, if the Francis parameterisation of year-class strengths is used, R_{mean} or B_{mean}). (For a two-stock model you can use the alternative parameterisation above).
2. In a size-based model, the number of years in the constant-recruitment simulations used to determine the equilibrium state.
3. If you want the initial abundance to be able to differ from the equilibrium abundance, then $B_{initial}$ for each stock, or $R_{initial}$ for each stock (optionally, relative to R_0).

4. If you want the initial age or size distribution to be able to differ from the equilibrium age or size distribution, then $C_{initial, i}$ (or $C_{initial_male, i}$ and $C_{initial_female, i}$) for each age or size class i of each stock.

4.6 Applying ogives

An ogive is a function with a different value for each age or size class (i.e., for each column of the partition). Ogives are used frequently throughout the CASAL population section: for selectivity curves (Section 4.4.6), rates of migration (Section 4.4.4), and maturation rates (Section 4.4.3).

Ogives have a number of different parametric forms in CASAL and you can use any of these for any ogive parameter. Some common parameterisations are `logistic`, `knife_edge`, `double_normal`, and the most flexible parameterisation `allvalues` where each ogive element is specified separately. See Section 2.4 for instructions on specifying ogives in CASAL. Note also that some ogive forms can be shifted (see Section 4.4.6).

An ogive may be defined to apply just to some subgroup of fish. For example, `rates_male_logistic` would be used to describe a logistic migration ogive for males, and `male_mature_logistic` would be used for a logistic selectivity ogive to be applied only to mature males. See Sections 7.8, 7.9, and 7.12 for the permissible subgroup descriptors for maturation, migration, and selectivity ogives, respectively. In the following examples we use `subgroup` as a generic subgroup descriptor.

The usage of ogives depends on whether the model is age- or size-based. Ogives can be:

1. *Age-based in an age-based model*

The ogive is indexed by fish age, with indices from `min_age` to `max_age`.

For example, you might have an age-based selectivity that was logistic with 50% mark at age 5 and 95% mark at age 7. This would be defined by `subgroup_logistic`, $a_{50} = 5$, $a_{1095} = (7 - 5) = 2$. Then the value of the ogive at age $x = 3$ is $1/\left[1+19^{(a_{50}-x)/a_{1095}}\right] = 1/\left[1+19^{(5-3)/2}\right]$.

2. *Size-based in a size-based model*

The ogive is indexed by fish size class, with indices from 1 to `n_classes`. The value of the ogive for each size class is a function of the class midpoint. A plus size group has no midpoint, of course, so if you have a plus size group you need to assign it a nominal midpoint using the `plus_group_size` parameter (which is also used to calculate mean weight for the plus group, see Section 4.9).

For example, you might have size classes of 30–40, 40–50, 50–60, 60–70, and 70+ cm, and want a size-based selectivity that was logistic with 50% mark at 55 cm and 95% mark at 75 cm. This would be defined by `subgroup_logistic`, $a_{50} = 55$, $a_{1095} = (75 - 55) = 20$. Then the value of the ogive for the second size class is $1/\left[1+19^{(55-45)/20}\right]$.

3. Size-based in an age-based model

This allows you to add size-based model features to your age-based model, for example a size-based selectivity. The value of the ogive for each element of the partition is the integral of the size-based ogive over the distribution of fish sizes (which depends on age, and potentially on the other partition characters, the year, and the time step, see Section 4.8).

For example, you might have a size-based selectivity that was logistic with 50% mark at 55 cm and 95% mark at 75 cm. This would be defined by `subgroup size_based logistic, a50 = 55, a95 = (75 - 55) = 20`. Suppose the partition is divided by maturity, sex, and age, and that 3-year-old mature male fish in time step 2 have a mean size of 62 cm, and a normal size distribution with a c.v. of 0.2. Then the value of the ogive, for 3-year-old mature male fish in time step 2, is

$$\int L(x) s(x) dx,$$

where $L(x)$ is the logistic ogive $= 1/\left[1 + 19^{(55-x)/20}\right]$,

and $s(x)$ is the probability density function of the fish sizes,

$$s(x) = \frac{1}{\sqrt{2\pi} (62 \cdot 0.2)} \exp\left(-0.5 \left(\frac{x - 62}{62 \cdot 0.2}\right)^2\right).$$

CASAL calculates the above integral by a discrete approximation. It takes n_{quant} evenly spaced quantiles of the specified fish size distribution (defined as the quantiles of $((1 \dots n_{quant}) - 0.5)/n_{quant}$), evaluates the ogive at each, and calculates the average of the ogive values. By default $n_{quant} = 5$. This default will generally be adequate, unless your size-based ogives are very steep (e.g., knife-edge), in this case you may find that the resulting age-based ogives are quite discretised. Fix this problem by increasing the value of n_{quant} . Note that decreasing n_{quant} to 1 effectively bases the ogive on the mean size at age, and ignores the distribution of sizes at age (and reduces the computational cost considerably).

Note that the use of n_{quant} does not effect other uses of variation of length at age in the model, i.e., age/size observations (Section 4.8) or mean weight at size (Section 4.9).

Not all types of ogives can be used as size-based ogives in an age-based model. The permitted types are specified below.

So far, the use of size-based ogives in an age-based model where size-at-age varies from year to year is only implemented for selectivities — not proportions maturing, migration rates, etc. (Whereas if size-at-age does not vary between years, then you can use size-based versions of any kind of ogive.)

Note that the function values for some choices of parameters for some ogives can result in a computer numeric overflow error (i.e., the number calculated from parameter values is either too large or too small to be represented in computer memory). CASAL implements range checks on some parameters to test for a possible numeric overflow error before attempting to calculate function values. For example, the logistic ogive is implemented such that if $(a_{50} - x)/a_{95} > 5$) then the value of the ogive at x is zero, i.e., for $a_{50}=5, a_{95}=0.1$, then the value of

the ogive at $x=1$, without range checking would be 7.1×10^{-52} . With range checking, that value is 0 (as $(a_{50-x})/a_{t0.95} = 40 > 5$).

4.7 Ogives descriptions

The available ogives are:

constant

$$f(x) = C$$

The `constant` ogive has the estimable parameter C . This ogive can be shifted (trivially), and can be used as a size-based ogive in an age-based model.

knife_edge

$$\begin{aligned} f(x) &= 0, & (x < E) \\ &= 1, & (x \geq E) \end{aligned}$$

The `knife_edge` ogive has the non-estimable parameter E , and cannot be shifted. (It might seem straightforward to shift a `knife_edge` ogive, just by changing E , however this cannot work in a gradient-based minimiser, as the test of $(x < E)$ is not differentiable). The `knife_edge` ogive can be used as a size-based ogive in an age-based model.

allvalues

$$f(x) = V_x$$

The `allvalues` ogive has estimable parameters $V_{low} V_{low+1} \dots V_{high}$. Here, you need to provide an ogive value for each age or size class. The `allvalues` ogive cannot be shifted and cannot be used as a size-based ogive in an age-based model.

allvalues_bounded

$$\begin{aligned} f(x) &= 0, & (x < L) \\ &= V_x, & (L \leq x \leq H) \\ &= V_H, & (x > H) \text{ (not } f(x) = 1!) \end{aligned}$$

The `allvalues_bounded` ogive has non-estimable parameters L and H . The estimable parameters are $V_L V_{L+1} \dots V_H$. Here, you need to provide an ogive value for each age or size class. The `allvalues_bounded` ogive cannot be shifted and cannot be used as a size-based ogive in an age-based model.

logistic

$$f(x) = 1 / \left[1 + 19^{(a_{50-x})/a_{t0.95}} \right]$$

The `logistic` ogive has estimable parameters a_{50} and $a_{t0.95}$. The `logistic` ogive takes values 0.5 at $x = a_{50}$ and 0.95 at $x = a_{50} + a_{t0.95}$. It can be shifted and can be used as a size-based ogive in an age-based model.

logistic_capped

$$f(x) = a_{\max} / \left[1 + 19^{(a_{50}-x)/a_{t095}} \right]$$

The logistic_capped ogive has estimable parameters a_{50} , a_{t095} , and a_{\max} . When $a_{\max} = 1$, it is identical to the logistic ogive, and otherwise follows a logistic form with values $0.5 \times a_{\max}$ at $x = a_{50}$ and $0.95 \times a_{\max}$ at $x = a_{50} + a_{t095}$. The logistic_capped ogive can be shifted and can be used as a size-based ogive in an age-based model.

logistic_bounded

$$\begin{aligned} f(x) &= 0, & (x < a_{50} - a_{t095}) \\ &= 1, & (x > a_{50} + a_{t095}) \\ &= 1 / \left[1 + 19^{(a_{50}-x)/a_{t095}} \right], & \text{otherwise} \end{aligned}$$

The logistic_bounded ogive is included to allow CASAL to replicate the ogives in previous NIWA software (pmod). It has estimable parameters a_{50} and a_{t095} . The logistic_bounded ogive can be shifted and can be used as a size-based ogive in an age-based model.

double_logistic

$$f(x) = \frac{\min \left(a_{\max} / \left[1 + 19^{(a_{50}-x)/a_{t095}} \right], a_{\max} / \left[1 + 19^{(x-(a_{50}+b_{50}))/b_{t095}} \right] \right)}{1 + 19^{\left(a_{50} - \left(\frac{a_{50}b_{t095} + a_{t095}(a_{50}+b_{50})}{a_{t095} + b_{t095}} \right) \right) / a_{t095}}$$

The double_logistic ogive has estimable parameters a_{50} , a_{t095} , b_{50} , b_{t095} , and a_{\max} . The ogive is evaluated as the minimum of a logistic increasing curve (defined by a_{50} and a_{t095}) and a logistic decreasing curve (defined by $a_{50}+b_{50}$ and b_{t095}). The maximum occurs at the intercept of the two logistics, and has value a_{\max} . The double_logistic ogive can be shifted and can be used as a size-based ogive in an age-based model.

logistic_product

$$f(x) = \frac{a_{\max} / \left(\left[1 + 19^{(a_{50}-x)/a_{t095}} \right] \left[1 + 19^{(x-(a_{50}+b_{50}))/b_{t095}} \right] \right)}{\max \left(1 / \left(\left[1 + 19^{(a_{50}-x)/a_{t095}} \right] \left[1 + 19^{(x-(a_{50}+b_{50}))/b_{t095}} \right] \right) \right)}$$

The logistic_product ogive has estimable parameters a_{50} , a_{t095} , b_{50} , b_{t095} , and a_{\max} . The ogive is the product of two logistic ogives, where the first is increasing (defined by a_{50} and a_{t095}) and the second decreasing (defined by $a_{50}+b_{50}$ and b_{t095}). The logistic_product has maximum value of a_{\max} , at the function maximum. This value is determined by approximation, and is defined as the maximum value of the 100 step sequence between $a_{50}-a_{95}$ and $a_{50}+b_{50}+b_{t095}$. This approximation should usually be accurate to within 0.1%. The logistic_product ogive can be shifted and can be used as a size-based ogive in an age-based model.

double_normal

$$\begin{aligned} f(x) &= 2^{-[(x-a_1)/s_L]^2}, & (x \leq a_1) \\ &= 2^{-[(x-a_1)/s_R]^2}, & (x > a_1) \end{aligned}$$

The double_normal ogive has estimable parameters a_1 , s_L , and s_R . It has values 1 at $x = a_1$, and 0.5 at $x = a_1 - s_L$ or $x = a_1 + s_R$. The double_normal ogive can be shifted and can be used as a size-based ogive in an age-based model.

double_normal_capped

$$\begin{aligned} f(x) &= a_{\max} \times 2^{-[(x-a_1)/s_L]^2}, & (x \leq a_1) \\ &= a_{\max} \times 2^{-[(x-a_1)/s_R]^2}, & (x > a_1) \end{aligned}$$

The double_normal_capped ogive has estimable parameters a_1 , s_L , s_R , and a_{\max} .

When $a_{\max} = 1$, it is identical to the double_normal ogive, and otherwise follows a double normal form with values a_{\max} at $x = a_1$, and $0.5 \times a_{\max}$ at $x = a_1 - s_L$ or $x = a_1 + s_R$. The double_normal_capped ogive can be shifted and can be used as a size-based ogive in an age-based model.

double_normal_plateau

$$\begin{aligned} f(x) &= a_{\max} \cdot 2^{-[(x-a_1)/s_L]^2}, & (x \leq a_1) \\ &= a_{\max}, & (a_1 < x \leq a_1 + a_2) \\ &= a_{\max} \cdot 2^{-[(x-(a_1+a_2))/s_R]^2}, & (x > a_1 + a_2) \end{aligned}$$

The double_normal_plateau ogive has estimable parameters a_1 , a_2 , s_L , s_R , and a_{\max} .

When $a_{\max}=1$ and $a_2=0$, it is identical to the double_normal ogive, and otherwise follows a double normal form with values a_{\max} at $a_1 < x \leq a_1+a_2$, and $0.5 \times a_{\max}$ at $x = a_1 - s_L$ or $x = a_1+a_2 + s_R$. The double_normal_plateau ogive can be shifted and can be used as a size-based ogive in an age-based model.

double_normal_coleraine

$$\begin{aligned} f(x) &= \exp\left[-(x-a_1)^2/\sigma_L^2\right], & (x \leq a_1) \\ &= \exp\left[-(x-a_1)^2/\sigma_R^2\right], & (x > a_1) \end{aligned}$$

The double_normal_coleraine ogive has estimable parameters a_1 , σ_L^2 , and σ_R^2 . This ogive can be shifted and can be used as a size-based ogive in an age-based model. It is designed to replicate the double normal ogive implemented in the Coleraine stock assessment model software (see Hilborn et al. 2001 for detail).

logistic_producing

$$\begin{aligned}
 f(x) &= 0, & (x < L) \\
 &= \lambda(L), & (x = L) \\
 &= (\lambda(x) - \lambda(x-1)) / (1 - \lambda(x-1)), & (L < x < H) \\
 &= 1, & (x \geq H) \\
 \text{where } \lambda(x) &= 1 / \left[1 + 19^{(a_{50}-x)/a_{1095}} \right]
 \end{aligned}$$

The logistic_producing ogive has the non-estimable parameters L and H , and has estimable parameters a_{50} and a_{1095} . The logistic_producing ogive cannot be shifted and cannot be used as a size-based ogive in an age-based model.

For maturation ogives, $f(x)$ represents the proportion maturing, *not* the proportion mature. If a logistic_producing maturation ogive is specified then (in the absence of other influences) the proportion mature will follow a logistic curve with parameters a_{50} , a_{1095} .

increasing

$$\begin{aligned}
 f(x) &= 0, & (x < L) \\
 &= f(x-1) + \pi_x (1 - f(x-1)), & (L \leq x \leq H) \\
 &= f(H), & (x > H) \text{ (note: not 1)}
 \end{aligned}$$

The increasing ogive has non-estimable parameters L and H . The estimable parameters are $\pi_L \pi_{L+1} \dots \pi_H$ (but if these are estimated, they should always be constrained to be between 0 and 1). The increasing ogive cannot be shifted and cannot be used as a size-based ogive in an age-based model. Note that the increasing ogive is similar to the allvalues_bounded ogive, but is constrained to be non-decreasing.

increasing_capped

$$\begin{aligned}
 f(x) &= 0, & (x < L) \\
 &= f(x-1) + \pi_x (C - f(x-1)), & (L \leq x \leq H) \\
 &= f(C), & (x \geq H)
 \end{aligned}$$

The increasing_capped ogive has non-estimable parameters L , H , and C . The estimable parameters are $\pi_L \pi_{L+1} \dots \pi_{H-1}$. Note that the maximum is π_{H-1} , not π_H as for the increasing ogive. As for the increasing ogive, if these are estimated then they should always be constrained to be between 0 and 1. The increasing_capped ogive cannot be shifted and cannot be used as a size-based ogive in an age-based model. Note that the increasing_capped ogive is similar to both the increasing ogive, but is constrained to be non-decreasing up to a specified cap.

4.8 Calculation of size-at-age (in an age-based model)

In an age-based model, fish size does not feature in the partition, but size-at-age is still an element of the model. See Section 4.4.5 for a discussion of fish growth in a size-based model.

Size-at-age is based on a growth curve which specifies the mean size at a given age. There are two alternative growth curves in CASAL:

1. von Bertalanffy, where size at age is defined as, $\bar{s}(age) = L_{inf} (1 - \exp(-k(age - t_0)))$
2. Schnute, where size at age is defined as,

$$\bar{s}(age) = \begin{cases} \left[y_1^b + (y_2^b - y_1^b) \frac{1 - \exp(-a(age - \tau_1))}{1 - \exp(-a(\tau_2 - \tau_1))} \right]^{1/b} & \text{if } a \neq 0, b \neq 0 \\ y_1 \exp \left[\ln(y_2/y_1) \frac{1 - \exp(-a(age - \tau_1))}{1 - \exp(-a(\tau_2 - \tau_1))} \right] & \text{if } a \neq 0, b = 0 \\ \left[y_1^b + (y_2^b - y_1^b) \frac{age - \tau_1}{\tau_2 - \tau_1} \right]^{1/b} & \text{if } a = 0, b \neq 0 \\ y_1 \exp \left[\ln(y_2/y_1) \frac{age - \tau_1}{\tau_2 - \tau_1} \right] & \text{if } a = 0, b = 0 \end{cases}$$

The von Bertalanffy curve is parameterised by L_{inf} , k , and t_0 ; the Schnute curve (Schnute 1981) by y_1 and y_2 , which are the mean sizes at reference ages τ_1 and τ_2 , and a and b (when $b = 1$, this reduces to the von Bertalanffy with $k = a$). All these parameters can depend on sex, stock, and/or growth-path. (But note, all the parameters should depend on the same thing. You can't supply L_{inf} by sex and k for both sexes combined.)

The model can incorporate changes in size-at-age during the year — i.e., growth between fish birthdays — by incrementing *age* as specified by the `annual_cycle.growth_props` parameter (see Section 4.3).

Optionally, you can give CASAL mean-size-at-age data which it can use instead of a growth curve. For one or more years, you provide the mean size at each age. For years for which you did not provide data, CASAL uses the average of all years for which you did provide data. You can provide mean size-at-age data for the projection period (see Section 6.3). In MCY/CAY simulations, CASAL uses the average of all years for which you did provide data. Be clear that CASAL does not calculate a best fit to the data; rather, it uses exactly the data provided. This option is only implemented for models where there is no growth between fish birthdays, i.e., `growth_props 0`.

Optionally, you can specify distributions of sizes at age, as well as the mean size at age. These size distributions are used to fit size frequency and age/size observations (Section 5.6), to calculate mean-weights-at-age (Section 4.9), and to convert size-based ogives to age-based (Section 4.6). Two distributional forms are implemented, normal and lognormal. In either case you need to give the c.v. of size-at-age, which can depend on sex, stock, and growth-path — but not on age.

If you use a growth curve, then as an optional feature, CASAL allows you to model annual growth variation. There is an annual growth variable for each year, indicating “how many

average years growth” fish achieve in that year. The mean size of a fish of age a in year y is then given by $f(e_{ay})$, where f is the growth curve, e_{ay} the ‘effective’ age, given by

$$e_{ay} = \sum_{y-a+I_{aged}}^{y-1+I_{aged}} r_i$$

where $I_{aged} = 1$ if fish have been aged yet in the year or 0 else, and the r_i are annual growth variables (in year y all fish grow as much as they would in r_y years of average growth). If within-year growth is included in the model, and proportion p of annual growth has occurred by a given time step, then the effective age at that time step is $e_{ay} + pr_{y+1_{aged}}$. Be careful to avoid off-by-one errors when you specify the r_i . They can cover any range of consecutive years. (Note that different annual growth variables for different stocks is not implemented.) If you use this feature, you can only use size-based ogives for selectivity, not for any other model feature such as migration rates.

So, to specify size at age, you need to tell CASAL the following:

1. Which growth curve is to be used — Schnute or von Bertalanffy. The parameters of the growth curve (which can depend on sex, stock, growth-path).
2. Alternatively, mean-size-at-age data for one or more years.
3. Whether size distributions around the mean are to be used. If so, with what distribution, and what c.v. (which again can depend on sex, stock, growth-path).
4. If annual growth variation is used, the growth variable for each year.

Be careful about the scale of the parameters (i.e., L_{inf} for von Bertalanffy growth) — this should be in units compatible with the size-weight relationship (Section 4.9). For example, if you provide your catch in tonnes and your size-weight relationship on a scale that converts a length in centimetres to a weight in tonnes, then your growth curve should be specified in centimetres.

4.9 Calculation of mean weight

In size-based models, you need to provide CASAL with size-weight parameters a and b , which can depend on sex and stock. It then calculates the mean weight for each size class as

$$\text{mean weight} = a \times \text{size}^b,$$

where size is approximated by (upper bound + lower bound of size class) / 2. If there is a plus group, you need to specify a nominal mean size for it using the `plus_group_size` parameter.

In age-based models, you again need to provide CASAL with size-weight parameters a and b , which can depend on sex and stock. If you don’t specify a distribution for sizes-at age (see Section 4.8), then the mean weight for a given partition element is calculated as,

$$\text{mean weight} = a \times (\text{mean size at age})^b$$

where the mean size at age can depend on the other partition characters, the time step, and the year (Section 4.8).

If you do specify a size distribution, then the mean weight at age is calculated over that distribution, using the following formula, which is exact for lognormal distributions, and a good approximation for a normal distribution (if the c.v. is not large),

$$\text{mean weight} = a \times (\text{mean size at age})^b \times (1 + cv^2)^{\frac{b(b-1)}{2}}$$

where *cv* is the c.v. of sizes-at-age for that element of the partition.

Be careful about the scale of *a* — this is easily specified incorrectly. If you provide your catch in tonnes, and your growth curve in centimetres, then *a* should be on the right scale to convert a length in centimetres to a weight in tonnes. Within the fin fisheries at NIWA, *a* is more often expressed on a scale to convert length in centimetres to weight in kilograms, and the user needs to divide this figure by 1000. Also note that the command `@size_weight` has the optional subcommand `verify_size_weight` that can be used to help check that the units specified are plausible.

4.10 Weightless model (running CASAL as a numbers only model)

You may wish to use a model which does not involve fish weight at all, but models the number of individuals instead (as can be the case in some shellfish models). For this type of model, abundance and catch data need to refer to numbers of fish, not biomass. If this is the case, set the `@weightless_model` switch to `true` (the default is `false`). The effect of this command is to assume a size-weight relationship of $w = 1$, i.e., each fish is assumed to have a nominal ‘weight’ of 1 tonne, irrespective of size. Note that CASAL will still label catches and abundances as if they were biomass (e.g., CASAL will report SSBs) but in each case, these values can be read directly as numbers of individuals.

Note that if `@weightless_model true`, then specifying any size-weight parameters in the `population.csl` file will generate an error.

4.11 Maturity, in models without maturity in the partition

When maturity is not a character in the partition, processes may still depend on maturity. You must then make the assumption that the proportion of mature fish in each element of the partition remains constant over time. You need to provide CASAL the proportion of mature fish in each size or age class, which can depend on sex (but not on stock). (Also, note that you are providing the proportion of mature fish, not the proportion of maturing fish as in Section 4.4.3).

Once you have done this, you can calculate SSB as a mature biomass (Section 4.3) and calculate fits to observations which relate to maturity (Section 5.6).

You can also have migrations which move only immature or only mature fish. If you migrate fish on the basis of maturity when maturity is not a partition character, be aware that the model does not know that the arriving fish are all mature, or all immature. So, if you migrate only mature fish into the spawning area, you need to tell CASAL that the SSB includes all fish in the relevant area — because it does not know that there are no immature fish present. See the subcommand `spawning_use_total_B` (Section 7.2).

You may not want to include maturity in the model in any shape or form, but CASAL still insists that you give it information on proportions mature. In this case, just set the proportion of mature fish to 1 for all age/size classes, as follows,

```
@maturity_props  
all constant 1
```

Then the SSB is simply the total biomass in the spawning area at the appropriate time.

5. THE ESTIMATION SECTION

5.1 Role of the estimation section

The tasks carried out by the estimation section are:

1. Get the *point estimate*, i.e., the least-squares fit, maximum likelihood estimate (MLE), or maximum posterior density estimate (MPD) (see Section 5.3).
2. *Profile* selected parameters, i.e., find, for each of a series of values of a parameter, allowing all other free parameters to vary, the minimum value of the objective function (Section 5.4). This is called either a likelihood or posterior profile (profiling is not appropriate for weighted least-squares estimation).
3. For Bayesian estimation only, generate an *MCMC* sample from the posterior distribution (Section 5.5).
4. For maximum likelihood or Bayesian estimation, calculate the approximate *covariance* matrix of the parameters as the inverse of the minimiser's approximation to the Hessian, and the corresponding correlation matrix (Section 5.3).

A key decision is between least-squares, likelihood, and Bayesian estimation, which define the *objective function* as a weighted sum of squares, negative log-likelihood, and negative log-posterior respectively (Section 0).

5.2 Specifying the free parameters

You need to tell CASAL which of the estimable parameters are to be freed by using `@estimate` commands (see Section 8). An `@estimate` command-block looks like this,

```
@estimate
parameter initialization.B0
lower_bound 1000
upper_bound 100000
prior uniform
```

See Section 2.4 for instructions on how to generate the parameter name. You have to specify at least one free parameter. You still provide values for the free parameters as normal, these are used as the starting values for the minimiser (unless you provide alternative starting values using `casal -i`, see Section 2.1).

All parameters are estimated within bounds. For each free parameter (scalar, vector, or ogive), you need to specify the bounds, and, in a Bayesian analysis, the prior (Section 5.7.5). Note that the bounds and prior on an ogive refer to the ogive free parameters, not the actual values of the ogive.

You need to estimate all the estimable parameters of an ogive if you estimate any, but you can fix some of them if you want by setting the lower and upper bound equal. Similarly, if you want to estimate only some elements of a vector, fix the others by setting the bounds equal.

Relativity constants q are a bit of a special case, because no starting value is provided in the *nuisance* method (see Section 5.7.2) whereas all other free parameters always need starting values. But you still need to provide an `@estimate` block for each q , containing the bounds on the q , using the following format,

```
@estimate
parameter q[label].g
lower_bound 1e-6
upper_bound 1e-2
```

where `label` is the label of the q in the observations blocks. If you're still uncertain how to do this then look at the example in Section 13.

If you want to estimate two (scalar, vector, or ogive) parameters and to constrain them to be the same, you need to use the `same` subcommand. This might arise if, for example, you wanted to use the same migration ogive for two different migrations involving different stocks. Only use one `@estimate` block, for one or other of the two parameters. Put in the same command, using the name of the other parameter as the argument. For example,

```
@estimate
parameter growth[1].g
same growth[2].g
...
```

means that the `g` parameters for the first two growth episodes are both estimated, but constrained to be equal. (Don't put in a second `@estimate` block for `growth[2].g` with same `growth[1].g`.)

5.3 Point estimation

Point estimation is invoked with `casal -e`, and also used in several other tasks. Mathematically, it is an attempt to find a minimum of the objective function. CASAL approaches this optimisation problem using a quasi-Newton minimiser built into `Betadiff`, which is a slightly modified implementation of the main algorithm of Dennis Jr. & Schnabel (1996).

The minimiser has three kinds of (non-error) exit status:

1. *Successful convergence* (suggests you have found a local minimum, at least).
2. *Failure to converge* (you have not reached a local minimum, though you may deem yourself to be 'close enough' at your own risk).
3. *Convergence unclear* (the minimiser has clunked to a halt. You may have found a local minimum, although you should check by restarting the minimiser at the final values of the free parameters).

You can choose the maximum number of quasi-Newton iterations and objective function evaluations allotted to the minimiser. If it exceeds either limit, it exits with a convergence failure. We urge you to use large numbers of evaluations and iterations (at least the defaults of 300 and 1000) unless you successfully reach convergence with less. You can also specify the starting point of the minimiser using `casal -i`.

We want to stress that this is a local optimisation algorithm trying to solve a global optimisation problem. What this means is that, even if you get a 'successful convergence' message, your solution may be only a local minimum, not a global one. To diagnose this problem, try doing multiple runs from different starting points and comparing the results, or

(probably better) doing profiles of one or more key parameters and seeing if any of the profiled estimates is actually better than the original point estimate. Otherwise you may have reached a false (local) minimum.

The approximate covariance matrix of the free parameters can be calculated as the inverse of the minimiser's approximation to the Hessian, and the corresponding correlation matrix is also calculated. These results are printed if you use the `print.covariance` parameter (along with the eigenvalues of the Hessian, see Section 9.1). Be aware (i) that the Hessian approximation develops over many minimiser steps, so if the minimiser has only run for a small number of iterations the covariance matrix can be a very poor approximation, (ii) in any case the inverse Hessian is not a good approximation to the covariance matrix of the free parameters, and should not be used, for example, to construct confidence intervals. Also note that if a free parameter has equal lower and upper bounds, it will have entries of '0' in the covariance matrix and 'NaN' in the correlation matrix.

Multi-phase estimation is allowed, and is implemented in a manner similar to that by AD-Model Builder (Otter Research Limited 2000). In this case, some free parameters are initially held fixed, and a minimisation is carried out. Next, some or all of the fixed parameters are freed, and another minimisation is carried out, etc. Sensible starting values should be used for the fixed parameters. Apparently this can be quicker and/or more effective than estimating all the parameters in a single minimisation. (The main idea is that the 'key' parameters should be freed first and the 'nuisance' parameters last, although there is little known about the actual performance improvements that may be expected from this approach.) If this feature is used, then each parameter should be allocated a 'phase'; the default phase is 1. The phase 1 parameters are freed first, then the phase 2 parameters, etc. You can specify that a different maximum number of iterations and/or evaluations is to be used for the 'intermediate' phases, i.e., all but the last. It would probably be advisable to use rather less effort for the intermediate phases than for the final phase.

An option (`casal -E`) is provided in which the point estimate is calculated using finite difference gradients instead of automatic differentiation. This was implemented for three reasons:

1. You can use finite differences to check your results if you suspect the automatic differentiation is misbehaving.
2. If you become aware of problems in the automatic differentiation section of Betadiff and can't fix them, you will need to switch to finite differences.
3. You may find that finite differences is faster than automatic differentiation, with comparable accuracy, for problems with small numbers of free parameters. (Although, for large problems, ~100 free parameters, it can be many times slower.)

This option is implemented only for simple point estimation, and not for profiling or for the initial point estimate in MCMC runs.

5.4 Likelihood or posterior profiles

If profiles are requested (`casal -p`), CASAL will first calculate a point estimate, then, for each scalar parameter to be profiled, fix its value at a sequence of n evenly spaced numbers between specified bounds l and u , and calculate a point estimate at each value. By default $n = 10$, and $(l, u) = (\text{lower bound on parameter} + (\text{range} / (2n)), \text{upper bound on parameter} - (\text{range} / (2n)))$. Each minimisation starts at the final parameter values from an adjacent value

of the parameter being profiled. The program reports the objective function for each parameter value, and all the parameter estimates. The initial point estimate is also inserted into the profile (note that this serves as a check that none of the other points along the profile have a better objective function value than the initial ‘minimum’).

You specify which parameters are to be profiled, and optionally n , l , and u values for each. Only scalar parameters can be profiled.

You can also supply the initial point estimate using `casal -i`, so that CASAL doesn’t need to do the first minimisation. Be aware that you are supplying the point estimate, not the minimiser starting point to get to the point estimate (as in other situations where `casal -i` is used).

If you have specified multi-phase estimation (see Section 5.3), it is only used for the initial point estimate. Subsequent minimisations are done single-phase, as they should start reasonably close to the endpoint and so shouldn’t need multiple phases.

If you are doing a Bayesian analysis and want likelihood profiles rather than posterior profiles, then either switch to the likelihood objective function (using `@estimator likelihood`) for the duration or make all the priors uninformed.

If you get an implausible profile, it may be a result of not using enough iterations in the minimiser. In this case, increase `max_iters` and/or `max_evals` and retry.

5.5 Bayesian estimation

CASAL can:

1. Use a Monte Carlo Markov Chain to generate a sample from the posterior distribution of the free parameters (`casal -m`); and output the sampled values to a file, (optionally only every n th set of values).
2. If the run is interrupted, recover the results from the file and continue the run from where it left off, appending the results to the file (`casal -a`).
3. Combine the results of one or more chains into a single posterior sample by removing samples from the ‘burn-in’ periods and concatenating the results; allow the user to reduce the size of the resulting sample by sub-sampling; and optionally, apply prior reweighting in the sub-sampling process, i.e., apply probability weights to generate a sample from a posterior based on a different prior (`casal -C`). The sub-sampling may be either systematic (every n th point) or randomly (with replacement). The former is recommended (to minimise autocorrelation) except with prior reweighting, when the latter must be used.
4. For a posterior sample, calculate the values of various output quantities at each sample point and export these so that they can be plotted and/or summarised using an external package (use `casal -v`).

Two major steps are best done by an external package, as CASAL has no post-processing capabilities. CASAL cannot:

1. Produce MCMC convergence diagnostics (use a package such as BOA, see <http://www.public-health.uiowa.edu/boa>).

2. Plot/summarize the posterior distributions of the output quantities (use a general-purpose statistical or spreadsheet package such as S/S-Plus/R or Microsoft Excel).

Bayesian methodology and MCMC are both large and complex topics, and we do not describe either properly here. See Gelman et al. (1995) and Gilks et al. (1998) for details of both Bayesian analysis and MCMC methods. In addition, see Punt & Hilborn (2001) for an introduction to quantitative fish stock assessment using Bayesian methods.

This section only briefly describes the MCMC algorithms used in CASAL. See Section 2.2 for a better description of the sequence of CASAL commands used in a full Bayesian analysis.

CASAL uses a straightforward implementation of the Metropolis algorithm (Gelman et al. 1995, Gilks et al. 1998). The Metropolis algorithm attempts to draw a sample from a Bayesian posterior distribution, and calculates the posterior density π , scaled by an unknown constant. The algorithm generates a ‘chain’ or sequence of values. Typically the beginning of the chain is discarded and every N th element of the remainder is taken as the posterior sample. The chain is produced by taking an initial point x_0 and repeatedly applying the following rule, where x_i is the current point:

1. Draw a candidate step s from a proposal distribution J , which should be symmetric i.e., $J(-s) = J(s)$.
2. Calculate $r = \min(\pi(x_i + s) / \pi(x_i), 1)$.
3. Let $x_{i+1} = x_i + s$ with probability r , or x_i with probability $1 - r$.

An initial point estimate is produced before the chain starts, which is done so as to calculate the approximate covariance matrix of the free parameters (as the inverse Hessian), and may also be used as the starting point of the chain.

The user can specify the starting point of the point estimate minimiser using `casal -i`. Don’t start it too close to the actual estimate (either by using `casal -i`, or by changing the parameter values in `population.csl`) as it takes a few iterations to form a reasonable approximation to the Hessian.

There are three options for the starting point of the Markov Chain:

1. Start from the point estimate.
2. Start from a random point near the point estimate (the point is generated from a multivariate normal distribution, centred on the point estimate, with covariance equal to the inverse Hessian times a user-specified constant). This is done to prevent the chain from getting ‘stuck’ at the point estimate.)
3. Start from a point specified by the user with `casal -i`.

The chain moves in natural space, i.e., no transformations are applied to the free parameters. The default proposal distribution is a multivariate normal centred on the current point, with covariance matrix equal to a matrix based on the approximate covariance produced by the minimiser, times some `stepsize` factor. The following steps define the initial covariance matrix of the proposal distribution:

1. The covariance matrix is taken as the inverse of the approximate Hessian from the quasi-Newton minimiser.
2. This covariance matrix is modified so as to decrease all correlations greater than `max_cor` down to `max_cor`, and similarly to increase all correlations less than `-max_cor` up to `-max_cor` (the `max_cor` parameter defaults to 0.8). This should help to avoid getting 'stuck' in a lower-dimensional subspace.
3. The diagonal of the covariance matrix is then modified, so that any nonzero element which is less than 0.01% of the difference between its lower and upper bound is rounded up to that value. This allows a free parameter to move in the MCMC even if its variance is very small according to the inverse Hessian.
4. The `stepsize` (a scalar factor applied to the covariance matrix to improve the acceptance probability) is chosen by the user. The default is $2.4d^{-0.5}$ where d is the number of free parameters, as recommended by Gelman et al. (1995), though experience has shown that this is often too high, leading to a very low acceptance rate.

The proposal distribution can also change adaptively during the chain, using two different mechanisms. Both are offered as means of improving the convergence properties of the chain. It is important to note that any adaptive behaviour must finish before the end of the burn-in period, i.e., the proposal distribution must be finalised before the kept portion of the chain starts (CASAL enforces this). The adaptive mechanisms are as follows:

1. You can request that the step size change adaptively at one or more sample numbers. At each adaptation, the step size is doubled if the acceptance rate since the last adaptation is more than 0.5, or halved if the acceptance rate is less than 0.2. (See Gelman et al. 1995 for justification, but the basic idea is that the convergence of the chain can be expected to be poor if the acceptance probability of each step is very small or if the expected size of the step is very small.) The new step size is recorded in the `objectives` file.
2. You can request that the entire covariance matrix change adaptively at one or more sample numbers. At each adaptation, it is replaced with a matrix based on the sample covariance of an earlier section of the chain. The theory here is that the covariance of a portion of chain could potentially be a better estimate of the covariance of the posterior distribution than the inverse Hessian.

The procedure used to choose the sample of points is as follows. First, all points on the chain so far are taken. All points in an initial user-specified period are discarded. The assumption is that the chain will have started moving during this period - if this is incorrect and the chain has still not moved by the end of this period, it is a fatal error and CASAL stops. The remaining set of points must contain at least some user-specified number of transitions — if this is incorrect and the chain has not moved this often, it is again a fatal error. If this test is passed, the set of points is systematically subsampled down to 1000 points (it must be at least this long to start with).

The variance-covariance matrix of this subsample of chain is calculated. As above, correlations greater than `max_cor` are reduced to `max_cor`, correlations less than `-max_cor` are increased to `-max_cor`, and very small nonzero variances are increased. The result is the new variance-covariance matrix of the proposal distribution.

The step size parameter is now on a completely different scale, and must also be reset. It is set to a user-specified value (which may or may not be the same as the initial step size). We recommend that some of the step size adaptations are set to occur after this, so that the step size can be readjusted to an appropriate value which gives good acceptance probabilities with the new matrix.

All modified versions of the covariance matrix are printed to the standard output, but only the initial covariance matrix (inverse Hessian) is saved to the `objectives` file. (As a consequence, a lapsed chain cannot be continued using `-a` if adaptive covariance is used.) The number of covariance modifications by each iteration is recorded as a column on the `objectives` file.

The probability of acceptance for each jump is 0 if it would move out of the bounds, or 1 if it improves the posterior, or $(\text{new posterior}/\text{old posterior})$ otherwise.

You can specify how often the position of the chain is recorded using the `keep` parameter. For example, with `keep 10`, only every 10th sample is written to file.

You have the option to specify that some of the free parameters are fixed during MCMC. If the chain starts at the point estimate or at a random location, these fixed parameters are set to their values at the point estimate. If you specify the start of the chain using `-i`, these fixed parameters are set to the values in the file.

A multivariate t distribution is available as an alternative to the multivariate normal proposal distribution. If you request multivariate t proposals, you may want to change the degrees of freedom from the default of 4. As the degrees of freedom decrease, the t distribution becomes more heavy tailed. This may lead to better convergence properties.

Having produced one or more Markov chains and looked at the diagnostics, you should reload all the chain output files into CASAL and use them to generate a single posterior sample (using `-C`). At this stage, the first `burn_in` iterations for each chain are discarded (so, with `keep 10`, `burn_in 1000`, the first 1000 recorded samples are discarded for each chain). Unless a very large value of `keep` was originally chosen, it will be necessary to further reduce the size of the posterior sample (possibly down to several hundred) such that it can be analysed in a reasonable amount of time. This is done by sub-sampling. You specify the size of the sub-sample to be produced (or else no sub-sampling is done). You have the option to generate a systematic sub-sample (i.e., every n th point is kept) or a random sub-sample (the former is recommended except with prior reweighting, when the latter must be used).

Given a posterior (sub)sample, CASAL can calculate a list of output quantities for each sample point (see Section 6.2). These quantities can be dumped into a file (using `casal -v`) and read into an external software package where the posterior distributions can be plotted and/or summarised.

The posterior sample can also be used for projections (Section 6.3) and stochastic yield calculations (Section 6.5). The advantage of this is that the parameter uncertainty, as expressed in your posterior distribution, can be included into the risk and yield estimates.

It is possible to investigate the results that you would have got if you had used a different prior. This is called prior reweighting and is done by calculating the ratio of the new prior to the original prior for each point in the posterior sample, then using these ratios as probability weights when generating a random (not systematic) sub-sample with `casal -C`. Prior reweighting is applicable only if the new prior is zero in every part of the parameter space for which the original prior was zero. Also, it is likely to be numerically unstable unless the new

prior is very small in every part of the parameter space for which the original prior was very small.

5.6 Observations

The objective function is based on the goodness-of-fit of the model to your observations. In the current release of CASAL, most observations are different kinds of time series, i.e., data which were recorded for one or more years, in the same format each year. Examples of time series data types include relative abundance indices, commercial catch length frequencies, survey numbers-at-age, etc.,.

Generally, time series must relate to a specified time step, and a specified area if the model is spatial, and one or more years in which they were recorded. These are the exceptions; (a) catch-at data (see below) can be based on more than one fishery and hence can cover multiple areas or time steps, and (b) age-at-maturation data (see below) are not associated with a year, time step, or area.

5.6.1 Types of observations

Each time series of observations belongs to one of the following types:

1. *Abundance*: including survey biomass indices and CPUE. Can be absolute abundance or relative abundance. Can be expressed as biomass or numbers of fish.
2. *Catch-at*: including commercial catch proportions-at-age and proportions-at-size. Can be split by sex.
3. *Numbers-at*: including survey numbers-at-age and numbers-at-size. Can be absolute numbers or relative numbers. Can be split by sex.
4. *Proportions-at*: including survey proportions-at-age and proportions-at-size. Can be split by sex.
5. *Proportions mature*: i.e., data on the proportion of fish, by age or size class, which are mature. If you are using age frequency observations in an age-based model, you need to say which age classes are included and whether the last age class is a plus group. For example, your partition might include ages from 1 to 20+ but you might have observations only for 2, 3, 4, 5, and 6+ aged fish. The same applies to size frequency observations in a size-based model.
6. *Proportions migrating*: i.e., data on the proportion of fish, by age or size class, which went on a particular migration. More specifically, for each age/size class, the number of migrating fish (in the source area of the migration, including all stocks) divided by the total number of fish (in the source area of the migration, including all stocks). Can be provided for either sex, or both combined.
7. *Age size*: i.e., observations of the ages and sizes of individual fish — primarily used to fit size-at-age parameters in age-based models. See below.
8. *Age-at-maturation*: a specialised observation type, used for modelling orange roughy and potentially other fish with similar characteristics. The age at which a mature orange roughy became mature can be estimated by examining the otolith after

capture. The age-at-maturation data type allows CASAL to use this information for estimation of maturation parameters. See below.

For each time series, as well as the above, you need to provide CASAL the following information:

1. A label. (This should be unique, neither “Bpre” or “Bpost”, and should not contain a full stop.)
2. The years in which they were observed.
3. For *catch-at* observations, the fishery or fisheries they cover.
4. For all observations except *catch-at* and *age-at-maturation*:
 - The area in which they were observed (in a multi-area model).
 - After what proportion of the mortality in the time step they occurred. (This is a useful option as it allows you to insert an observation partway through a time step, which can sometimes avoid splitting the time step into two.)
 - The name of the selectivity ogive which should be applied, if any (trawl survey data should use the selectivity of the research vessel; CPUE data could arguably use the selectivity of the commercial fleet).
5. If you are using age frequency observations in an age-based model, you need to say which age classes are included and whether the last age class is a plus group. For example, you might have observations of 2, 3, 4, 5, and 6+ aged fish. The same applies to size frequency observations in a size-based model.
6. If you are using a set of size frequency observations in an age-based model, you need to provide the size classes. (In a size-based model, the size bins used in the observations must be consecutive groupings of the bins defined in the population section of the model.)
7. For *catch-at* and *proportions-at* observations, you need to specify whether the observations and fits for each year should sum to 1.

If you say that proportions must sum to 1, the CASAL calculates the expected values using a denominator constructed from the number of fish in the age or size classes included in the range specified in the observations.

Otherwise, if you do not force the proportions to sum to 1, then the expected values are calculated using a denominator of the total number of fish in all age or size classes in the partition. This approach was intended for backwards compatibility with previous least-squares estimation software, and may not always work in likelihood analysis.

The setting `sum_to_one true` is the default and is recommended. If the observations don't sum to 1 (and this is determined for each year independently), a warning is printed and the observations (for that year) are rescaled to sum to 1. (Note the test for observations summing to 1 is implemented with a tolerance of ± 0.01 .)

8. For relative observations, you need to provide the label of the relativity constant q . Several time series can share the same q (see Section 5.7.2). You can also provide a curvature parameter b for non-linear relationships (see Section 5.7.2).

9. For *at-age* data, whether ageing error should be applied (Section 5.7.7). It may be the case that you have defined an ageing error but don't want to apply it to one or more of your time series, perhaps because the ageing error is meant to be included in the likelihood.
10. For weighted least-squares estimation, you need to provide the weight u for the time series and the (single) c.v. c for each year (Section 5.7.1).
11. For likelihood or Bayesian analysis, you need to provide the error distribution and its parameters (the error distributions are listed in Section 5.7.2). For variability parameters (c.v.s, standard deviations, and effective sample sizes N), there can either be one value for all years, or one value per year, or (for *at-age* or *at-size* data), one value per age/size class per year. Note that c.v.s are expressed as a proportion not a percentage, for example, if you put a c.v. of 40 then you probably meant 0.4. See also specifying the process error (Section 5.7.3).
12. And of course you have to provide the observation values. Abundance values are straightforward — one number per year. *At-age* or *at-size* data are a bit more complicated. There is one row of numbers for each year, one column per age/size class. If the observations are sexed, then there are male columns followed by female columns, rather than separate male and female tables. *Age-size* and *age-at-maturation* data are input as several rows of data, one per variable: age, size and potentially sex for *age-size*, and age at capture, age at maturation and potentially sex for *age-at-maturation*

5.6.2 Age/size observations

Age/size data are observations of the ages and sizes of individual fish. They are primarily used to fit *size-at-age* parameters in age-based models.

Age-size data cannot currently be used in growth path models, and can never be used in *size-based* models. An error message will be issued if *age-size* observations are used without valid *size-at-age* parameters (i.e. no `@size_at_age_dist` or with `@size_at_age.cv = 0`). *Age-size* observations can only be used in Bayesian or likelihood analysis (not weighted least squares).

The data include a list of ages, a list of sizes, and (in a sexed model) a list of sexes, plus information on when, where, and how the observations were collected. So, the i^{th} elements of the lists contain the age, size, and sex of the i^{th} fish observed.

There are several possible sampling regimes, i.e. assumptions about how the observed fish were sampled from the general population of fish available at that time and place. The options are:

- `random`: fish were a simple random sample from the available population
- `random_at_sex`: fish were a simple random sample within each sex
- `random_at_size`: fish were a simple random sample within each size class
- `random_at_sex_and_size`: fish were a simple random sample within each size class of each sex
- `random_at_age`: fish were a simple random sample within each age class
- `random_at_sex_and_age`: fish were a simple random sample within each age class of each sex

We believe the `at_age` and `at_sex_and_age` options are quite unlikely to be true, yet they are widely applied in fisheries (for example, in the Coleraine (Hilborn et al. 2001) stock modelling software). Probably the `at_size` and `at_sex_and_size` options are most likely to hold for most NIWA finfish programmes.

In age-size data, there should be no observations for which the age is outside the age range defined for the partition (this will generate a fatal error message), nor any non-integer ages. Observations with ages below the minimum age in the partition should be removed. Observations where an age exceeds the maximum age in the partition could either be included in the plus group (i.e., with the observed age changed to that of the plus group) or removed, depending on the circumstances. If there is no plus group in the partition they should be removed. For `random_at_age` or `random_at_sex_and_age` samples they may be either included or removed in the plus group. For all other sample types they could be included in the plus group.

In addition, the user can specify a selectivity ogive which was applied in the sampling process, perhaps due to the sampling gear that was used, or the areal availability of fish at that place or time. The ogive can be age- or size-based: the choice has direct bearing on the likelihood of the observations.

Under some sampling regimes, a size-based selectivity has no effect on the likelihood and hence should not be used (since it adds computational time). This occurs when the character on which the selectivity acts was not randomly chosen in the sample: for instance, if 10 fish of each sex were chosen from each size class, then a size-based selectivity will have had no effect (except perhaps to make it easier/harder to find the 10 fish!). In general, a size-based selectivity has no effect under the `random_at_sex_and_size` sampling method, and it has no effect under the `random_at_size` method unless the selectivity is specified by sex. If you attempt to use a size-based selectivity in this situation, CASAL will issue a warning and will not apply the ogive. In cases where the size-based selectivity does have an effect, CASAL will issue a warning that the selectivity adds to the computational time, and will apply it as requested.

Similarly, under some sampling regimes, an age-based selectivity has no effect on the likelihood, and should not be used. This applies under the `random_at_sex_and_age` sampling method, and under the `random_at_age` method unless the selectivity is specified by age. If you attempt to use an age-based selectivity in this situation, CASAL will issue a warning and not apply the ogive.

The user must additionally specify the year, time step, and proportion of mortality when the observations were collected, the area in which they were collected, and whether ageing error is to be applied (by default it is applied, if ageing error parameters are supplied in the input parameter files).

The user doesn't need to specify an error distribution or its parameters: instead, the choice of size-at-age distribution sampling regime, selectivity ogive, and ageing error determines the likelihood equation. The appropriate likelihood for a single observation, (a,l,s) , depends on the nature of the sample.

With a random sample covering a single stock,

$$L = P(a,l,s) = \left[\sum_{a'} N_{a's} M_{a'a} f_{a's}(l) \right] / \left[\sum_{a's'} N_{a's'} \right],$$

where $N_{a's}$ is the number of fish of true age a' and sex s (in the specified stock and area, and after the specified selectivity, if any, is applied), $M_{a'a}$ is the probability that a fish of true age a' is observed as age a , and $f_{a's}(l)$ is the probability density function describing the distribution of sizes for a given (true) age a' and sex s . When there is no ageing error the numerator of the above equation simplifies to $N_{as}f_{as}(l)$ and the denominator to $\sum_{a's'} N_{a's'}$.

If there is a mixture of stocks in the area sampled, then the N and f terms are stock-dependent and the numerator and denominator are each summed over stocks

For all the other sample types the likelihood is a conditional probability, and is calculated as a fraction whose numerator is the same as for $P(a,l,s)$ and whose denominator is given in Table 2.

Table 2: Age/size likelihoods for the different sample types.

Sample	Conditional probability	Denominator	
		With ageing error	Without
random_at_sex	$L = P(a,l s)$	$\sum_{a'} N_{a's}$	same
random_at_age	$L = P(l,s a)$	$\sum_{a's'} N_{a's'} M_{a'a}$	$\sum_{s'} N_{as'}$
random_at_size	$L = P(a,s l)$	$\sum_{a's'} N_{a's'} f_{a's'}(l)$	same
random_at_sex_and_age	$L = P(l a,s)$	$\sum_{a'} N_{a's} M_{a'a}$	N_{as}
random_at_sex_and_size	$L = P(a l,s)$	$\sum_{a'} N_{a's} f_{a's}(l)$	same

If the user specifies a selectivity for an age/size observation, this is easy to deal with if the selectivity is age-based. If $N'_{a's}$ is the number at true age a' and sex s before the selectivity is applied then $N_{a's} = N'_{a's} S_s(a')$, where S is the selectivity function.

It's more complicated with a size-based selectivity because we have also to distinguish between the distribution of size at age before [$f'_{a's}(l)$] and after [$f_{a's}(l)$] the selectivity is applied (note: it is the former which is defined by the model parameters). The appropriate equations are

$$f_{a's}(l) = S_s(l) f'_{a's}(l) / \int_{l'} S_s(l') f'_{a's}(l') dl',$$

and

$$N_{a's} = N'_{a's} \int_{l'} S_s(l') f'_{a's}(l') dl'$$

The integrals in these equations are calculated by discrete approximation (using 5 points), in the same way as size-based ogives are converted to age-based ogives.

Fits and residuals are not displayed for age-size data, instead CASAL shows the contribution to the objective function of each individual age-size pair.

5.6.3 Age-at-maturation observations

This type of observation makes sense only for species like orange roughy, in which it is possible to tell (by examining a mature fish) the age it was when it matured. It is available

only in an age-based model with maturity in the partition and no more than one maturation episode per year. The use of this observation requires the important assumption that mortality is independent of maturity status (i.e., the mortality experienced in a given year by fish of a given age is independent of whether the fish is mature or not). This observation differs from most others in that it is not associated with any particular area, year, time step, or ogive.

Suppose we have a sample of n fish and let A_{sj} be the age at which the j th fish was sampled and A_{mj} its age at maturation (the age which it matured). If the fish is not mature we will signal this by setting $A_{mj}=0$. It is not necessary that this be a fully random sample, but it is necessary that, amongst fish of the same age, the probability of selection does not depend on the age at maturation. That is, the sample must be random, conditional on the age at sampling. The likelihood that CASAL associates with these observations is the conditional likelihood $P(A_{mj}|A_{sj})$. (Note that if the sample were fully random it would be sensible to calculate the joint likelihood, $P(A_{sj},A_{mj})=P(A_{mj}|A_{sj})P(A_{sj})$. Users can, in effect, achieve this by also providing the A_{sj} as a `proportions_at` observation.)

We further assume that mortality is independent of maturity status. This will often not be strictly true. For example, it is false if migration to a fishing ground is dependent on maturity status. The assumption is necessary because without it, the calculation of the likelihood, though still possible in principle, would require fundamental structural changes to CASAL. CASAL does not check this assumption so it is up to the user to decide whether it is warranted.

For orange roughy, the age at maturation cannot be determined until, say, k years after maturation (because the maturation mark in the otolith is not clearly apparent until additional otolith material has been deposited — see Francis & Horn 1997). This means that there would be no observations with $A_s - A_m < k$. The user may specify a value k .

If there is no ageing error the likelihood is as follows:

$$P(A_{mj} = a | A_{sj} = A) = \begin{cases} \frac{O_a U_{a-1}}{U_{A-1} + \sum_{a'=a_{\min}}^{A-1} O_{a'} U_{a'-1}} & a_{\min} \leq a \leq A - k \\ 0 & A - k < a < A \\ \frac{U_{A-1} + \sum_{a'=A-k}^{A-1} O_{a'} U_{a'-1}}{U_{A-1} + \sum_{a'=a_{\min}}^{A-1} O_{a'} U_{a'-1}} & a = 0 \end{cases}$$

where O_a is the proportion maturing at age a , a_{\min} is the minimum age in the partition, $U_a = \prod_{a'=a_{\min}}^a (1 - O_{a'})$, and $U_{a_{\min}-1} = 1$. Note that this is independent of the year in which the observation was made (which is not true if mortality depends on maturity status).

Ageing error can have a substantial effect on this type of observation (it makes the maturation ogive appear to be less steep than it really is) so it is important to be able to allow for it. However, it is not straightforward to allow for error in both A_{mj} and A_{sj} (in fact, if the sample is not fully random we're not sure it's possible). Thus, we allow only for error in the A_{mj} ; we do not allow for error either in the A_{sj} or in the detection of maturity. With ageing error, the adjusted likelihood is given by

$$P'(A_{mj} = a | A_{sj} = A) = \begin{cases} \frac{\sum_{b=a_{\min}}^{A-k} M_{ba} P_{bA}}{\sum_{a'=a_{\min}}^{A-k} \sum_{b=a_{\min}}^{A-k} M_{ba'} P_{bA}} & a_{\min} \leq a \leq A-k \\ 0 & A-k < a < A \\ \frac{U_{A-1} + \sum_{a'=A-k}^{A-1} O_{a'} U_{a'-1}}{U_{A-1} + \sum_{a'=a_{\min}}^{A-1} O_{a'} U_{a'-1}} & a = 0 \end{cases}$$

where $P_{bA} = P(A_{mj} = b | A_{sj} = A)$ and M_{ba} is the probability that a fish with true maturity age b is observed as having matured at age a (this is the ageing-error misclassification matrix — see Section 5.7.7).

If the calculated likelihood for an observation is equal to zero then CASAL replaces this with 10^{-6} to avoid errors from taking the logarithm of zero.

Fits and residuals are not displayed for age-at-maturation data, instead CASAL shows the contribution to the objective function of each individual age-at-maturation observation.

For these observations, the user must specify;

- age at maturation for each fish sampled
- age at capture for each fish sampled
- optionally, the sex of each fish sampled
- k , the number of years after which maturity can be detected
- whether ageing error is used.

Note that users can moderately improve the performance of CASAL in cases where there are observations on fish with an identical set of maturation ages, sampled ages, and sexes in sexed observations, by placing such observations immediately adjacent to each other in the `age_at_maturation` subcommands within the `estimation.csl` file. In this situation (i.e., when an observation on a fish is identical to the fish immediately preceding) CASAL simply copies the contribution of that observation to the likelihood rather than attempting to recalculate it.

5.7 The objective function

In *maximum likelihood* estimation, the objective function is a negative log-likelihood,

$$\text{Objective}(\mathbf{p}) = -\sum_i \log[L(\mathbf{p} | O_i)]$$

where \mathbf{p} is a vector of the free parameters, L the likelihood function and O_i the i th observation.

In *Bayesian* estimation, the objective function is a negative log-posterior,

$$\text{Objective}(\mathbf{p}) = -\sum_i \log[L(\mathbf{p} | O_i)] - \log[\pi(\mathbf{p})]$$

where π is the joint prior density of the parameters \mathbf{p} .

In *weighted least-squares* estimation, the objective function is a weighted sum of squares on the log-scale,

$$\text{Objective } (\mathbf{p}) = \sum_i w_i \left(\log \left(\max(P_i, k_p) \right) - \log \left(\max(O_i, k_o) \right) \right)^2$$

where w_i is the weight (see Section 5.7.1), and P_i the predicted (fitted) value, of O_i , and k_p and k_o are small robustifying constants.

Under any estimation method, penalties can be added to the objective function (see Section 5.7.6). You will usually want to use penalties to ensure that the exploitation rate constraints on your fisheries are not breached (otherwise there is nothing to prevent the model from having abundances so low that the recorded catches could not have been taken). A penalty to force the YCS to average to 1 (i.e., to have mean 1) may also be necessary.

5.7.1 Weighted least-squares

In weighted least-squares estimation, each observation has a weight. The default option is for all weights to be equal (i.e., ordinary least-squares). We also implement the “Cordue” weighting scheme (see, for example, Cordue 2000). In this framework, you should specify a weight u_i , and a c.v. c_{iy} for each year y , of each time series i . Then the weights on the k individual observations from time series i , for year y of n_s years, are

$$w_{iy} = \frac{u_i n_i}{k c_{iy}^2 \sum_j \frac{1}{c_{ij}^2}}$$

The original “Cordue” formulation was further modified by ‘sources’. Several time series collected in similar ways (e.g., from the same set of trawl surveys) could share the same source. See Cordue (2000) for details. This feature is not implemented in CASAL.

5.7.2 Likelihoods

CASAL has five different kinds of likelihoods:

1. Those used for proportions data, where the proportions should sum to 1 over the columns of the partition. This includes commercial and survey proportions-at-size and proportions-at-age data, but not proportions mature or proportions migrating.
2. Those used for proportions data, where the proportion can be between 0 and 1 in each cell, and need not sum to 1 over the columns of the partition. This includes proportions mature and proportions migrating.
3. Those used for absolute index data. These likelihoods are used for absolute abundance and can also be used for proportions mature and proportions migrating.
4. Those used for relative index data, including relative abundance and relative survey numbers-at-age.
5. Specialized likelihoods used for age-size data or age-at-maturation data — see Section 5.6 for details.

See also Section 5.7.3 for detail about process error.

Likelihoods for proportions data (with proportions summing to 1 across columns of the partition)

These likelihoods are used for commercial and survey proportions-at-size and proportions-at-age data, but not proportions mature or proportions migrating. They apply to data which are distributed across columns of the partition, typically summing to 1 across columns.”

Let \mathbf{O} be the observations for a single year in a proportions time series, expressed as a vector of n proportions summing to 1; let \mathbf{E} be the corresponding fitted values; let N be the “effective sample size” parameter. Then you can use the following likelihoods, which are expressed on the objective-function scale of $-\log(L)$:

1. Multinomial

$$-\log(L) = N \sum_i (O_i + r) \log \left[\frac{O_i + r}{E_i + r} \right]$$

where r is a non-negative robustifying constant (a nonzero value of r is recommended if only to prevent division by zero errors).

2. Fournier

$$-\log(L) = 0.5 \sum_i \log(E'_i) - \sum_i \log \left[\exp \left(\frac{-(O_i - E_i)^2}{2E'_i/N'} \right) + 0.01 \right]$$

where $E'_i = (1 - E_i)E_i + 0.1/n$ and $N' = \min(N, 1000)$.

This is a robustified multivariate normal (it would be the usual multivariate normal with $\sigma^2 = (1 - E_i)E_i/N$ if the $0.1/n$ and 0.01 terms are omitted and if $N' = N$). See Fournier et al. (1990).

3. Coleraine. As per Fournier above, but replace E'_i with $O'_i = (1 - O_i)O_i + 0.1/n$. A recent reference is Starr et al. (1999).

4. (Robustified) lognormal

$$-\log(L) = \sum_{i=1}^n \left(\log(\sigma_i) - \log \left(\exp \left(-0.5 \left(\frac{\log(O_i/E_i)}{\sigma_i} + 0.5\sigma_i \right)^2 \right) + r \right) \right)$$

where $\sigma_i = \sqrt{\log(1 + c_i^2)}$, the c_i 's are c.v.s by age/size class, and r is a robustifying constant.

The robustification term r is intended to reduce the influence of outliers, in the same way as the robustified Fournier likelihoods above. We recommend $r = 0.01$, though smaller values (0.001, 0.0001, ...) could be used for a lesser robustifying effect.

Likelihoods for proportions data (with proportions between 0 and 1 in each cell)

These likelihoods are used for proportions mature and proportions migrating data. They apply to data which can be between 0 and 1 in each cell, not necessarily summing to 1 across columns.

So far there is just one likelihood in this category, the binomial .

Let \mathbf{O} be the observations for a single year in a time series, expressed as a vector of n proportions between 0 and 1; let \mathbf{E} be the corresponding fitted values; let N be the “effective sample size” parameter. A single effective sample size may be specified for all observations in a given year, or different sample sizes may be specified for each observation in the year. Then the likelihood, which is expressed on the objective-function scale of $-\log(L)$, is:

$$-\log(L) = \sum_{i=1}^n 0.5 \left(\frac{(E_i + r)(1 - E_i + r)}{N_i} \right) + \sum_{i=1}^n 0.5 \left((O_i - E_i) / \sqrt{\frac{(E_i + r)(1 - E_i + r)}{N_i}} \right)^2$$

where r is a non-negative robustifying constant (a nonzero value of r is recommended if only to prevent division by zero errors).

This likelihood corresponds to the assumption that, for each observation in each year, a simple random sample of size N_i was taken from the partition and that the observation value was the proportion of fish in the sample who migrated/matured.

Likelihoods for absolute index data

Let \mathbf{O} be the observations for a single year in an time series of absolute indices, expressed as a vector of n elements (with $n = 1$ for abundance indices, $n > 1$ for proportions mature or proportions migrating); let \mathbf{E} be the corresponding fitted values; express the variability of each observation O_i in terms of its c.v. c_i (or in one case, its standard deviation s_i). Then you can use the following likelihoods, which are expressed on the objective-function scale of $-\log(L)$:

1. Normal

$$-\log(L) = \sum_{i=1}^n \left(\log(c_i E_i) + 0.5 \left(\frac{O_i - E_i}{c_i E_i} \right)^2 \right)$$

This reflects the distributional assumption that O_i has the normal distribution, with mean E_i and c.v. c_i .

2. Normal parameterised by standard deviation rather than c.v.

$$-\log(L) = \sum_{i=1}^n \left(\log(s_i) + 0.5 \left(\frac{O_i - E_i}{s_i} \right)^2 \right)$$

This reflects the distributional assumption that O_i has the normal distribution, with mean E_i and standard deviation s_i .

3. Lognormal:

$$-\log(L) = \sum_{i=1}^n \left(\log(\sigma_i) + 0.5 \left(\frac{\log(O_i/E_i)}{\sigma_i} + 0.5\sigma_i \right)^2 \right)$$

$$\text{where } \sigma_i = \sqrt{\log(1 + c_i^2)}.$$

This reflects the distributional assumptions that O_i has the lognormal distribution, that the mean of O_i is E_i and the c.v. of O_i is c_i .

4. Normal-log

$$-\log(L) = \sum_{i=1}^n \left(\log(\sigma_i) + 0.5 \left(\frac{\log(O_i/E_i)}{\sigma_i} \right)^2 \right)$$

$$\text{where } \sigma_i = \sqrt{\log(1 + c_i^2)}.$$

This reflects the distributional assumption that $\log(O_i)$ has the normal distribution, that the mean of $\log(O_i)$ is $\log(E_i)$ and the c.v. of O_i is c_i .

We make the distinction between lognormal and normal-log because they represent subtly different assumptions. With the lognormal, O has mean E and hence the mean of $\log(O)$ is less than $\log(E)$: whereas with the normal-log, $\log(O)$ has mean $\log(E)$ and hence the mean of O is more than E .

Relativity constants q ; likelihoods for relative index data

The log-likelihoods of relative observations depend on the error distribution and the way in which q 's are treated in the model. There are two approaches to modelling q 's:

1. The q 's can be treated as 'nuisance' parameters. For each set of values of the free parameters, the model uses the values of the q 's which minimise the objective function. These optimal q 's are calculated algebraically (see Section 5.7.4). If one of the q 's falls outside the bounds specified by the user, it is set equal to the closest bound.

This approach reduces the size of the parameter vector and hence should improve the performance of the estimation method. It is the default in CASAL. However, it is not correct when calculating a sample from the posterior in a Bayesian analysis (except asymptotically, see Walters & Ludwig 1994) and we offer the following alternative;

2. The q 's can be treated as ordinary free parameters.

For both options, we need to evaluate the contribution of O to the negative log-likelihood for a given value of q . Let O be the observations for a single year in an time series of relative indices, expressed as a vector of n elements (with $n = 1$ for relative abundance indices, $n > 1$ for relative numbers-at-age or at-size). Let E be the corresponding fitted values and q the relativity constant for the time series. It is possible for two or more time series collected in similar ways to use the same q . Each observation O_i varies about qE_i — express the variability

of O_i in terms of its c.v. c_i (or in one case, its standard deviation s_i). Here are the likelihoods, which are expressed on the objective-function scale of $-\log(L)$:

1. Normal

$$-\log(L) = \sum_{i=1}^n \left(\log(c_i q E_i) + 0.5 \left(\frac{O_i - q E_i}{c_i q E_i} \right)^2 \right)$$

This reflects the distributional assumption that O_i has the normal distribution, with mean $q E_i$ and c.v. c_i .

2. Lognormal

$$-\log(L) = \sum_{i=1}^n \left(\log(\sigma_i) + 0.5 \left(\frac{\log(O_i/q E_i)}{\sigma_i} + 0.5 \sigma_i \right)^2 \right)$$

where $\sigma_i = \sqrt{\log(1 + c_i^2)}$.

This reflects the distributional assumptions that O_i has the lognormal distribution, that the mean of O_i is $q E_i$ and the c.v. of O_i is c_i .

3. Normal-log

$$-\log(L) = \sum_{i=1}^n \left(\log(\sigma_i) + 0.5 \left(\frac{\log(O_i/q E_i)}{\sigma_i} \right)^2 \right)$$

where $\sigma_i = \sqrt{\log(1 + c_i^2)}$.

This reflects the distributional assumption that $\log(O_i)$ has the normal distribution, that the mean of $\log(O_i)$ is $\log(q E_i)$ and the c.v. of O_i is c_i .

4. Robustified lognormal

$$-\log(L) = \sum_{i=1}^n \left(\log(\sigma_i) - \log \left(\exp \left(-0.5 \left(\frac{\log(O_i/q E_i)}{\sigma_i} + 0.5 \sigma_i \right)^2 \right) + r \right) \right)$$

where $\sigma_i = \sqrt{\log(1 + c_i^2)}$ and r is a robustifying constant.

This modification to the lognormal is intended to reduce the influence of outliers, and is analogous to the robustified normal distributions of Fournier et al. (1990). We recommend $r = 0.01$, though smaller values (0.001, 0.0001, ...) could be used for a lesser robustifying effect. We believe that this likelihood may be most appropriate for relative numbers-at-age or at-size but not for relative abundance indices (its effects are equivalent to dropping observations with large normalised residuals, which may be undesirable for abundance data).

Optionally, a curvature parameter can be used, in which case the error distribution of O_i is centred on $q(E_i/\max(E))^{1/b}$ rather than qE_i . This is intended for modelling hyper-depletion or hyper-stability in CPUE (see Harley et al. 2001). Note that the interpretation of q changes as the expected values are rescaled. We also note that this option has not yet been fully tested in CASAL.

5.7.3 Process error

In a likelihood-based or Bayesian analysis, you can specify a ‘process error’ for each set of observations. This has the effect of increasing the error in the data (by increasing c.v.s or standard deviations, or decreasing effective sample sizes), and hence of decreasing the weight given to the data in the fitting process.

For data where the likelihood is parameterised by the c.v., you can specify the process error for a given set of observations as a c.v., in which case all the c.v.s c_i are changed to

$$c'_i = \sqrt{c_i^2 + c_{process_error}^2} .$$

Similarly, if the likelihood is parameterised by the standard deviation,

$$\sigma'_i = \sqrt{\sigma_i^2 + \sigma_{process_error}^2} ,$$

and by the effective sample size,

$$N'_i = \frac{1}{1/N_i + 1/N_{process_error}} .$$

In all three cases, the process error has more effect on small errors than on large ones. Be clear that a large $N_{process_error}$ means a small process error.

CASAL allows you to estimate process error, though whether you should is another matter. If you want to make several sets of observations share the same process error, use the same subcommand in the `estimate` block (see Section 5.2).

5.7.4 Calculating nuisance q 's

This section describes the equations used to calculate nuisance q 's (see Section 5.7.2). From the user's point of view, the essence is that you can use nuisance q 's in the following situations:

1. With least-squares.
2. With maximum likelihood.
3. With Bayesian estimation, providing that your prior on the q is one of the following:
 - Uniform
 - Uniform-log
 - Lognormal with observations distributed lognormal, robustified lognormal, or normal-log.

For weighted least-squares, nuisance q 's are calculated as per Cordue (2000),

$$\hat{q} = \exp \left(\frac{1}{\sum w_i} \sum_i \left[w_i \log \left(\frac{\max(O_i, k_o)}{\max(P_i, k_p)} \right) \right] \right)$$

where the summation is over the n observations O_i sharing the same q .

For ordinary least-squares, this reduces to

$$\hat{q} = \exp \left(\frac{1}{n} \sum_i \log \left(\frac{\max(O_i, k_o)}{\max(P_i, k_p)} \right) \right).$$

The equations used for calculating nuisance qs in maximum likelihood or Bayesian analysis are indexed in Table 3.

Table 3: Equations used to calculate nuisance q 's. (* = no analytic solution found.)

Distribution of observations	Maximum likelihood	Bayesian with specified prior on q			
		Uniform	Uniform-log	Normal	Lognormal
Normal	(1)	(1)	(4)	*	*
Lognormal	(2)	(2)	(5)	*	(6)
Normal-log	(3)	(3)	(7)	*	(8)

Note that q 's are calculated for robustified lognormal likelihoods as if they were ordinary lognormal likelihoods.

The equations and their derivations follow. Let $\sigma_i = \sqrt{\log(1 + c_i^2)}$ throughout, and let n be the number of observations in the time series. The case of multiple time series sharing the same q is addressed at the end of the subsection.

First, consider maximum likelihood estimation. When the (O_i) are assumed to be normally distributed,

$$-\log(L) = \sum_i \log(c_i q_i E_i) + 0.5 \sum_i \left(\frac{O_i - q E_i}{c_i q E_i} \right)^2$$

The value of q which minimises the objective function is found by solving $\partial/\partial q(-\log(L)) = 0$.

$$\frac{\partial}{\partial q}(-\log(L)) = \frac{n}{q} + \frac{1}{q^2} \sum_i \frac{O_i}{c_i^2 E_i} - \frac{1}{q^3} \sum_i \left(\frac{O_i}{c_i E_i} \right)^2$$

hence

$$\hat{q} = \frac{-S_1 + \sqrt{S_1^2 + 4nS_2}}{2n} \tag{1}$$

where $S_1 = \sum_i (O_i/c_i^2 E_i)$ and $S_2 = \sum_i (O_i/c_i E_i)^2$.

When the (O_i) are assumed to be lognormally distributed,

$$\begin{aligned}
 -\log(L) &= \sum_i \log(\sigma_i) + 0.5 \sum_i \left(\frac{\log(O_i) - \log(qE_i) + 0.5\sigma_i^2}{\sigma_i} \right)^2, \\
 \frac{\partial}{\partial q}(-\log(L)) &= \frac{-1}{q} \sum_i \left(\frac{\log(O_i/E_i) - \log(q) + 0.5\sigma_i^2}{\sigma_i^2} \right), \\
 \hat{q} &= \exp\left(\frac{0.5n + S_3}{S_4} \right)
 \end{aligned} \tag{2}$$

where $S_3 = \sum_i (\log(O_i/E_i)/\sigma_i^2)$ and $S_4 = \sum_i (1/\sigma_i^2)$.

When the (O_i) are assumed to be distributed normal-log, the equations are similar,

$$\begin{aligned}
 -\log(L) &= \sum_i \log(\sigma_i) + 0.5 \sum_i \left(\frac{\log(O_i) - \log(qE_i)}{\sigma_i} \right)^2, \\
 \frac{\partial}{\partial q}(-\log(L)) &= \frac{-1}{q} \sum_i \left(\frac{\log(O_i/E_i) - \log(q)}{\sigma_i^2} \right), \\
 \hat{q} &= \exp\left(\frac{S_3}{S_4} \right)
 \end{aligned} \tag{3}$$

Next consider Bayesian estimation, where we must also specify a prior for q .

The effects of the prior on the equations are to replace likelihood L by posterior P throughout, to add $-\log(\pi(q))$ to the equation for $-\log(P)$ and $\partial/\partial q(-\log(\pi(q)))$ to the equation for $\partial/\partial q(-\log(P))$.

This last term is 0 for a uniform prior on q , $1/q$ for a log-uniform prior, $\frac{q - \mu_q}{(\mu_q c_q)^2}$ for a normal

prior, and $\frac{1}{q} \left(1.5 + \frac{\log(q) - \log(\mu_q)}{\sigma_q^2} \right)$ for a lognormal prior,

where μ_q and c_q are the mean and c.v. of the prior on q and $\sigma_q = \sqrt{\log(1 + c_q^2)}$. Clearly, if the prior is uniform, the equation for \hat{q} is the same as for maximum likelihood estimation.

When the (O_i) are assumed to be normally distributed and the prior is log-uniform, equation (1) becomes

$$\hat{q} = \frac{-S_1 + \sqrt{S_1^2 + 4(n+1)S_2}}{2(n+1)} \quad (4)$$

but we cannot solve for \hat{q} with either a normal or lognormal prior.

When the O_i are assumed to be lognormally distributed and the prior is log-uniform, equation (2) becomes

$$\hat{q} = \exp\left(\frac{0.5n - 1 + S_3}{S_4}\right) \quad (5)$$

and if the prior is lognormal,

$$\hat{q} = \exp\left(\frac{0.5n - 1.5 + \log(\mu_q)/\sigma_q^2 + S_3}{S_4 + 1/\sigma_q^2}\right), \quad (6)$$

but we cannot solve for \hat{q} with a normal prior.

When the (O_i) are assumed to be distributed normal-log and the prior is log-uniform,

$$\hat{q} = \exp\left(\frac{S_3 - 1}{S_4}\right) \quad (7)$$

and if the prior is lognormal,

$$\hat{q} = \exp\left(\frac{-1.5 + \log(\mu_q)/\sigma_q^2 + S_3}{S_4 + 1/\sigma_q^2}\right) \quad (8)$$

but again we cannot solve for \hat{q} with a normal prior.

The above equations have been written for a single time series (O_i) . Suppose now that there are m time series, all with the same q , and all with the same error distribution. This has little effect on the above equations. All we need to do is to extend the summations in S_1 , S_2 , S_3 , and S_4 over all observations in the m time series, and let $n = \sum_j n_j$.

5.7.5 Priors

In a Bayesian analysis, you need to give a prior for every free parameter. There are no defaults.

Note that when some of these priors are parameterised in terms of mean, c.v., and standard deviation, these refer to the parameters of the distribution before bounds are applied. The moments of the prior after the bounds are applied may differ.

For a single scalar parameter p , you can choose between the following priors (expressed in terms of their contribution to the objective function):

1. Uniform,

$$-\log(\pi(p)) = 0.$$

2. Uniform-log (i.e., $\log(p) \sim \text{uniform}$),

$$-\log(\pi(p)) = \log(p).$$

3. Normal with mean μ and c.v. c ,

$$-\log(\pi(p)) = 0.5 \left(\frac{p - \mu}{c\mu} \right)^2.$$

4. Normal with mean μ and standard deviation σ ,

$$-\log(\pi(p)) = 0.5 \left(\frac{p - \mu}{\sigma} \right)^2.$$

5. Lognormal with mean μ and c.v. c . $s = \sqrt{\log(1 + c^2)}$, is the standard deviation of $\log(p)$.

$$-\log(\pi(p)) = \log(p) + 0.5 \left(\frac{\log(p/\mu)}{s} + \frac{s}{2} \right)^2.$$

6. Normal-log with $\log(p)$ having mean m and standard deviation s ,

$$-\log(\pi(p)) = \log(p) + 0.5 \left(\frac{\log(p) - m}{s} \right)^2.$$

7. Beta with mean μ and standard deviation σ , and range parameters A and B .

$$-\log(\pi(p)) = (1 - m) \log(p - A) + (1 - n) \log(B - p),$$

$$\text{where } \nu = \frac{\mu - A}{B - A}, \text{ and } \tau = \frac{(\mu - A)(B - \mu)}{\sigma^2} - 1$$

and then $m = \tau\nu$, and $n = \tau(1 - \nu)$. Note that the beta prior is undefined when $\tau \leq 0$.

Vectors of parameters can be independently (but not necessarily identically) distributed according to any of the above forms, in which case the joint negative-log-prior for the vector is the sum of the negative-log-priors of the components. Values of each parameter need to be specified for each element of the vector.

In addition, for a vector \mathbf{p} of n identically distributed parameters (for example, YCS) the following priors are allowed:

1. Multivariate normal from a stationary AR(1) process with parameters

$\mu = E(p_i)$, $\sigma = \text{sqrt}(\text{Var}(p_i))$, and $\rho = \text{Cor}(p_i, p_{i+1})$,

$$-\log(\pi(\mathbf{p})) = \frac{(p_1 - \mu)^2}{2\sigma^2} + \frac{\sum_{i=2}^n (p_i - \rho p_{i-1} - \mu(1 - \rho))^2}{2\sigma^2(1 - \rho^2)} + n \log(\sigma) + 0.5(n-1) \log(1 - \rho^2)$$

In other words, there are $(n - 1)$ i.i.d. normal variates z_i with mean μ and variance σ^2 , such that $p_i = \rho p_{i-1} + (\sqrt{1 - \rho^2}) z_i$.

If $\rho = 0$, then the p_i 's are i.i.d. normal.

2. Multivariate normal-log, where $\log(\mathbf{p})$ forms a stationary AR(1) process as per 1. above, with parameters

$m = E(\log(p_i))$, $s = \text{sqrt}(\text{Var}(\log(p_i)))$, and $r = \text{Cor}(\log(p_i), \log(p_{i+1}))$,

$$-\log(\pi(\mathbf{p})) = \frac{(\log(p_1) - m)^2}{2s^2} + \frac{\sum_{i=2}^n (\log(p_i) - r \log(p_{i-1}) - m(1 - r))^2}{2s^2(1 - r^2)} + n \log(s) + 0.5(n-1) \log(1 - r^2) + \sum_i p_i$$

3. Multivariate normal-log with mean 1, where $E(p_i) = 1$ and $\log(\mathbf{p})$ forms a stationary AR(1) process as for the multivariate normal above, with parameters

$s = \text{sqrt}(\text{Var}(\log(p_i)))$ and $r = \text{Cor}(\log(p_i), \log(p_{i+1}))$,

$$-\log(\pi(\mathbf{p})) = \frac{(\log(p_1) + 0.5s^2)^2}{2s^2} + \frac{\sum_{i=2}^n (\log(p_i) - r \log(p_{i-1}) + 0.5s^2(1 - r))^2}{2s^2(1 - r^2)} + n \log(s) + 0.5(n-1) \log(1 - r^2) + \sum_i p_i$$

(i.e., $m = -0.5s^2$)

5.7.6 Penalties

Penalties can be added to any objective function. You will usually want to use a *catch limit penalty* for each fishery to ensure that the exploitation rate constraints on your fisheries are not breached (otherwise there is nothing to prevent the model from having abundances so low that the recorded catches could not have been taken). A *vector average penalty* to force YCS to average to 1 is also very common.

For each penalty, you need to specify a *multiplier*. The objective function is increased by this multiplier times the penalty as described below. In some cases you will want to make the multiplier quite large, so as to prohibit some model behaviour.

So far, the penalties implemented in CASAL are:

1. *Ogive smoothing penalty:*

Applied to an `allvalues` or `allvalues_bounded` ogive parameter (Section 4.6). Sum of squares of r th differences. This encourages the ogive to be like a polynomial of degree $(r - 1)$. For compatibility with previous NIWA software, you can choose to exclude indices outside a given set of bounds (these indices are ignored during differencing).

2. *Catch limit penalty:*

Sum of squares of (actual catch less specified catch), optionally on a log scale, for a single fishery. These are intended to avoid parameter values that cause the specified fishing pressure limits to be exceeded. The penalty is only applied if some fishing pressure limit has been exceeded (since inaccuracy in the iterative solution for F in the Baranov equation leads to actual catches slightly less than specified, and you don't want to penalise that).

3. *Vector average penalty:*

Applied to a vector parameter. Square of $(\text{mean}(\text{vector}) - k)$, or of $(\text{mean}(\log(\text{vector})) - l)$, or of $(\log(\text{mean}(\text{vector}) / m))$. Encourages the vector to average arithmetically to k or m , or geometrically to $\exp(l)$. Typically used for YCS with $k = 1$ or $m = 1$ or $l = 0$, to encourage the YCS to centre on 1.

4. *Vector smoothing penalty:*

Applied to a vector parameter. Sum of squares of r th differences. This encourages the vector to be like a polynomial of degree $(r - 1)$. Note a range of the vector to be "smoothed" can be specified (and if not, the smoother is applied to the entire vector), but this must be specified by an index of the vector and must be between 1 and the length of the vector, inclusive.

5. *Element difference penalty:*

Applied to two vector parameters. Square of $(\text{vector}_1[i] - \text{vector}_2[i])$. Encourages the i th elements of the two vectors to be equal.

6. *YCS difference penalty:*

Applied to the YCS of two different stocks. Squared difference between the YCS values for a given year in the two stocks. Used to encourage the two stocks to have the same YCS for that year. If the Haist YCS parameterisation is used, then the penalty applies to the YCS (as one would expect) and not the Y 's.

7. *Similar q s penalty:*

Applied to two relativity constants q (Section 5.7.2). Square of $(\log(q_i) - \log(q_j))$. This is intended to encourage q_i and q_j to be similar, because they belong to observations collected in similar ways.

8. *Ogive comparison penalty:*

Applied to two ogive parameters (Section 4.6). Sum of squares of $\max(\text{ogive}_1 - \text{ogive}_2, 0)$. Encourages ogive_1 to be at or below ogive_2 . Typically ogive_1 is a selectivity for males, ogive_2 is a selectivity for females. This is intended to encourage female selectivities to be greater than those of males at the same age/size. For compatibility with previous NIWA software, you can choose to exclude indices outside a given set of bounds (these indices are dropped off before comparing). Note that this penalty may not be applied to size-based ogives in age-based models.

9. *Ogive difference penalty:*

Applied to two ogive parameters (Section 4.6). Square of $(\text{ogive}_1 - \text{ogive}_2)$ for a single size or age class. This is intended to encourage the two ogives to take the same value for that class. Note that this penalty may not be applied to size-based ogives in age-based models.

5.7.7 Ageing error

In age-based models, we allow ageing error in at-age observations to be modelled explicitly. After E (expected) values are calculated for at-age observations, misclassification rates are applied to them, which has the effect of ‘smearing’ the age frequencies. The resulting ‘smeared’ age frequencies are used in calculating the objective function.

Ageing error is optional, and if it is used, it may be omitted for any individual time series. However, CASAL does not yet implement changes in ageing error over time, or different ageing error regimes for different time series.

The ageing error models implemented in CASAL are as follows:

1. *Off by one:*

Proportion p_1 of fish of each age a are misclassified as age $a - 1$ and proportion p_2 are misclassified as age $a + 1$. Fish of age $a < k$ are not misclassified. If there is no plus group in the population model, then proportion p_2 fish of the oldest age class will ‘fall off the edge’ and disappear.

2. *Normal:*

Fish of age a are classified as ages which are normally distributed with mean a and constant c.v. c . As above, if there is no plus group in the population model, some fish of the older age classes may disappear. If c is high enough, some fish of the younger age classes may ‘fall off the other edge’ too.

3. *Misclassification matrix:*

A complete misclassification matrix M is provided, such that M_{ij} is the probability that a fish of age i will be classified as age j . Rows need not sum to 1, but a warning will be issued if they don’t.

Note that the expected values (fits) reported by CASAL for an individual time series with ageing error, have had the ageing error applied.

5.8 Residuals

CASAL can generate three kinds of residuals, (1) the usual residuals (i.e., observed less fitted), (2) *Pearson* residuals, and (3) *normalised* residuals. There are defined in CASAL as,

Let O be an observation and F the corresponding fit ($= qE$ for relative observations), then:

1. *Residuals* are defined as $(O - F)$.
2. *Pearson residuals* attempt to express the residual relative to the variability of the observation, and are defined as $(O - F) / \text{std.dev.}(O)$, where $\text{std.dev.}(O)$ is calculated as,
 - $F \times cv$ for normal, lognormal, robustified lognormal, and normal-log error distributions.
 - s for normal-by-standard deviation error distributions.
 - $\sqrt{\frac{F(1-F)}{N}}$ for multinomial or binomial likelihoods (regardless of the robustifying constant).
 - $\sqrt{\frac{F'}{N'}}$ for Fournier likelihoods (where $F'_i = (1 - F_i)F_i + 0.1/n$ and $N' = \min(N, 1000)$, on the basis that they would be equivalent to a multivariate normal with this standard deviation if the final (+0.01) term was omitted.)
 - $\sqrt{\frac{O'}{N'}}$ for Coleraine error likelihoods (similarly).
3. *Normalised residuals* attempt to express the residual on a standard normal scale, and are defined as,
 - Equal to the Pearson residuals for normal error distributions.
 - $(\log(O/F) + 0.5\sigma^2) / \sigma$ for lognormal (including robustified lognormal) error distributions, where $\sigma = \sqrt{\log(1 + cv^2)}$.
 - $(\log(O/F) / \sigma)$ for normal-log error distributions, again with $\sigma = \sqrt{\log(1 + cv^2)}$.
 - And are otherwise undefined.

5.9 Generate simulated observations

CASAL can generate simulated observations from a parameter fit, i.e., generate simulated observations which are randomly distributed (according to the error assumptions defined for the observations) around fits calculated from one or more sets of the 'true' parameter values. This is a form of parametric bootstrap.

One use of this feature is to investigate the uncertainty in CASAL parameter estimates, using a bootstrapping approach:

1. Get one or more sets of free parameters, either using an assumed set of values, a point estimate or a sample from the posterior distribution
2. Use CASAL to generate many sets of simulated observations, on the assumption that the free parameter estimates are the true values
3. For each set of simulated observations, generate a simulated estimate of the free parameters (replacing the real observations with the randomised observations)
4. The variability in the simulated estimates is a bootstrap estimate of the uncertainty in the estimation process.

This approach allows the user to assess estimator performance in varying conditions. For example, the simulated estimates could be carried out using a simplified estimation procedure (perhaps fixing some previously free parameters), and the effect of this simplification on estimator performance could then be investigated.

The way in which the above process could be undertaken might be:

1. Estimate the free parameters using `-e`, `-E` or `-m`. Generate a file of free parameter values using the usual format (described in Section 2.3)
2. Run CASAL in simulator mode (`-s`), supplying the file of free parameter values with the command line switch `-i`, using the same input parameter files (`population.csl` and `estimation.csl`). The results are files of simulated observations.
3. CASAL creates one file for each set of simulated observations. The total number of files is equal to the number of free parameter sets supplied, multiplied by the number of simulations per parameter set (the later is the number that you supply as an argument to `-s` on the CASAL command line).
4. Each file contains a complete set of observations, using the standard CASAL syntax.
5. All commands and subcommands will be unchanged from the original estimation parameter file, except for the original observation values which will be replaced with randomised values. (although the subcommands will appear in alphabetical order, and comments and white space will have been removed.)
6. A comment is appended to the top of each file, listing the free parameter values that were used to generate it.
7. Outside of CASAL, attach (append or prepend) a ‘stub’ estimation parameter file to each file of simulated observations. This stub file should contain all the estimation parameters *other than the observations*, and include a list of parameters to be estimated, etc., using the usual CASAL syntax.
8. Using each (stub + simulated observations) file along with standard input parameter files (`population.csl` and `output.csl`), carry out an estimate of the free parameters using `-e`, `-E`, or even `-m` or `-Y`.

The remainder of this section describes the method used to produce the simulated observations for a single set of ‘true’ free parameter values.

First, the model is run using the true free parameter values, and a set of fits is produced for each set of observations. If a set of observations uses ageing error, then ageing error is applied to the fits as per normal. If there are relative observations, then the catchability coefficient q is applied to the fits as per normal.

Second, each set of observations is randomised, based on

- the fitted values
- the type of likelihood specified (note, likelihoods must be provided — CASAL cannot be used as a simulator if estimation is by least squares)
- the variability parameters (c.v. N , or σ). Variability is increased by the process error associated with that time series, if any (see Section 5.7.3). If the process error parameter is a free parameter, then the ‘true value’ of the process error is used.

Age-size and age-at-maturation observations cannot be simulated. You need to remove these observations before using CASAL as a simulator.

The following text describes the process of generating simulated observations for each type of likelihood.

1. Normal likelihood parameterised by c.v.: Let E_{yi} be the fitted value for observation i in year y and c_{yi} be the corresponding c.v. (adjusted by process error if applicable). Each simulated observation value S_{yi} is generated as an independent normal deviate with mean E_{yi} and standard deviation $E_{yi} c_{yi}$.
2. Normal likelihood parameterised by standard deviation: Let E_{yi} be the fitted value for observation i in year y and s_{yi} be the corresponding standard deviation (adjusted by process error if applicable). Each simulated observation value S_{yi} is generated as an independent normal deviate with mean E_{yi} and standard deviation s_{yi} .
3. Log-normal likelihood: Let E_{yi} be the fitted value for observation i in year y and c_{yi} be the corresponding c.v. (adjusted by process error if applicable). Each simulated observation value S_{yi} is generated as an independent lognormal deviate with mean and standard deviation (on the natural scale, not the log-scale) of E_{yi} and $E_{yi} c_{yi}$ respectively. The robustification parameter r is ignored.
4. Normal-log likelihood: Let E_{yi} be the fitted value for observation i in year y and c_{yi} be the corresponding c.v. (adjusted by process error if applicable). Each simulated observation value S_{yi} is generated as an independent lognormal deviate, such that the mean of $\log(S_{yi})$ is $\log(E_{yi})$ and the c.v. of S_{yi} is c_{yi} .
5. Multinomial likelihood: This is only allowed if the same N value is used for all observations of the same time series in the same year. Let E_{yi} be the fitted value for observation i in year y , for i between 1 and n , and let N_y be the equivalent sample size for that year (rounded up to the next whole number, and adjusted by process error if applicable). Any robustification is ignored. The following process is carried out for each year y :
 - a. A sample of N data values from 1 to n is generated using the multinomial distribution, using sample probabilities proportional to the values of E_{yi} .
 - b. Each simulated observation value S_{yi} is calculated as the proportion of the N sampled values equalling i .

- c. The simulated observation values S_{yi} are then rescaled so that their sum is equal to the sum of E_{yi} . (The sum of the fitted values for the year may not be equal to 1 if `sum_to_one = False` and the age/size range of the observations does not cover all columns in the partition.)
6. Coleraine or Fournier likelihood: These are not ‘proper’ likelihoods in the technical sense, and we do not use them as distributions for generating simulated values. Instead, as they are analogous to the multinomial likelihood, we apply the above procedure for the multinomial, using the supplied value of the N parameter.
7. Binomial likelihood: Let E_{yi} be the fitted value for observation i in year y , for i between 1 and n , and N_{yi} the corresponding equivalent sample size (rounded up to the next whole number, and adjusted by process error if applicable). Any robustification is ignored. The following process is carried out for each observation i in each year y :
 - a. A sample of N_{yi} independent binary variates is generated, equalling 1 with probability E_{yi} .
 - b. The simulated observation value S_{yi} is calculated as the sum of these binary variates divided by N_{yi} .

6. THE OUTPUT SECTION

This section contains three main topics.

1. Section 6.1 describes the printouts from CASAL, which are dumped to standard output (and can be redirected to a file, and imported into S/S-Plus/R using the functions defined in Section 12).
2. Section 6.3 describes projections in CASAL.
3. Sections 6.4 and 6.5 describe yield calculations in CASAL, including deterministic MSY, various yield per recruit statistics, MCY, CAY, and CSP.

Information about MCMC output file formats is given earlier in Section 2.1.

6.1 Printouts from CASAL

CASAL prints out a bunch of different things to standard output. Some of them appear automatically, others you have to ask for. We should emphasize that the exact content of these outputs can be expected to change without notice. The best way to find out exactly what CASAL prints is to run it and find out.

The main types of printouts are:

- An initial header, giving the command by which CASAL was invoked, the date, the version numbers of the key source files used to build that copy of CASAL, the version number of CASAL itself, user login, and machine name.
- The names of any additional output files that were generated, such as MCMC output dumps.
- The results of the particular task asked for. If you run the model or estimate the parameters (`casal -e, -E, -r`), CASAL will print out the free parameters, the objective function and its components. If you profile some parameters (`casal -p`), CASAL will print out the objective function value and the free parameter estimates, for each value of each profiled parameter. If you are doing MCMC (`casal -m`), CASAL will print the initial point estimate, the approximate covariance matrix, the lower and upper bounds on the free parameters during MCMC, the start of the chain, and any changes in step size, etc.
- Printouts from the population section of CASAL. You can ask for printouts of the requests sent to the population section by the estimation and output sections and the corresponding results. You can request printouts of the initial state, the final state, the state after every year or every step. These are mostly useful in debugging, i.e., you can inspect them to figure out whether the population dynamics are what was intended. The most important of these printouts is 'population_section' which gives a text explanation of how CASAL interprets the population section in the `population.csl` file — always look at this printout when you develop a model to ensure that you have correctly specified the model.
- These population printouts only appear if you ask for them in the `output.csl` file. You may want to use the `-q` switch to suppress these printouts, because a major job

can generate a huge amount of them. (You can achieve the same result by turning all these requests off in `output.csl`.)

- Printouts from the estimation section of CASAL. You can ask CASAL to print out the parameter lists generated from your `population.csl` and `estimation.csl` files, a good way of checking that your files were read as intended. You can ask for fits, residuals, and standardised residuals. For debugging purposes, you can ask CASAL to print out the objective, parameters, or fits every time the objective function is calculated (so they are printed at each step of a minimization). You can ask for a text explanation of how CASAL thinks your estimation section works — always look at this printout when you develop a model. Also check out the list of parameters that were never accessed by CASAL. Presence of a parameter on the list may indicate that the parameter name was spelt incorrectly. (Or it may just mean that the parameter is not used for the task you were doing, for example `max_iters` is not used by `casal -r`). All these estimation printouts only appear if you ask for them in the `output.csl` file.
- Output quantities. These are model outputs calculated from the parameters. They can be produced for any set of free parameters, whether it comes from a model run (`casal -r`), a point estimate (`-e`, `-E`), for values sampled from a Bayesian posterior (`-v`), or for projections (`-P`). The output quantities CASAL can produce are listed in Section 6.2. Output quantities only appear if you ask for them in the `output.csl` file.

6.2 Output quantities

A variety of CASAL outputs are classed as ‘output quantities’. They can be produced for any set of free parameters, whether it comes from a model run (`casal -r`), a point estimate (`-e`, `-E`), for values sampled from a Bayesian posterior (`-v`, see Section 2.1), or for projections (`-P`, see Section 6.3).

Output quantities produced by model runs or point estimates are printed in a verbose format. They are marked with asterisks (*) in the output, which clearly identifies them for reading into statistical packages such as S-Plus. When output quantities are produced for samples from a Bayesian posterior or for projections, many sets of quantities are generated, so the results are dumped to a file in a columnar format instead.

Output quantities include the following:

1. The values of parameters. You can ask for ‘all free parameters’, and/or you can list the names of parameters, which need not be free. If you ask for an ogive, CASAL supplies the values of the ogive rather than the ogive arguments. If you ask for a size-based ogive in an age-based model, CASAL supplies the values of the ogive at the sizes given in the output parameter `print_sizebased_ogives_at`. If the size-based ogive is a selectivity, then probably a better way to extract its values is to use the `selectivity_at` pseudo-observations class (see below).
2. The arguments of ogive parameters (as opposed to the values, which see (1) above).
3. Spawning stock biomasses, for each stock in each model year (SSBs, see Section 4.3).

4. Recruitments, as absolute numbers of fish of each stock, by the year in which they recruit (see Section 4.4.2).
5. YCS, as deviates, by the year in which they spawn (see Section 4.4.2).
6. ‘True YCS’, defined as $YCS \times CR \times SR$, by the year in which the fish spawn (see Section 4.4.2),
7. The climate variable T by year (see Section 4.4.2).
8. Actual catches, optionally by stock (see Section 4.4.6).
9. Fishing pressures, by fishery, for each year (see Section 4.4.6).
10. B_0 , R_0 , B_{mean} , R_{mean} , $B_{initial}$, and $R_{initial}$ for each stock (see Section 4.5).
11. Nuisance q ’s (see Section 5.7.2).
12. The ‘stock crash’ quantity used to calculate stock risk (see Section 6.3).
13. Fits, residuals, Pearson residuals, and normalised residuals.
14. Pseudo-fits (see below).

Pseudo-fits are a bit of a special case. A *pseudo-fit* is an output defined as the fits to a set of *pseudo-observations*, fake observations which did not occur. This seems like an odd way of doing things, but in fact enables us to produce a number of useful outputs. For example, the total biomass for each model year, in a particular area and time step, can be generated as a pseudo-fit to an abundance series. If you want to see the selected biomass, add a selectivity to the pseudo-observations. If you want mature biomass, specify that the pseudo-observations include mature fish only. Similarly you can generate a combined biomass over all areas, biomass of a particular stock, total numbers rather than biomass, etc. You can also inspect the age or length composition of the fish by using numbers-at, proportions-at, or catch-at pseudo-observations. You cannot use pseudo-fits to age/size data, however.

One kind of “observation”, `selectivity_at`, can only ever be used as a pseudo-observation. It is intended for extracting the values of selectivity ogives, for each age/size class in the partition, in a particular year, time step, area, etc. It is particularly useful for extracting the values of a size-based ogive in an age-based model, because it converts them into values-at-age. This provides a one-step method for finding the actual ogive values being used by the model.

To ask for pseudo-fits, you just need to include the pseudo-observations in your `output.csl` file, in the same way that you include real observations in your `estimation.csl` file. The only differences are:

1. Don’t use relative observation types, i.e., `relative_abundance` and `relative_numbers_at`. Use the absolute equivalents instead.
2. Don’t supply the actual observation values — there aren’t any.
3. Don’t supply an error distribution, c.v.s, effective N ’s, weights, etc.

For example, insert these commands in the `output.csl` file to get outputs of total biomass across all areas, halfway through the mortality in time step 2, for all model years (1970 to 2000).

```
@abundance total_biomass
# output quantity: total biomass in all areas
biomass true
all_areas true
step 2
proportion_mortality 0.5
years 1970 1971 1972 1973 1974 1975 1976 1977 ... 1998 1999 2000
```

For projected output quantities (see Section 6.3.2), the range of years should extend into the future (up to year `final`).

6.3 Projections

Projection is the process of running the model forwards into the future, using randomised recruitments and hypothetical catches. CASAL does this in three situations:

1. Calculation of current annual yield (CAY) (see Section 6.5.1).
2. Calculation of current surplus production (CSP) (see Section 6.5.2).
3. Producing projected fishery performance estimators (FPIs), such as stock risk, or expected biomass in 5 years time.

All three situations use the same method for generating projections, which is described in Section 6.3.1. The calculation of FPIs is discussed in Section 6.3.2.

6.3.1 Carrying out projections

Projections can either be *point-based* (i.e., using a single point estimate of the free parameters), or *sample-based* (using a sample from the posterior distribution, typically generated by MCMC using `casal -m`, or `casal -C`).

For point-based projections CASAL does a large number of simulations, each using the same parameters. The simulations will differ only in terms of the randomised recruitments. Year class strengths will be randomised for the cohorts which will recruit in the ‘projection period’, i.e., the years `current+1` to `final`. You can also choose to randomise YCS for cohorts which have recently recruited (perhaps because there is no information about the abundance of these cohorts). If there is an explicit climate-recruitment relationship, CASAL uses the climate data T up until the last year for which it is provided (which could be as late as the assessment year, or might even be a forecast for the future) and then randomises T for years after that.

For sample-based projections CASAL does one simulation for each posterior sample point. Each simulation will use a different set of parameters and a different set of randomised recruitments. YCS and T 's will be randomised as above (the only difference is that the user might not need to randomise some recent YCS if their uncertainty was incorporated in the posterior distribution).

In either case, the ‘projected expectation’ of a quantity refers to an average over all the simulations.

When doing projections so as to calculate fishery performance indicators, you need to specify future catches for each fishery in each year. CASAL does not implement adaptive harvest strategies in projections. If you want to assess a different catch scenario, you need to change the future catches in the data file and rerun the program.

You can choose between four methods of randomising the YCS:

1. *Lognormal*: The randomised YCS are lognormally distributed, with mean 1, and specified standard deviation and autocorrelation on the log-scale. $YCS_i = \exp(X_i)$, where (X_i) are generated as a Gaussian AR(1) process with standard deviation σ_R and mean $-0.5\sigma_R^2$ (so that the mean of YCS_i is 1), and autocorrelation ρ . (Set $\rho = 0$ if you don’t want autocorrelation.)
2. *Lognormal-empirical*: The randomised YCS are lognormally distributed as per (1) above. The only difference is that the standard deviation parameter is chosen to give variability equal to that of the estimated YCS. CASAL uses $\sigma_R =$ the standard deviation of the log of the estimated YCS. Optionally, the calculation of σ_R can be based on a sub-range of the estimated YCS (since not all YCS are well estimated and some may even be fixed).
3. *Empirical*: The randomised YCS are resampled from the estimated YCS. Again, they can optionally be resampled from a sub-range of the estimated YCS.
4. *None*: All the randomised YCS are 1. Used for deterministic projections.

If a nonzero autocorrelation parameter ρ is used with lognormal or lognormal-empirical randomisations, then the randomised values must depend on the last fixed value YCS_f . This can get a bit ‘messy’.

Let $\mu_R = -0.5\sigma_R^2$, $X_f = \log\left(\frac{YCS_f - \mu}{\sigma_R^2}\right)$ and (Z_i) be standard normal random deviates, then

$$X_1 = \mu_R + \sigma_R\left(\rho X_f + \sqrt{1 - \rho^2} Z_1\right), \text{ and } X_{i+1} = \mu_R + \sigma_R\left(\rho X_i + \sqrt{1 - \rho^2} Z_{i+1}\right).$$

Now if the user specifies a very small or zero σ_R , probably in an effort to generate constant $YCS_i = 1$, and a nonzero ρ , and YCS_f is substantially different from 1, then the above formula gives an unexpected result, the YCS_i are not 1, but decay exponentially from YCS_f towards 1. This is because under these assumptions the value of YCS_f is highly implausible. CASAL avoids this situation by erroring out if $|X_f| > 5$ with “last non-random year has implausible value”. The user can fix the error by setting $\rho = 0$, increasing σ_R , or turning off randomisation (using method none).

We provide the same four methods for randomising the T_s associated with a climate-recruit relationship (see Section 4.4.2). The only difference is that the randomised T_s need not come from a distribution with mean 1. For lognormal randomisation, or no randomisation, the mean of the T_s is specified by the user. For lognormal-empirical randomisation, it is the mean of the estimated T_s .

So, to define the method of doing projections, you need to tell CASAL the following:

1. If projections are point based, the number of projections to be done.
2. Which is the first year for which YCS are randomised? The default is ($\text{current} - y_{\text{enter}} + 1$), which is the first year for which YCS are not provided. But you can specify an earlier year if they want to randomise abundance of some recently recruited cohorts. (Note that this is the year in which the fish spawn, not the year in which they recruit.)
3. Future catches for each fishery in each projected year. (This is only necessary for producing FPIs, not for CAY or CSP.)
4. The methods used to randomise YCS and T s, and the relevant parameters.

Of the above, only (1) is specified in the `output.csl` file. All the others are in `population.csl` since they relate to the recruitment variability of the population, and the catches.

6.3.2 Calculating projected fishery performance estimators (FPIs)

There are many fishery performance estimators (FPIs) commonly used in current New Zealand stock assessment. These include:

- Stock risk.
- $E(B_{\text{current}+k} / B_{\text{current}})$.
- $E(B_{\text{current}+k} / B_{\text{initial}})$.
- $E(B_{\text{current}+k} / B_0)$.
- $P(B_{\text{current}+k} > B_{\text{current}})$.

It is impractical to code all conceivable FPIs in CASAL, and it is more useful to dump the results of each individual projection into a text file, where you can use to generate your own FPIs in an external package such as S-Plus or Excel. Then if you want to calculate a different set of FPIs, you can do it without needing to redo the projections in CASAL.

Projected abundances and catches are requested from CASAL as output quantities (see Section 6.2). Call `casal -P` to run projections and generate the requested output quantities for the projected years. Use `-i filename` to pass CASAL a parameter estimate or a list of samples from the posterior. Projected actual catches and SSBs can be requested using the `quantities.actual_catches` and `quantities.SSBs` switches in `output.csl`. Various kinds of projected abundances can be requested by asking for abundance ‘pseudo-fits’ covering a range of years extending into the future. The projected results will then be sent to the output quantities file (the user must specify a filename as the argument of `-P`), which can be imported and processed by another package. There will be one row per projection.

When projections are point based, `casal -P` will also print out the expectation of each output quantity. This is intended as a shortcut so that some FPIs, such as $E(B_{\text{current}+k} / B_0)$, can be calculated without using a second software package. (Just divide the expected SSB for year $\text{current} + k$ by B_0 .) On the other hand some FPIs cannot be calculated using this method, such as $P(B_{\text{current}+k} > B_{\text{current}})$. You will need to use an external package to calculate these FPIs. The stock risk is a commonly used output quantity, defined as the probability that the SSB will fall below 20% B_0 in the projection period (for each stock). To allow a shortcut method

for calculating stock risk, we provide a projected output quantity `stock_crash`, which is defined as 1 if the SSB falls below 20% B_0 in the projection period or 0 otherwise (for each stock). Then you can read off the stock risk as the expectation of `stock_crash`.

6.4 Deterministic yield calculations

CASAL implements two kinds of deterministic yields, *per-recruit* analyses (Section 6.4.1) and *deterministic MSY* (Section 6.4.2). They are deterministic in the sense that they are based on simulations which use non-random recruitment with $YCS = CR(T) = 1$ (and hence the recruitment in year y is $R_y = R_0 \times SR(SSB_{y-y_{enter}})$). The calculations are based on a single set of parameters (i.e., a point estimate), supplied with `-i`.

Deterministic yields can be calculated only for single-stock models in CASAL (they may be implemented for multiple-stock models at a later stage).

Deterministic yield calculations are based on simulations at a constant *mortality rate* F . This “mortality rate” can be defined in several ways:

1. If there is only one fishery, then the mortality rate can be defined as an exploitation rate, which is the catch divided by a pre-fishery measure of biomass B_{pre} . Thus, the catch for each year is FB_{pre} . By default, B_{pre} is defined as the unselected mature biomass in all areas combined, in the time step of the fishery, before any mortality is applied, but you can change this definition. The important thing is that it must come before the start of the mortality episode in which the catch is taken.
2. If there are multiple fisheries, then the mortality rate must be defined as an exploitation rate as above. Again, B_{pre} is defined as the unselected mature biomass in all areas combined, in the time step of the fishery, before any mortality is applied, but you can change this definition. You have to provide a *catch split*, i.e., the proportion of the annual catch which must come from each fishery. Once the catch for the year has been calculated, it is split between fisheries according to this ‘catch split’.
3. Alternatively, if there is only one fishery and the Baranov catch equation is used, then you can opt to define the mortality rate as the instantaneous mortality rate of the Baranov equation. This is a more conventional method and may be required for comparability with other modelling work.

Note that for options 1 and 2, it may be impossible to take the catch even when $F < 1$, or alternatively it may be possible to take the catch even when $F > 1$, depending on the definition of B_{pre} .

B_{pre} is also used in CAY calculations (Section 6.5.1). The catch split is also used for MCY/CAY calculations and for CSP (Section 6.5.2).

6.4.1 Yield per recruit analyses

Per-recruit analyses are based on yield per recruit (YPR) and/or SSB per recruit (SPR). You can ask for any or all of the following:

- Data to plot a YPR curve (YPR versus mortality rate) or an SPR curve (SPR versus mortality rate).
- F_{max} , the mortality rate which maximizes YPR.

- $F_{0.1}$, the mortality rate at which the slope of the YPR curve is 0.1 times its slope at the origin (Gulland & Boerema 1973).
- $F_{x\%}$, the mortality rate at which the SPR is $x\%$ of its unfished value (Clark 1991).

$F_{0.1}$ should be calculated only if the mortality rate is an instantaneous rate (an $F_{0.1}$ based on exploitation rates could be calculated, but it is not clear that this would be a “safe” rate of fishing, in the way that the $F_{0.1}$ based on instantaneous rates is believed to be).

Each calculation of YPR or SPR works as follows. A single simulation run is done, starting from an unfished equilibrium state, and running until the catch and SSB stabilize. Having reached convergence, the total annual catch, SSB, and annual number of recruits are recorded, and YPR (total annual catch divided by number of recruits) and/or SPR (SSB divided by number of recruits) are calculated.

Traditionally, per-recruit analyses are done without a stock-recruitment relationship. However, it makes no difference either way, so long as the model divides by the actual number of recruits when calculating per-recruit statistics. However, we have found that finding the deterministic equilibrium with a high fishing pressure and a strong stock-recruitment relationship can take many, many simulated years. We recommend turning the stock-recruitment relationship off for per-recruit analyses to speed up calculations.

You need to provide an initial guesstimate of F , which is used to start off the minimiser for the estimates of F_{max} , $F_{0.1}$, and $F_{x\%}$. Providing a value in the right ballpark will help the minimiser find a more accurate solution.

6.4.2 Deterministic MSY

MSY_{det} is the maximum constant annual catch (using the specified catch split if there is more than one fishery) which can be sustained under deterministic recruitment. The corresponding mortality rate is F_{MSYdet} , and the corresponding SSB is B_{MSYdet} . Both MSY_{det} and B_{MSYdet} are expressed as percentages of B_0 .

Simulations for deterministic MSY work in the same way as the per-recruit simulations in Section 6.4.1. For each simulation run with mortality F , the equilibrium total annual catch C_F and spawning stock biomass SSB_F are calculated. CASAL searches over mortality rates to find F_{MSYdet} , the value of F that maximizes C_F . Then MSY_{det} and B_{MSYdet} are C_F and SSB_F respectively (expressed as percentages of B_0).

As well as calculating the MSY, you can request data with which to plot a yield versus SSB curve. You need to tell CASAL the mortality rates F at which SSB and yield are to be calculated.

The results of a deterministic MSY analysis depend heavily on the stock-recruitment relationship used. You have to specify one, even if it is ‘none’. Note that you can specify a different stock-recruitment relationship for simulations from the one used in ordinary model runs, using the `SR_simulation` and `steepness_simulation` parameters (Section 7.5).

You need to provide an initial guesstimate of F_{MSYdet} , which is used to start off the minimiser. Providing a value in the right ballpark will help the minimiser find a more accurate solution.

6.5 Stochastic yield estimates

CASAL implements two kinds of stochastic yields, MCY and CAY (Section 6.5.1) and current surplus production (CSP, Section 6.5.2). They are stochastic in the sense that they are based on simulations which use randomised recruitments. They can be either point-based or sample-based. Unlike deterministic yields, they can be calculated for multi-stock models.

6.5.1 MCY/CAY

Calculation of these yields is based on (and extends) the current NIWA procedures described by Francis (1992). Simulations are carried out to maximise yields, under either constant-catch or constant-mortality-rate harvesting, subject to the constraint that SSB should not fall below pB_0 more than proportion q of the time (defaulting to the traditional $p = 0.2$, $q = 0.1$).

By default, the risk constraint in CASALs MCY/CAY analysis specifies that the spawning stock biomass falls below pB_0 less than $q \times 100\%$ of the time. There is also an option to replace B_0 by a different reference biomass, which must be the spawning stock biomass (SSB) of the stock for some year between *initial* and *current*. For example, if the stock was believed to be in good condition in 1985, you could specify that the spawning stock biomass falls below pB_{1985} less than $q \times 100\%$ of the time (where B_{1985} is the spawning stock biomass in 1985).

CAY calculations are based on simulations at a constant *mortality rate* F . As per deterministic yields, this mortality rate can be defined either as an exploitation rate — catch/pre-fishery biomass B_{pre} — or, if there is only one fishery and the Baranov equation is used, as the instantaneous mortality rate of the Baranov equation. For both MCY and CAY calculations, if there are multiple fisheries, you have to provide a *catch split*, i.e., the proportion of the annual catch which must come from each fishery. These issues are discussed in more detail in Section 6.4. Note that B_{pre} is also used in deterministic yield calculations (Section 6.4) and the catch split is also used for deterministic yield calculations and for CSP (Section 6.5.2).

For each of a series of *harvest rates*, H (either a constant catch or a constant mortality rate) many simulation runs are carried out. Each simulation starts from a state which has stabilised under harvest rate H with deterministic recruitment (as per the deterministic simulations in Section 6.4). The run extends over $n_{discard} + n_{keep}$ years with stochastic recruitment. You need to choose both $n_{discard}$ and n_{keep} . Hopefully in the long term we will determine good default values. You need to choose a value of $n_{discard}$ which is large enough to allow the population to stabilize under harvest rate H by the end of $n_{discard}$ years. Francis (1992) recommends $n_{keep} =$ the approximate maximum age of the species $= \log_e(100) / M$ (the natural mortality rate). We print $E(SSB_{n_{discard}})$ and $E(SSB_{n_{discard} + 1})$ as diagnostics. If the two are about equal, then $n_{discard}$ may be large enough. Try also using different values of $n_{discard}$ and n_{keep} and seeing if it makes any difference to the results.

With one stock, for each run, CASAL will calculate, over the final period of n_{keep} years, the mean catch taken over all fisheries C_{av} , the mean SSB B_{av} , and the proportion P_{risk} of years in which the SSB falls below pB_0 . These quantities will then be averaged over all runs with harvest rate H to calculate $C_{av}(H)$, $B_{av}(H)$, $P_{risk}(H)$. The program then searches for the “optimal” harvest rate H_{opt} , which is the value of H which maximises $C_{av}(H)$, subject to the constraint that $P_{risk}(H) \leq q$. Note that the search may take quite a while, depending on how many simulations you do, and you may want to interrupt it once it reaches a solution which is good enough for your purposes. You may alternatively want to search manually, interactively supplying a sequence of harvest rates. In this case, you will be prompted to input a trial H , CASAL will print $C_{av}(H)$, $B_{av}(H)$, and $P_{risk}(H)$, you will be prompted for a new H , etc. When

you are satisfied, enter a negative value, meaning ‘stop here’. The last value of H you provided will be used to calculate yields.

CASAL then calculates yields. Constant-catch simulations give $MCY = H_{opt}$ (a target catch in tonnes) or if, according to the current assessment, the stock is depressed (i.e., $E(\text{current SSB} / B_0) < 0.2$, where the expectation is over the parameter sets provided), then $MCY = H_{opt} \times E(\text{SSB}_{\text{current}} / 0.2B_0)$ (see Section 4.5 of Francis 1992). This adjusted value is sometimes labelled the ‘current MCY’ to distinguish it from the ‘long-term MCY’ H_{opt} . Whichever MCY is calculated, $B_{MCY} = B_{av}(H_{opt})$. Constant-mortality-rate simulations give $F_{CAY} = H_{opt}$, $MAY = C_{av}(H_{opt})$, and $B_{MAY} = B_{av}(H_{opt})$. The calculation of CAY for next year requires a 1-year projection (Section 6.3). It is the expected catch in the projected year, under a mortality rate of F_{CAY} .

If an exploitation rate constraint is broken during a deterministic simulation, then the catch level is clearly too high. We don’t do the following stochastic simulation, and instead take $C_{av} = 0$, $B_{av} = 0$, $P_{risk} = 1$.

For multiple stocks, we had to reinvent the definitions of MCY and CAY. The quantities C_{av} , B_{av} , and P_{risk} are calculated for each stock separately. The “optimal harvest rate” H_{opt} is now the value of H which maximizes $\sum_s C_{av,s}$, subject to the constraint that $P_{risk,s}(H) \leq q$ for all s (where s indexes the stocks). There is no obvious way to split the MCY between stocks, and nor is it clear how, if at all, the MCY should be modified if one or more of the stocks is depressed. However we can calculate the B_{MCY} for each stock, $B_{MCY,s} = B_{av,s}(H_{opt})$. For CAY simulations, $F_{CAY} = H_{opt}$, $MAY_s = C_{av,s}(H_{opt})$, $B_{MAY,s} = B_{av,s}(H_{opt})$, and next years CAY_s is the expected catch in the projected year from stock s , under a mortality rate of F_{CAY} .

Where there are multiple stocks with one TAC per stock and no multi-stock fisheries, the natural approach is to calculate MCYs and CAYs separately for each stock using a catch split in which all the catch comes from a single stock.

CASAL provides four methods for generating random recruitments for the simulation period — “lognormal”, “lognormal-empirical”, “empirical”, and “none”. These methods are described in Section 6.3.1. The only major difference is that there is no ‘last non-random year’. All the YCS and T ’s are random.

As well as recruitment variability, simulations can incorporate uncertainty in several different ways:

1. For sample-based simulations, the Bayesian posterior is meant to express the uncertainty in the free parameters. One simulation run is done for each sample from the posterior (c.f. point-based simulations where many simulations are done using the single set of parameters). Note that if either “empirical” method of randomising recruitment is used for sample-based simulations then the recruitment variability will differ between individual simulations.
2. For point-based MCY simulations, the uncertainty associated with virgin biomass can be incorporated (as in the stock assessment software *pmod*, R.I.C.C. Francis, unpublished). For each year in the i th simulation run with target catch H , the actual catch taken will be $H\varepsilon_i$. This is intended to simulate what would happen if the true virgin abundance was B but was thought to be $B\varepsilon_i$. The ε_i are assumed to be i.i.d., either lognormal or normal with negative values increased to 0, with mean 1 and default c.v. 0.2 (as assumed by Francis 1992, where the normal distribution was used). These errors are not used in sample-based simulations because uncertainty in virgin abundance is meant to be incorporated in the posterior.

3. For both point- and sample-based CAY simulations, the annual stock-assessment uncertainty can be incorporated (as in the stock assessment software *pmod*, R.I.C.C. Francis, unpublished). For each year y in the i th simulation run with target mortality rate H , the catch will be calculated using $F = H\varepsilon_y$. This is intended to simulate what would happen if the true abundance was B but was thought to be $B\varepsilon_y$. The ε_y are assumed to be i.i.d., either lognormal or normal with negative values increased to 0, with mean 1 and default c.v. 0.2 (as assumed by Francis (1992), where the normal distribution was used). [Not implemented for the case where F is an instantaneous mortality rate.]

The same random numbers are used for the simulation runs at each harvest rate H . This increases comparability (e.g., between $C_{av}(H_1)$ and $C_{av}(H_2)$) and removes random noise from the $C_{av}(H)$ and $P_{risk}(H)$ curves.

The results of an MCY or CAY analysis depend heavily on the stock-recruitment relationship used. You have to specify one, even if it is ‘none’. Note that you can specify a different stock-recruitment relationship for simulation from the one used in ordinary model runs, using the `SR_simulation` and `steepness_simulation` parameters (Section 7.5).

You need to provide initial guesstimates of MCY and F_{CAY} , which are used to start off the minimiser. Providing a value in the near the ‘true value’ will help the minimiser find a solution faster.

6.5.2 Current surplus production (CSP)

CASAL defines the current surplus production (CSP) as the catch in year `current+1` which would make the projected expectation of post-fishery biomass B_{post} in year `current+1` equal to that in year `current`. The calculation of CSP is hence based on 1-year projections (Section 6.3), so you must set `final` to at least `current+1`.

CASAL defines the post-fishery biomass B_{post} as the unselected mature biomass in all areas combined, in the time step of the last fishery, after all mortality has been applied. You can change this definition if you want. If there are multiple fisheries then you have to specify a catch split (the proportion of the catch that must come from each fishery in year `current+1`). The same catch split will be used for deterministic yield calculations (Section 6.4) and for MCY/CAY (Section 6.5.1).

For a multiple-stock model, you have the choice of two approaches. You can either request an overall CSP (as above) or a CSP for each stock, in which case CASAL does the following for each stock s :

1. Redefines B_{post} to only include fish of stock s .
2. Finds the total catch in year `current+1` which would make the projected expectation of B_{post} in year `current+1` equal to that in year `current`.
3. Returns the expected amount of that catch which comes from stock s .

It is possible that there will be no CSP, i.e., even if no catch is taken, there is a drop in expected B_{post} .

You need to provide an initial guesstimate of CSP, which is used to start off the minimiser. Providing a value near the ‘true value’ will help the minimiser find a more accurate solution.

7. THE POPULATION.CSL FILE

The population parameters are specified in the `population.csl` file. See Section 4 for information about the population section, and Section 2.4 for instructions on writing a CASAL data file.

7.1 Defining the partition

@size_based	Should the model be size-based rather than age-based?
Type	Switch
Default	False (i.e., age-based)
Effects	Defines the model as either age-based or size-based.
@n_classes	Number of size classes
Conditions	Must be specified in a size-based model. Ignored in an age-based model.
Type	Integer
Effects	Defines the number of size classes
@class_mins	Size class lower limits (plus the upper limit of the last class if it is not a plus group)
Conditions	Must be specified in a size-based model. Ignored in an age-based model.
Type	Constant vector
Effects	Defines the lower limits of each of the <code>n_classes</code> size classes. If there is no plus group then an additional value defines the upper limit of the last size class.
@min_age, @max_age	Minimum and maximum age limits
Conditions	Must be specified in an age-based model. Ignored in a size-based model.
Type	2 x integer
Effects	Defines the minimum and maximum fish age classes.
@plus_group	Should a plus age or size group be used?
Type	Switch
Default	True (use a plus group)
Effects	Defines the last age or size class as a plus group.
@plus_group_size	Mean size of plus group
Conditions	Must be specified in a size-based model with a plus group. Otherwise ignored.
Type	Constant
Effects	Defines the nominal size of the plus group. Used for ogives and for mean weight calculations.
@sex_partition	Is the partition sex-structured?
Type	Switch
Default	False (the partition is not sex-structured)
Effects	Defines whether sex is a character in the partition
@mature_partition	Is the partition structured by maturity?
Type	Switch
Default	False (the partition is not structured by maturity)
Effects	Defines whether maturity is a character in the partition
@n_areas	Number of areas in the partition
Type	Integer
Default	1 (a single-area model)
Effects	If <code>n_areas=1</code> , then area is not a character in the partition. Otherwise, <code>n_areas</code> is the number of areas in the partition.

@area_names	Area names
Conditions	Necessary if <code>n_areas > 1</code> , otherwise ignored
Type	Vector of strings
Effects	Defines the text label to be associated with each area.
@n_stocks	Number of stocks in the partition
Type	Integer
Default	1 (a single-stock model)
Effects	If <code>n_stocks=1</code> , then stock is not a character in the partition. Otherwise, <code>n_stocks</code> is the number of stocks in the partition.
@stock_names	Stock names
Conditions	Necessary if <code>n_stocks > 1</code> , otherwise ignored
Type	Vector of strings
Effects	Defines the text label to be associated with each stock.
@n_growthpaths	Number of growth-paths in the partition
Type	Integer
Default	1 (not a growth-path model)
Effects	If <code>n_growthpaths=1</code> , then growth-path is not a character in the partition. Otherwise, <code>n_growthpaths</code> is the number of growth-paths in the partition.
@exclusions_char1	Partition exclusion term 1
@exclusions_val1	Partition exclusion value 1
@exclusions_char2	Partition exclusion term 2
@exclusions_val2	Partition exclusion value 2
Type	String vector
Conditions	All partition exclusion commands must be used if an exclusion is defined
Effects	Defines what combinations of characters are excluded from the partition. There is no row in the partition for which the character <code>exclusions_char1[i]</code> takes the value named <code>exclusions_val1[i]</code> and the character <code>exclusions_char2[i]</code> takes the value named <code>exclusions_val2[i]</code>
Example	If no females are allowed in area “home”, use entries of “sex”, “female”, “area”, and “home” respectively.
Notes	Exclusions are never necessary but can improve the model’s execution speed.

7.2 Defining the annual cycle and the time sequence

@initial	Initial assessment year
Type	Integer
Effects	Defines the first year of the assessment period.
@current	Current assessment year
Type	Integer
Effects	Defines the last year of the assessment period, excluding the projection period (if there is one).
@final	Final projection year
Type	Integer
Effects	Defines the last year of the projection period.
@annual_cycle	Annual cycle block command
Effects	Defines any following commands as <code>@annual_cycle</code> subcommands

time_steps	Number of time steps
Command	annual_cycle
Type	Integer
Effects	Defines the number of time steps in the annual cycle
recruitment_time	Time step in which recruitment occurs
Command	annual_cycle
Type	Integer
Effects	Defines the time step in which recruitment occurs
recruitment_areas	Area in which where recruitment occurs, for each stock
Command	annual_cycle
Conditions	Necessary if n_areas>1, otherwise ignored
Type	Vector of strings
Effects	Defines the area in which recruitment occurs, for each stock
Notes	Each entry should be an area label as per area_names. You need one entry per stock, even if all stocks recruit in the same area.
spawning_time	Time step for recording SSB
Command	annual_cycle
Type	Integer
Effects	Defines the time step in which the value of the spawning stock biomass is recorded
spawning_part_mort	Proportion of mortality in the time step before recording SSB
Command	annual_cycle
Type	Constant
Default	0.5
Effects	Defines the proportion of the time step's mortality episode to apply before recording SSB.
Notes	Should be a real number in [0,1]. Has no effect if there is no mortality in the time step.
spawning_areas	Area for recording SSB, for each stock
Command	annual_cycle
Conditions	Either spawning_areas or spawning_all_areas is necessary if n_areas > 1, otherwise ignored.
Type	Vector of strings
Effects	Defines the area in which to record SSB for each stock.
Notes	Each entry should be an area label as per area_names. Alternatively set spawning_all_areas.
spawning_all_areas	Is SSB recorded for all areas combined?
Command	annual_cycle
Conditions	Either spawning_areas or spawning_all_areas is necessary if n_areas>1, otherwise ignored. Only usable in a single-stock model.
Type	Switch
Effects	Defines that SSB is recorded for all areas combined.
Notes	Alternatively use spawning_areas.
spawning_ps	Spawning proportions by age/size class
Command	annual_cycle
Conditions	Specify either spawning_ps or spawning_p, but not both
Type	Estimable vector
Effects	Defines the factor applied to mature biomass to get the SSB for each stock.
Notes	This must have length equal to the number of stocks. In the special case with only one stock, then you can use either spawning_ps OR spawning_p.

spawning_p	Spawning proportion
Command	annual_cycle
Conditions	Specify either spawning_ps or spawning_p, but not both
Type	Estimable
Effects	Defines the factor applied to mature biomass to get the SSB for each stock.
spawning_use_total_B	Should SSB be defined as total biomass rather than mature biomass?
Command	annual_cycle
Conditions	Can be specified only if maturity is not a partition character
Type	Switch
Default	In single-area models, false (SSB = mature biomass). Otherwise, no default.
Effects	Defines the SSB as the total biomass in the spawning area rather than the mature biomass.
Notes	The only reason to use this, we believe, is if you have migrated 'mature' fish to a spawning area when maturity is not a partition character. You know all the fish in the spawning area are mature, but CASAL doesn't (because the state does not keep track of maturity).
n_growths	Number of growth episodes per year
Command	annual_cycle
Conditions	Can be used only in a size-based model
Type	Integer
Effects	Defines the number of growth episodes per year
growth_times	Time step in which each growth episode occurs
Command	annual_cycle
Conditions	Can be used only in a size-based model
Type	Integer
Effects	Defines the time step in which each growth episode occurs
Notes	Number of entries = annual_cycle.n_growths
aging_time	Time step when age is incremented
Command	annual_cycle
Conditions	Can be used only in an age-based model
Type	Integer
Effects	Defines the time step when ageing occurs
growth_props	Proportion of growth that has occurred by each time step
Command	annual_cycle
Conditions	Can be used only in an age-based model
Type	Constant vector
Default	Vector of zeros, i.e., no growth between fish birthdays
Effects	Defines the proportion of the year's growth which has occurred by the start of each time step.
Notes	The mean size of fish of age <i>a</i> years (rounded down) in the <i>i</i> th time step is calculated as if their age was ($a + \text{growth_props}[i]$). growth_props[aging_time] must be 0, and the entries of growth_props have to be non-decreasing between fish birthdays.
Example	If the first entry of growth_props is 0.5, then, in time step 1, the mean size of 2-year-old fish is calculated as if they were age 2.5.
M_props	Proportion of natural mortality that occurs in each time step
Command	annual_cycle
Type	Constant vector
Effects	Defines the proportion of the year's natural mortality which occurs in each time step
baranov	Should fishing mortality be applied simultaneously with natural mortality using the Baranov equation, rather than instantaneously?
Command	annual_cycle
Type	Switch

Default	False (i.e., use instantaneous mortality)
Effects	Defines fishing and natural mortality to be applied simultaneously using the Baranov equation, rather than as half a time step's natural mortality, then instantaneous fishing mortality, then the remaining half of the natural mortality.
Notes	You cannot use the Baranov equation if you have more than one fishery in the same area in the same time step.

midmortality_partition Method to calculate mortality within the time step

Command	annual_cycle
Type	String
Default	weighted_sum if fishing mortality is instantaneous, weighted_product if the Baranov equation is used
Effects	Defines how the partition is calculated partway through a mortality episode. Must be either weighted_sum or weighted_product.
Notes	See Section 4.3. The defaults are usually sensible.

fishery_names Names of the fishery

Command	annual_cycle
Type	Vector of strings
Effects	Defines the text labels of the fisheries.

fishery_times Time step when each fishery occurs

Command	annual_cycle
Type	Constant vector
Effects	Defines the time step in which each fishery occurs.
Notes	One entry per entry of fishery_names.

fishery_areas Area when each fishery occurs

Command	annual_cycle
Conditions	Necessary if n_areas>1, otherwise ignored
Type	Vector of strings
Effects	Defines the area in which each fishery occurs.
Notes	Each entry should be an area label as per area_names. One entry per entry of fishery_names.

n_migrations Number of migrations in each year

Command	annual_cycle
Conditions	Can be used only if n_areas>1
Type	Integer
Effects	Defines the number of migrations in each year

migration_names Names of the fishery

Command	annual_cycle
Conditions	Can be used only if n_migrations>0
Type	Vector of strings
Effects	Defines the text labels of the migrations.
Notes	Number of entries should be n_migrations.

migration_times Time step of each migration

Command	annual_cycle
Conditions	Can be used only if n_migrations>0
Type	Constant vector
Effects	Defines the time step in which each migration occurs.
Notes	Number of entries should be n_migrations.

migrate_from Area from which each migration departs

Command	annual_cycle
Conditions	Can be used only if n_migrations>0
Type	Vector of strings
Effects	Defines the source area of each migration

Notes	Each entry should be an area label as per <code>area_names</code> . Number of entries should be <code>n_migrations</code> .
migrate_to	Area where each migration arrives
Command	<code>annual_cycle</code>
Conditions	Can be used only if <code>n_migrations > 0</code>
Type	Vector of strings
Effects	Defines the destination area of each migration
Notes	Each entry should be an area label as per <code>area_names</code> . Number of entries should be <code>n_migrations</code> .
n_maturations	Number of maturation episodes in each year
Command	<code>annual_cycle</code>
Conditions	Can be used only if maturity is a character in the partition
Type	Integer
Effects	Defines the number of maturation episodes in each year
maturity_times	Time step of each maturation episode
Command	<code>annual_cycle</code>
Conditions	Can be used only if <code>n_maturations > 0</code>
Type	Constant vector
Effects	Defines the time step in which each maturation occurs.
Notes	Number of entries should be <code>n_maturations</code> .
disease_mortality_time	Time step to apply disease mortality
Command	<code>annual_cycle</code>
Default	None
Conditions	Should only be used if disease mortality is to be applied
Type	Integer
Effects	If defined, then the disease mortality is applied in that time step

7.3 Defining recruitment

@y_enter	Number of years after which a year class enters the partition
Type	Integer
Effects	Defines the number of years after which a year class enters the partition.
Notes	In an age-based model, there is probably a 'correct' value of <code>y_enter</code> . You will be warned if you use any other value.
@standardise_YCS	Use the Haist parameterisation for YCS?
Type	Switch
Default	False (do not use the Haist parameterisation)
Effects	Defines YCS to use the Haist parameterisation in Section 4.4.2.
Note	Do not use both <code>standardise_YCS</code> and <code>use_mean_YCS</code> . You will also need to set <code>@recruitment.first_free</code> and <code>@recruitment.last_free</code> for each stock.
@use_mean_YCS	Use the Francis parameterisation for YCS?
Type	Switch
Default	False (do not use the Francis parameterisation)
Effects	Defines YCS to use the Francis parameterisation in Section 4.4.2.
Note	Do not use both <code>standardise_YCS</code> and <code>use_mean_YCS</code> . You will also need to set <code>@recruitment.first_free</code> and <code>@recruitment.last_free</code> for each stock.
@recruitment	Recruitment block command
Label	the name of a stock
Effects	Defines any following commands as <code>@recruitment</code> subcommands for the stock
Notes	Omit the stock label if there is only one stock in the model.

YCS	Year class strengths for the stock
Command	recruitment [stock_name]
Type	Estimable vector
Effects	Defines the year class strengths for the stock
Notes	Entries should correspond to the years given in YCS_years. If @standardise_YCS is set to true, the Haist parameterisation of YCS is used. And If @use_mean_YCS is set to true, the Francis parameterisation of YCS is used.
YCS_years	Years for which YCS are provided
Command	recruitment [stock_name]
Type	Constant vector
Effects	Defines the years for which the YCS are provided
Notes	For each entry of YCS there should be a corresponding year in YCS_years. The years should cover the range $initial-y_enter+n_rinitial$ to $current-y_enter$, or a consecutive subset of these years if the Francis parameterisation is used.
n_rinitial	Number of years for which $R_{initial}$ is to be used as the YCS
Command	recruitment [stock_name]
Type	Integer
Default	0
Effects	If $n_rinitial=0$, do not use $R_{initial}$ as a YCS. If $n_rinitial>0$, use $R_{initial}$ as the YCS for the first $n_rinitial$ years. You then do not need to supply YCS for these year classes.
SR	Stock-recruitment relationship
Command	recruitment [stock_name]
Type	String
Default	none
Effects	Defines the stock-recruitment relationship. Should be BH (Beverton-Holt), Ricker, or none.
steepness	Steepness parameter of the stock-recruitment relationship
Command	recruitment [stock_name]
Conditions	Only used if SR=BH or Ricker
Type	Estimable
Effects	Defines the steepness parameter h of the stock-recruitment relationship.
CR	Climate-recruitment relationship
Command	recruitment [stock_name]
Type	String
Default	none
Effects	Defines the stock-recruitment relationship. Should be exponential, arctan, logistic, none, identity, or linear-combination.
Ts	Climate variable T
Command	recruitment [stock_name]
Conditions	Only used if there is a climate-recruitment relationship
Type	Constant vector
Effects	Defines the climate variable T for each year
Notes	Entries should correspond to the years given in Ts_years.
Ts_years	Years for which the climate variable T is provided
Command	recruitment [stock_name]
Conditions	Only used if there is a climate-recruitment relationship
Type	Constant vector
Effects	Defines the years for which the T s are provided
Notes	For each entry of T s there should be a corresponding year in Ts_years. The years should at least cover the range $initial-y_enter$ to $current-y_enter$, or can extend further forwards at your option.

CR_alpha, CR_beta Climate-recruitment parameters alpha and beta

Command	recruitment [stock_name]
Conditions	Only used if there is a logistic, arctan, or exponential climate-recruitment relationship
Type	Estimable
Effects	Defines the values of the α and β parameters in the climate-recruitment relationship

CR_beta2 Climate-recruitment parameter beta2

Command	recruitment [stock_name]
Conditions	Only used if there is a logistic climate-recruitment relationship
Type	Estimable
Effects	Defines the value of the β_2 parameter in the logistic climate-recruitment relationship

CR_p Climate-recruitment parameter p

Command	recruitment [stock_name]
Conditions	Only used if there is a linear-combination climate-recruitment relationship
Type	Estimable
Effects	Defines the values of the p parameter in the linear-combination climate-recruitment relationship

initial_size_mean Mean size at recruitment (both sexes)**initial_size_cv c.v. of size at recruitment (both sexes)**

Command	recruitment [stock_name]
Conditions	Only used in a size-based model. Use either <code>initial_size_mean</code> and <code>initial_size_cv</code> , or the sex-based equivalents.
Type	Estimable
Effects	Define the mean and c.v. of the size of recruiting fish

initial_size_mean_male Mean size at recruitment (male)**initial_size_mean_female Mean size at recruitment (female)****initial_size_cv_male c.v. of size at recruitment (male)****initial_size_cv_female c.v. of size at recruitment (female)**

Command	recruitment [stock_name]
Conditions	Only used in a size-based, sexed model. Alternatively use <code>initial_size_mean</code> and <code>initial_size_cv</code> .
Type	Estimable
Effects	Define the mean and c.v. of the size of recruiting fish, by sex.

p_male Proportion of recruits that are male

Command	recruitment [stock_name]
Conditions	Only used in a sex-based model.
Type	Estimable
Default	0.5
Effects	Defines the proportion of recruits that are male

growthpaths Proportion of recruits on each growth-path

Command	recruitment [stock_name]
Conditions	Only used for growth-path models.
Type	Constant vector
Effects	Defines the proportions of recruits that follow each of the growth-paths
Notes	Should be a vector with one element per growth-path, summing to 1.

first_free, last_free Range of free YCS defining R_0 , with the Haist or Francis YCS parameterisation

Command	recruitment [stock_name]
Conditions	Can be used only if either <code>standardise_YCS</code> or <code>use_mean_YCS</code> is true.

Type	2 x integer
Default	initial-y_enter+n_rinitial, current-y_enter
Effects	Defines the range of years [first_free ... last_free] over which the average recruitment is defined to be R_0 .
Notes	Check that the YCS have come out right by using the quantities.YCS output quantity parameter.

7.4 Defining recruitment variability

@randomisation_method Randomisation method for recruitment variability in stochastic simulations and projections

Type	String
Default	No default
Effects	Defines the randomisation method to use in stochastic simulations or projections. Should be lognormal, lognormal-empirical, empirical, or none.

@first_random_year For projections, the first year for which YCS are randomised

Type	Integer
Default	The default is the first year for which YCS are not specified, i.e., current-y_enter+1.
Effects	Defines the first year for which YCS are randomised in projections.
Notes	Ignored unless you are doing projections. You may want to set it earlier than the default if the last one or more YCS are poorly estimated (and you are not operating in a Bayesian framework). You cannot set it later than the default. This is the year in which the fish are spawned, not the year in which they recruit.

Recruitment subcommands (specified separately for each stock) follow:

sigma_r	Standard deviation on the log scale of randomised YCS
T_sigma_r	Standard deviation on the log scale of randomised climate data T
Command	recruitment [stock_name]
Conditions	Only used if randomisation_method=lognormal. T_sigma_r is only used if there is a climate-recruitment relationship.
Type	Constant
Effects	sigma_r defines the standard deviation of log-YCS (σ_r) parameter for use in projections and stochastic simulations. T_sigma_r likewise for the climate variable T .
Notes	See also the rho, T_rho, T_mean parameters
rho	Lag-1 log-scale autocorrelation of randomised YCS
T_rho	Lag-1 log-scale autocorrelation of randomised climate data T
Command	recruitment [stock_name]
Conditions	Only used if randomisation_method=lognormal or lognormal-empirical. T_rho is only used if there is a climate-recruitment relationship.
Type	Constant
Default	0 (no autocorrelation)
Effects	rho defines the autocorrelation of log-YCS (ρ) parameter for use in projections and stochastic simulations. T_rho likewise for the climate variable T .
year_range	Year range from which randomised YCS are resampled
T_year_range	Year range from which randomised climate data T are resampled
Command	recruitment [stock_name]

Conditions	Only used if <code>randomisation_method=empirical</code> or <code>lognormal-empirical</code> .
Type	Constant vector
Default	all years in the original YCS, or all years in the original T 's
Effects	Defines the year range from which the resampling (nonparametric bootstrap) methods draw their samples, in projections and stochastic simulations.
Notes	Provide just two numbers, the first and last years, for each parameter. The program resamples YCS and T 's from these ranges of years (inclusive)
T_mean	Mean on the linear scale of randomised climate data T
Command	<code>recruitment [stock_name]</code>
Conditions	Only used if <code>randomisation_method=lognormal</code> or <code>none</code> and there is a climate-recruitment relationship.
Type	Constant
Effects	<code>T_mean</code> defines the mean of randomised climate data T , to be used in projections and stochastic simulations.
Notes	You need to provide <code>T_mean</code> even if there is no randomisation of climate data. Then all 'randomised' T 's are simply set to <code>T_mean</code> . See also the <code>T_sigma_r</code> parameter.

7.5 Defining recruitment in yield simulations

These are recruitment subcommands (specified separately for each stock):

simulation_SR	The stock-recruitment relationship to be used in yield simulations
Command	<code>recruitment [stock_name]</code>
Type	String
Default	If <code>SR</code> is explicitly provided, <code>simulation_SR</code> defaults to <code>SR</code> . If <code>SR</code> is not explicitly defined, the default is <code>none</code> .
Effects	Defines the stock-recruitment relationship used in yield simulations (Sections 6.4, 6.5). Should be <code>BH</code> (Beverton-Holt), <code>Ricker</code> , or <code>none</code> .
simulation_steepness	Steepness parameter of the stock-recruitment relationship to be used in yield simulations
Command	<code>recruitment [stock_name]</code>
Conditions	Only used if <code>simulation_SR=BH</code> or <code>Ricker</code> .
Type	Estimable
Effects	Defines the steepness parameter h of the stock-recruitment relationship used in yield simulations (Sections 6.4, 6.5).

7.6 Defining growth (in a size based model)

@growth	Growth block command
Conditions	Only used in a size-based model.
Effects	Defines any following commands as <code>@growth</code> subcommands for an episode. The i th <code>@growth</code> block relates to the i th growth episode in a year.
stock	Stock that the growth episode applies to
Command	<code>growth [i]</code>
Conditions	Only used in a size-based model.
Type	String
Default	(All stocks)
Effects	If supplied, defines the single stock that the growth episode applies to. If not supplied, the growth episode applies to all stocks equally.
type	Growth model used by the growth episode
Command	<code>growth [i]</code>
Conditions	Only used in a size-based model.

Type	String
Effects	Defines the growth model used. So far CASAL only has one growth model, <code>basic</code> .
g	Reference growths for the basic growth model
l	Reference sizes for the basic growth model
cv	c.v. for the basic growth model
minsigma	Lower bound on sigma for the basic growth model
Command	<code>growth[i]</code>
Conditions	Only used in a size-based model. Alternatively use the equivalents suffixed by <code>_male</code> and <code>_female</code> , <code>_mature</code> and <code>_immature</code> , or <code>_male_mature</code> , <code>_male_immature</code> , <code>_female_mature</code> , <code>_female_immature</code> below.
Type	Estimable vector, estimable vector, estimable, estimable (respectively)
Effects	Parameters of the <code>basic</code> growth model. <code>g</code> and <code>l</code> should be 2-vectors, with <code>g</code> containing reference growths <code>g_alpha</code> and <code>g_beta</code> and <code>l</code> containing reference sizes <code>l_alpha</code> and <code>l_beta</code> . <code>cv</code> is the c.v. around the reference line and <code>minsigma</code> the lower bound on standard deviation around the reference line.
g_male, g_female	Reference growths for the basic growth model
l_male, l_female	Reference sizes for the basic growth model
cv_male, cv_female	c.v. for the basic growth model
minsigma_male, minsigma_female	Lower bound on sigma for the basic growth model
Command	<code>growth[i]</code>
Conditions	Only used in a sexed, size-based model. Alternatively use the equivalents suffixed by nothing, <code>_mature</code> and <code>_immature</code> , or <code>_male_mature</code> , <code>_male_immature</code> , <code>_female_mature</code> , <code>_female_immature</code> .
Type	2 x estimable vector, 2 x estimable vector, 2 x estimable, 2 x estimable
Effects	Parameters of the <code>basic</code> growth model, specified by sex. See the unsexed versions above.
g_mature, g_immature	Reference growths for the basic growth model
l_mature, l_immature	Reference sizes for the basic growth model
cv_mature, cv_immature	c.v. for the basic growth model
minsigma_mature, minsigma_immature	Lower bound on sigma
Command	<code>growth[i]</code>
Conditions	Only used in a size-based model with maturity in the partition. Alternatively use the equivalents suffixed by nothing, <code>_male</code> and <code>_female</code> , or <code>_male_mature</code> , <code>_male_immature</code> , <code>_female_mature</code> , <code>_female_immature</code> .
Type	2 x estimable vector, 2 x estimable vector, 2 x estimable, 2 x estimable
Effects	Parameters of the <code>basic</code> growth model, specified by maturity. See the other versions above.
g_male_mature, etc.	Reference growths for the basic growth model
l_male_mature, etc.	Reference sizes for the basic growth model
cv_male_mature, etc.	c.v. for the basic growth model
minsigma_male_mature, etc.	Lower bound on sigma
Command	<code>growth[i]</code>
Conditions	Only used in a size-based model with sex and maturity in the partition. Alternatively use the equivalents suffixed by nothing, <code>_male</code> and <code>_female</code> , or <code>_mature</code> and <code>_immature</code> .
Type	4 x estimable vector, 4 x estimable vector, 4 x estimable, 4 x estimable

Effects Parameters of the basic growth model, specified by sex and maturity. See the other versions above.

7.7 Defining maturation (when maturity is in the partition)

@maturation	Maturation block command
Conditions	Only used in a model where maturity is a partition character.
Effects	Defines any following commands as maturation subcommands for an episode. The <i>ith</i> @maturation block relates to the <i>ith</i> maturation episode in a year.
stock	Stock that the maturation episode applies to
Command	maturation[i]
Conditions	Only used in a model where maturity is a partition character.
Type	String
Default	(All stocks)
Effects	If supplied, defines the single stock that the maturation episode applies to. If not supplied, the maturation episode applies to all stocks equally.
area	Area that the maturation episode applies to
Command	maturation[i]
Conditions	Only used in a model where maturity is a partition character.
Type	String
Default	(All areas)
Effects	If supplied, defines the single area that the maturation episode applies to. If not supplied, the maturation episode applies to all areas equally.
rates_all	Rates of maturation by age or size class
Command	maturation[i]
Conditions	Only used in a model where maturity is a partition character. Alternatively use sex-specific rates (<i>rates_male</i> , <i>rates_female</i>)
Type	ogive
Effects	Defines the rates of maturation by age/size class.
rates_male, rates_female	Rates of maturation by sex and age or size class
Command	maturation[i]
Conditions	Only used in a model where sex and maturity are partition characters. Alternatively use non-sex-specific rates (<i>rates_all</i>)
Type	2 x ogive
Effects	Defines the rates of maturation by sex and age/size class.

7.8 Defining maturity (when maturity is not in the partition)

@maturity_props	Maturity proportion block command
Conditions	Only used in a model where maturity is not a partition character.
Effects	Defines any following commands as @maturity_props subcommands
all	Maturity proportions by age/size class
Command	maturity_props
Conditions	Only used in a model where maturity is not a partition character. Alternatively use male/female.
Type	ogive
Effects	Defines the proportions mature by age/size class.
male, female	Maturity proportions by sex and age/size class
Command	maturity_props
Conditions	Only used in a sexed model where maturity is not a partition character. Alternatively use all.
Type	2 x ogive
Effects	Defines the proportions mature by sex and age/size class.

7.9 Defining migrations

@migration	Migration block command
Conditions	Only used in a multi-area model.
Label	the name of a migration
Effects	Defines any following commands as @migration subcommands.
stock	Stock that migrates
Command	migration[label]
Type	String
Default	(All stocks)
Effects	If supplied, defines the single stock that migrates. If not supplied, all stocks migrate.
migrators	Whether mature, immature, or both kinds of fish migrate
Command	migration[label]
Type	String
Default	all
Effects	Defines the fish that can migrate, either mature, immature, or all.
prop	Proportion of applicable fish that migrate
Command	migration[label]
Conditions	Use one of prop, rates_all, or rates_male and rates_female.
Type	Estimable
Default	1, unless one of the other rates parameters is defined.
Effects	Defines the proportion of applicable fish (of the right stock and maturity status) which migrate.
rates_all	Proportion of applicable fish that migrate, by age/size class
Command	migration[label]
Conditions	Use one of prop, rates_all, or rates_male and rates_female.
Type	ogive
Effects	Defines the proportion of applicable fish (of the right stock and maturity status) which migrate, by age/size class.
rates_male, rates_female	Proportion of applicable fish that migrate, by sex and age/size class
Command	migration[label]
Conditions	Only used in a sexed model. Use one of prop, rates_all, or rates_male and rates_female.
Type	2 x ogive
Effects	Defines the proportion of applicable fish (of the right stock and maturity status) which migrate, by sex and age/size class.
density_dependent	Are the migration rates density-dependent?
Command	migration[label]
Type	Switch
Default	False (no density dependence)
Effects	Defines the migration rates as density-dependent.
S, D	Source and destination density-dependence parameters
Command	migration[label]
Conditions	Both these parameters must be provided if density_dependent is set, and are ignored otherwise.
Type	2 x estimable
Effects	Define the dependence on source and destination abundance.
Notes	Typically $S \leq 0, D \geq 0$. A zero value means no dependence.
wave	Is this one wave of a 2-wave migration. If so, is it the first or second?
Command	migration[label]
Conditions	There should be a matching migration for the other wave.
Type	Integer
Default	Not a 2-wave migration.

Effects	If 1, defines this migration event as the first wave of a 2-wave migration. If 2, defines this migration event as the second wave of a 2-wave migration.
pwave	The proportion of fish in the first wave of a 2-wave migration
Command	<code>migration[label]</code>
Conditions	Only used in a multi-area model. There should be a matching migration for the other wave.
Type	Estimable
Default	Not a 2-wave migration.
Effects	Defines the proportion of fish in the first wave of the 2-wave migration this is part of.
Notes	If you are estimating this parameter, you need to use the same subcommand in your <code>estimation.csl</code> file, to ensure it takes the same value for both waves of the migration.

7.10 Defining natural mortality

@natural_mortality	Natural mortality block command
Label	The name of the stock (if there is more than one <code>@natural_mortality</code> command block, i.e., natural mortality is applied to different stock differently).
Effects	Defines any following commands as <code>@natural_mortality</code> subcommands.
all	The overall natural mortality rate
Command	<code>natural_mortality</code>
Conditions	Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	Estimable
Effects	Defines the natural mortality rate.
male, female	The male and female natural mortality rates
Command	<code>natural_mortality</code>
Conditions	Only used in a sex-based model. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	2 x estimable
Effects	Defines the natural mortality rate by sex.
avg, diff	The male/female average and male-female difference in natural mortality rates
Command	<code>natural_mortality</code>
Conditions	Only used in a sex-based model. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	2 x estimable
Effects	Defines the natural mortality rate by sex. The male rate is $(avg+diff/2)$. The female rate is $(avg-diff/2)$.
mature, immature	The mature and immature natural mortality rates
Command	<code>natural_mortality</code>
Conditions	Only used in a model with maturity in the partition. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and

	ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	2 x estimable
Effects	Defines the natural mortality rate by maturity.
male_mature, etc. Natural mortality rates by sex and maturity	
Command	natural_mortality
Conditions	Only used in a model with sex and maturity in the partition. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	4 x estimable
Effects	Defines the natural mortality rate by sex and maturity.
ogive_all The overall natural mortality rate as an ogive	
Command	natural_mortality
Conditions	Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	ogive
Effects	Defines the natural mortality rate by age/size class.
ogive_male, ogive_female The male and female natural mortality rates as ogives	
Command	natural_mortality
Conditions	Only used in a sex-based model. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	2 x ogive
Effects	Defines the natural mortality rate by sex and age/size class.
ogive_avg, ogive_diff The male/female average and male-female difference in natural mortality rates, as ogives	
Command	natural_mortality
Conditions	Only used in a sex-based model. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	2 x ogive
Effects	Defines the natural mortality rate by sex and age/size class. For each age/size class, the male rate is $(avg+diff/2)$ and the female rate is $(avg-diff/2)$.
ogive_mature, ogive_immature The mature and immature natural mortality rates, as ogives	
Command	natural_mortality
Conditions	Only used in a model with maturity in the partition. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	2 x ogive
Effects	Defines the natural mortality rate by maturity and age/size class.

ogive_male_mature, etc. Natural mortality rates by sex and maturity, as ogives

Command	natural_mortality
Conditions	Only used in a model with sex and maturity in the partition. Use one of the following: all, male and female, avg and diff, mature and immature, male_mature and etc., ogive_all, ogive_male and ogive_female, ogive_avg and ogive_diff, ogive_mature and ogive_immature, ogive_male_mature, etc.
Type	4 x ogive
Effects	Defines the natural mortality rate by sex, maturity, and age/size class.

7.11 Defining fishing mortality

@fishery

Fishery block command

Label	the name of a fishery
Effects	Defines any following commands as @fishery subcommands

catches

Catches by year

Command	fishery[fishery_name]
Type	Constant vector
Effects	Defines the catch for the fishery by year
Notes	Entries should correspond to the years given in <i>years</i> .

years

Years for which catches are provided

Command	fishery[fishery_name]
Type	Constant vector
Effects	Defines the years for which the catches are provided
Notes	For each entry of <i>catches</i> there should be a corresponding year in <i>years</i> . The years should be consecutive, and in the range <i>initial</i> to <i>current</i> . If the years start after <i>initial</i> , then all catches before the first year supplied are taken to be 0.

selectivity

Name of the selectivity to use

Command	fishery[fishery_name]
Type	String
Effects	Defines the label of the selectivity to use with this fishery, which should be an entry of <i>selectivity_names</i> .

F_max

Maximum fishing pressure (Baranov mortality)

Command	fishery[fishery_name]
Conditions	Only used if the Baranov equation is applied.
Type	Constant
Effects	Defines the maximum possible fishing pressure F_{max} .

U_max

Maximum fishing pressure (instantaneous mortality)

Command	fishery[fishery_name]
Conditions	Only used if the fishing mortality is applied instantaneously.
Type	Constant
Effects	Defines the maximum possible fishing pressure U_{max} .

Fs

Instantaneous mortality F by year

Command	fishery[fishery_name]
Condition	Only usable with the Baranov catch equation.
Type	Constant vector
Effects	Defines the F for the fishery by year (for years in which catches were not available)
Notes	Entries should correspond to the years given in <i>Fs_years</i> . Be clear about the meaning of F , it is an instantaneous mortality and is multiplied by the selectivity before it is applied.

Fs_years	Years for which <i>F</i>'s are provided
Command	fishery[fishery_name]
Type	Constant vector
Effects	Defines the years for which instantaneous mortalities <i>F</i> are provided
Notes	For each entry of <i>Fs</i> there should be a corresponding year in <i>Fs_years</i> . The years should be consecutive, in the range <i>initial</i> to <i>current</i> , and non-overlapping with <i>years</i> (which is the range of years for which catches are provided).
future_catches	Catches by year in the projection period
Command	fishery[fishery_name]
Type	Constant vector
Effects	Defines the catch for the fishery by year, in the projection period. Ignored unless you are doing projection.
Values	Entries should correspond to the years given in <i>future_years</i> .
future_years	Years for which catches are provided
Command	fishery[fishery_name]
Type	Constant vector
Effects	Defines the years for which the catches are provided, in the projection period. Ignored unless you are doing projection.
Notes	For each entry of <i>future_catches</i> there should be a corresponding year in <i>future_years</i> . The years should be consecutive, and in the range <i>current</i> +1 to <i>final</i> .

7.12 Defining disease mortality

@disease_mortality	Disease mortality block command
Effects	Defines any following commands as <i>@disease_mortality</i> subcommands.
DM	Disease mortality rate
Command	disease_mortality
Conditions	Must be supplied if <i>@annual cycle.disease_mortality_time</i> is specified
Type	Estimable
Effects	Defines the disease mortality rate.
selectivity	The selectivity ogive
Command	disease_mortality
Conditions	Must be supplied if <i>@annual cycle.disease_mortality_time</i> is specified
Type	ogive
Effects	Defines the selectivity by age/size class (S_{ij}), i.e., the disease mortality to apply to each age/size class in year k is $index[k] \times DM \times S_{ij}$
years	Years to apply the disease mortality
Command	disease_mortality
Conditions	Must be supplied if <i>@annual cycle.disease_mortality_time</i> is specified
Type	Constant vector
Effects	Defines the years in which to apply the disease mortality.
index	Relative value of the disease mortality by year
Command	disease_mortality
Conditions	Must be supplied if <i>@annual cycle.disease_mortality_time</i> is specified
Type	Estimable vector
Effects	Defines the relative value of DM to apply for each year in <i>@disease_mortality.years</i> , i.e., the disease mortality to apply to each age/size class in year k is $index[k] \times DM \times S_{ij}$

7.13 Defining selectivities

@selectivity_names	List of selectivity names
Type	Vector of strings
Effects	Lists the labels of all the selectivities in the model.
@selectivity	Selectivity block command
Label	the name of a selectivity
Effects	Defines any following commands as @selectivity subcommands
all	The selectivity ogive
Command	selectivity [name]
Conditions	Use one of the following: all, male and female, mature and immature, male_mature, etc.
Type	ogive
Effects	Defines the selectivity by age/size class.
male, female	The selectivity ogives by sex
Command	selectivity [name]
Conditions	Only used in a sex-based model. Use one of the following: all, male and female, mature and immature, male_mature, etc.
Type	2 x ogive
Effects	Defines the selectivity by sex and age/size class.
mature, immature	The selectivity ogives by maturity
Command	selectivity [name]
Conditions	Only used in a model with maturity in the partition. Use one of the following: all, male and female, mature and immature, male_mature, etc.
Type	2 x ogive
Effects	Defines the selectivity by maturity and age/size class.
male_mature, etc.	The selectivity ogives by sex and maturity
Command	selectivity [name]
Conditions	Only used in a model with sex and maturity in the partition. Use one of the following: all, male and female, mature and immature, male_mature, etc.
Type	4 x ogive
Effects	Defines the selectivity by sex, maturity, and age/size class.
shift_E	Exogenous selectivity shift variable E
Command	selectivity [name]
Type	Constant vector
Effects	Defines the value of the exogenous variable used to shift the fishery selectivity.
Notes	Entries should correspond to the years given in shift_years.
shift_years	Years for which exogenous selectivity shift variable E is provided
Command	selectivity [name]
Type	Constant vector
Effects	Defines the years for which the shift variable <i>E</i> is provided.
Notes	For each entry of shift_E there should be a corresponding year in shift_years. If there is no entry, no shift is carried out for that year.
shift_a	Exogenous selectivity shift parameter a
Command	selectivity [name]
Type	Estimable
Effects	Defines the value of the selectivity shift parameter

7.14 Setting the initial state

@n_equilibrium	Number of years of running the equilibrium model
Conditions	Only used for a size based model
Type	Integer
Effects	Defines the number of years that the equilibrium model is run when setting the initial state
Notes	Try some different values, if it makes a difference then you are probably too low.
@Rinitial_is_deviate	Is $R_{initial}$ supplied relative to R_0?
Conditions	Not used unless $R_{initial}$ is supplied.
Type	Switch
Default	False ($R_{initial}$ is supplied as an absolute number)
Effects	Defines if $R_{initial}$ is supplied relative to R_0 , rather than as an absolute number
@initialization	Initialization block command
Label	the name of a stock
Effects	Defines any following commands as @initialization subcommands for the stock
Notes	Omit the stock label if there is only one stock in the model.
B0	Equilibrium abundance B_0
Command	initialization[stock_name]
Conditions	Define either B0 for each stock or R0 for each stock or, in a two-stock model, ((B0_total or log_B0_total) and B0_prop_stock1) OR ((R0_total or log_R0_total) and R0_prop_stock1)
Type	Estimable
Effects	Defines the value of B_0
Notes	Not to be used if use_mean_YCS is true
R0	Equilibrium recruitment R_0
Command	initialization[stock_name]
Conditions	Define either B0 for each stock or R0 for each stock or, in a two-stock model, ((B0_total or log_B0_total) and B0_prop_stock1) OR ((R0_total or log_R0_total) and R0_prop_stock1)
Type	Estimable
Effects	Defines the value of R_0
Notes	Not to be used if use_mean_YCS is true
Bmean	Equilibrium abundance B_{mean} corresponding to R_{mean}
Command	initialization[stock_name]
Conditions	Define either B_{mean} for each stock or R_{mean} for each stock or, in a two-stock model, ((Bmean_total or log_Bmean_total) and Bmean_prop_stock1) OR ((Rmean_total or log_Rmean_total) and Rmean_prop_stock1)
Type	Estimable
Effects	Defines the value of B_{mean}
Notes	Can only be used when use_mean_YCS is true
Rmean	Expected recruitment in any year, R_{mean}
Command	initialization[stock_name]
Conditions	Define either Bmean for each stock or Rmean for each stock or, in a two-stock model, ((Bmean_total or log_Bmean_total) and Bmean_prop_stock1) OR ((Rmean_total or log_Rmean_total) and Rmean_prop_stock1)
Type	Estimable

Effects	Defines the value of R_{mean}
Notes	Can only be used when <code>use_mean_YCS</code> is true
Binitial	Initial abundance $B_{initial}$
Command	<code>initialization[stock_name]</code>
Conditions	Either provide <code>Binitial</code> for each stock, or <code>Rinitial</code> for each stock, or <code>Cinitial</code> for each stock, or <code>Cinitial_male</code> and <code>Cinitial_female</code> for each stock, or none of them.
Type	Estimable
Effects	Defines the value of $B_{initial}$
Rinitial	Initial recruitment $R_{initial}$
Command	<code>initialization[stock_name]</code>
Conditions	Either provide <code>Binitial</code> for each stock, or <code>Rinitial</code> for each stock, or <code>Cinitial</code> for each stock, or <code>Cinitial_male</code> and <code>Cinitial_female</code> for each stock, or none of them.
Type	Estimable
Effects	Defines the value of $R_{initial}$
Cinitial	Initial number in each age/size class $C_{initial}$
Command	<code>initialization[stock_name]</code>
Conditions	Either provide <code>Binitial</code> for each stock, or <code>Rinitial</code> for each stock, or <code>Cinitial</code> for each stock, or <code>Cinitial_male</code> and <code>Cinitial_female</code> for each stock, or none of them.
Type	ogive
Effects	Defines the values of $C_{initial, i}$, i.e., the initial number in each age or size class i .
Notes	You almost certainly want to use an <code>allvalues</code> ogive, as other options make little sense. The $C_{initial}$ value you set for the first age/size class is related to the ordering of ageing and recruitment in your annual cycle (see Section 4.5).
Cinitial_male, Cinitial_female	Initial number in each age/size class $C_{initial}$ for each sex
Command	<code>initialization[stock_name]</code>
Conditions	Either provide <code>Binitial</code> for each stock, or <code>Rinitial</code> for each stock, or <code>Cinitial</code> for each stock, or <code>Cinitial_male</code> and <code>Cinitial_female</code> for each stock, or none of them.
Type	2 x ogive
Effects	Defines the values of $C_{initial, i}$, i.e., the initial number in each age or size class i , for each sex.
Notes	You almost certainly want to use <code>allvalues</code> ogives as other options make little sense. The $C_{initial}$ value you set for the first age/size class is related to the ordering of ageing and recruitment in your annual cycle (see Section 4.5).
@B0_total, @log_B0_total, @R0_total, @log_R0_total	
Conditions	In a two-stock model only, you can use one of these parameters instead of supplying <code>B0</code> or <code>R0</code> . You must supply the appropriate one of <code>B0_prop_stock1</code> or <code>R0_prop_stock1</code> .
Type	Estimable
Effects	Defines the value of B_0 , $\log(B_0)$, R_0 , or $\log(R_0)$, summed across the two stocks.
@B0_prop_stock1, @R0_prop_stock1	
Conditions	In a two-stock model only, you can use one of these parameters instead of supplying <code>B0</code> or <code>R0</code> . You must supply one of <code>B0_total</code> , <code>log_B0_total</code> , <code>R0_total</code> , <code>log_R0_total</code> .
Type	Estimable
Effects	Defines the proportion of the equilibrium abundance or recruitment which is of the first stock.

7.15 Defining ogive preferences

@n_quant	Number of points at which to evaluate size-based ogives in an age-based model
Conditions	Only used in an age-based model which uses size-based ogives.
Type	Integer
Default	5
Effect	Defines the number of points used in the approximation to the integral of the ogive over the distribution of sizes at age.
Notes	The default should normally be adequate, unless you have ogives which change very steeply.

7.16 Defining size-at-age

@size_at_age_type	Size-at-age model type
Conditions	Only used in an age-based model.
Type	String
Effects	Defines the size-at-age model used. So far there is <code>von_Bert</code> and <code>Schnute</code> , for when growth curves are used, and <code>data</code> , for when size-at-age data are provided for one or more years.
Notes	<code>data</code> is only an option if you have not specified <code>annual_cycle.growth_props</code> .
@size_at_age_years	Years for which mean-size-at-age data are provided
Conditions	Only used in an age-based model with <code>size_at_age_type=data</code> .
Type	Constant vector
Effects	Defines the list of years for which mean-size-at-age data are provided.
Notes	Need not be consecutive. Can include years in the projection period.
@size_at_age_dist	Distribution of sizes-at-age around the mean
Conditions	Only used in an age-based model.
Type	String
Default	No variation of sizes-at-age around the mean
Effects	Defines the distribution of sizes-at-age around the mean. Can be <code>normal</code> or <code>lognormal</code> . If you don't choose either, there is no variation of size-at-age around the mean.
Notes	Also used for the likelihood of age-size observations (see Section 5.6). <code>c.v.s</code> are specified within the <code>@size_at_age</code> blocks.
@size_at_age	Size-at-age block command
Label	the name of a stock (optional)
Conditions	Only used in an age-based model.
Effects	Defines any following commands as <code>@size_at_age</code> subcommands for the stock
Notes	Omit the stock label if there is only one stock in the model, or if all the stocks have the same size-at-age.
k, t0, Linf	von Bertalanffy parameters
Command	<code>size_at_age[stock_name]</code>
Conditions	Only used in an age-based model with <code>size_at_age_type=von_Bert</code> . Use either <code>k, t0, Linf</code> or <code>k_male, t0_male, Linf_male, k_female, t0_female, Linf_female</code> .
Type	3 x estimable vector (<i>NOT</i> estimable constant)
Effects	Defines the von Bertalanffy parameters. If this is not a growth-path model, each parameter is a 1-vector. If this is a growth-path model, each parameter can either be a 1-vector holding the common value for all paths, or a vector with element <i>i</i> holding the value for path <i>i</i> .

Warning Make sure that you have specified the relationship in units that are compatible with the catches and the size-weight relationship (see section 4.8 and 4.9)

k_male, t0_male, Linf_male, k_female, t0_female, Linf_female von Bertalanffy parameters by sex

Command `size_at_age[stock_name]`
 Conditions Only used in an age-based model with `size_at_age_type=von_Bert`. Use either `k`, `t0`, `Linf` or `k_male`, `t0_male`, `Linf_male`, `k_female`, `t0_female`, `Linf_female`.
 Type 6 x estimable vector (*NOT* estimable constant)
 Effects Defines the von Bertalanffy parameters by sex. If this is not a growth-path model, each parameter is a 1-vector. If this is a growth-path model, each parameter can either be a 1-vector holding the common value for all paths, or a vector with element *i* holding the value for path *i*.
 Warning Make sure that you have specified the relationship in units that are compatible with the catches and the size-weight relationship (see section 4.8 and 4.9)

y1, y2, tau1, tau2, a, b Schnute parameters

Command `size_at_age[stock_name]`
 Conditions Only used in an age-based model with `size_at_age_type=Schnute`. Alternatively use the versions suffixed `_male`, `_female`.
 Type 6 x estimable vector (*NOT* estimable constant)
 Effects Defines the Schnute parameters. If this is not a growth-path model, each parameter is a 1-vector. If this is a growth-path model, each parameter can either be a 1-vector holding the common value for all paths, or a vector with element *i* holding the value for path *i*.
 Warning Make sure that you have specified the relationship in units that are compatible with the catches and the size-weight relationship (see section 4.8 and 4.9)

y1_male, y2_male, tau1_male, tau2_male, a_male, b_male, y1_female, y2_female, tau1_female, tau2_female, a_female, b_female Schnute parameters by sex

Command `size_at_age[stock_name]`
 Conditions Only used in an sex/age model with `size_at_age_type=Schnute`. Alternatively use the versions without suffixes `_male`, `_female`.
 Type 12 x estimable vector (*NOT* estimable constant)
 Effects Defines the Schnute parameters by sex. If this is not a growth-path model, each parameter is a 1-vector. If this is a growth-path model, each parameter can either be a 1-vector holding the common value for all paths, or a vector with element *i* holding the value for path *i*.
 Warning Make sure that you have specified the relationship in units that are compatible with the catches and the size-weight relationship (see section 4.8 and 4.9)

male_[year], female_[year] Mean-size-at-age of male and female fish in [year]

Command `size_at_age[stock_name]`
 Conditions Only used with `size_at_age_type=data` in a sexed model. Alternatively use the unsexed version `all_[year]`.
 Type 2 x constant vector
 Effects Defines the mean-size-at-age data for males and females in [year]. There should be both male and female data for each year in `size_at_age_years`, and each set of data should have one entry for each age class.
 Example If you have sexed mean-size-at-age data for 1997, 1998, and 1999, and `min_age=2`, `max_age=6`, you would put, e.g.,

```
@size_at_age_type data
@size_at_age_years 1997 1998 1999
@size_at_age
# age 2 3 4 5 6+
male_1997 20 40 50 55 58
female_1997 40 80 100 110 130
male_1998 18 36 48 52 55
```

	... (three more rows)
Warning	Make sure that you have specified the relationship in units that are compatible with the catches and the size-weight relationship (see section 4.8 and 4.9)
all_[year]	Mean-size-at-age of fish of both sexes in [year]
Command	size_at_age[stock_name]
Conditions	Only used with size_at_age_type=data. Alternatively use the sexed versions male_[year], female_[year].
Type	Constant vector
Effects	Defines the mean-size-at-age data in [year], applied to both sexes. There should be data for each year in size_at_age_years, and each set of data should have one entry for each age class.
Example	If you have mean-size-at-age data for 1997, 1998, and 1999, and min_age=2, max_age=6, you would put, e.g., <pre>@size_at_age_type data @size_at_age_years 1997 1998 1999 @size_at_age # age 2 3 4 5 6+ all_1997 20 40 50 55 58 all_1998 23 45 52 58 63 all_1999 18 36 48 52 55</pre>
Warning	Make sure that you have specified the relationship in units that are compatible with the catches and the size-weight relationship (see section 4.8 and 4.9)
cv	c.v. of sizes-at-age around the mean
Command	size_at_age[stock_name]
Conditions	Only used in an age-based model where size_at_age_dist has been supplied. Alternatively use the versions suffixed _male, _female.
Type	Estimable vector (<i>NOT</i> estimable constant)
Effects	Defines the c.v. of sizes around the mean. If this is not a growth-path model, the parameter is a 1-vector. If this is a growth-path model, the parameter can either be a 1-vector holding the common value for all paths, or a vector with element <i>i</i> holding the value for path <i>i</i> .
Notes	Also used for the likelihood of age-size observations (see Section 5.6).
cv_male, cv_female	c.v. of sizes-at-age around the mean, by sex
Command	size_at_age[stock_name]
Conditions	Only used in an sexed, age-based model where size_at_age_dist has been supplied. Alternatively use cv without suffixes.
Type	2 x estimable vector (<i>NOT</i> estimable constant)
Effects	Defines the c.v. of sizes around the mean, by sex. If this is not a growth-path model, each parameter is a 1-vector. If this is a growth-path model, each parameter can either be a 1-vector holding the common value for all paths, or a vector with element <i>i</i> holding the value for path <i>i</i> .
Notes	Also used for the likelihood of age-size observations (see Section 5.6).
@annual_growth	Use annual growth variation. Amount of an average year's growth that occurs in each year
Conditions	Only used in an age-based model with size_at_age_type not equal to data.
Type	Estimable vector
Default	No annual growth variation.
Effects	Defines that annual growth variation is to be used, and defines the proportion of average annual growth that occurs in each year.
Notes	Entries should correspond to the years given in annual_growth_years. The first entry must be 1, i.e., a year of ordinary growth. If this is not the case, prefix one more year onto annual_growth_years. Warning, watch out for off-by-one errors.

@annual_growth_years Years for which annual growths are provided

Conditions	Only used in an age-based model, where <code>annual_growths</code> is set.
Type	Constant vector
Effects	Defines the years for which the annual growth increments are provided
Notes	For each entry of <code>annual_growths</code> there should be a corresponding year in <code>annual_growth_years</code> . The years should be consecutive, and contained in the range <code>initial</code> to <code>current</code> . All annual growth increments not supplied are taken to be 1.

7.17 Defining the size-weight relationship

@size_weight Size-weight block command

Label	the name of a stock (optional)
Effects	Defines any following commands as <code>@size-weight</code> subcommands for the stock
Notes	Omit the stock label if there is only one stock in the model, or if all the stocks have the same size-weight.

type The size-weight relationship function

Command	<code>size_weight [stock_name]</code>
Type	String
Default	<code>basic</code>
Effects	Defines the type of size-weight relationship used in the model. The only option available is <code>basic</code> .

a, b The size-weight parameters *a* and *b*

Command	<code>size_weight [stock_name]</code>
Conditions	Only used with the <code>basic</code> size-weight relationship. Alternatively use the versions suffixed <code>_male</code> , <code>_female</code> .
Type	2 x constant
Effects	Defines the <i>a</i> and <i>b</i> parameters of the size-weight relationship
Notes	If you provide your catch in tonnes, and your growth curve in centimetres, then <i>a</i> should be on the right scale to convert a length in centimetres to a weight in tonnes.
Warning	Make sure that you have specified the relationship in units that are compatible with the catches and the growth parameters (see <code>verify_size_weight</code> and section 4.8 and 4.9)

a_male, b_male, a_female, b_female The size-weight parameters *a* and *b*, by sex

Command	<code>size_weight [stock_name]</code>
Conditions	Only used with the <code>basic</code> size-weight relationship. Alternatively use the un-suffixed versions.
Type	4 x constant
Effects	Defines the <i>a</i> and <i>b</i> parameters of the size-weight relationship, by sex
Notes	If you provide your catch in tonnes, and your growth curve in centimetres, then <i>a</i> should be on the right scale to convert a length in centimetres to a weight in tonnes.
Warning	Make sure that you have specified the relationship in units that are compatible with the catches and the growth parameters (see <code>verify_size_weight</code> and section 4.8 and 4.9)

verify_size_weight Verify the supplied size-weight relationship and units

Command	<code>size_weight [stock_name]</code>
Conditions	Only used with the <code>basic</code> size-weight relationship.
Type	Constant vector defining the three parameters (1) fish length (in cm), (2) lower bound on weight at this length (in kg), and (3) upper bound on weight at this length (in kg).

Default	The weight of a fish that is 25 cm is calculated and compared against a lower bound of 0.05 kg and an upper bound of 5 kg. If the weight is outside these bounds then a warning message is given.
Effects	If the fish weight at the given fish length is outside the bounds then the program halts with an error message. Otherwise this fish weight and length are given as part of the output.
Notes	The calculation of the fish weight is done under the assumption that the catch is in tonnes, and the growth curve in centimetres. (see section 4.8 and 4.9)

@weightless_model	Is this a model which does not involve fish weight?
Type	Switch
Default	False (model involves fish weight)
Effects	Defines all fish to have a nominal weight of 1.
Notes	As a result, CASAL interprets your catches as numbers of fish rather than tonnes of fish. "Biomass", including SSB and abundance observations, now also represents numbers of fish.

8. THE ESTIMATION.CSL FILE

The estimation parameters are specified in the `estimation.csl` file. See Section 5 for information about the estimation section, and Section 2.4 for instructions on writing a CASAL data file.

8.1 Defining the estimation method

@estimator	Choice of estimation method
Type	String
Effects	Defines the estimation method as Bayes, likelihood, or least_squares.
@weighting	Choice of least-squares weighting
Condition	@estimator=least_squares
Type	String
Effects	Defines the least-squares weighting as Cordue or none.
@k	Robustifying constant for least-squares weighting
Condition	@estimator=least_squares. Use either k, or ko and kp.
Type	Constant
Effects	Defines both the least-squares robustifying constants k_o and k_p .
@ko, @kp	Robustifying constants for least-squares weighting
Condition	@estimator=least_squares. Use either k, or ko and kp.
Type	2 x constant
Effects	Defines the least-squares robustifying constants k_o and k_p .

8.2 Defining point estimation

@max_iters	Maximum number of iterations in the minimiser
Condition	Only used in point estimation.
Type	Integer
Default	300
Effects	Defines the maximum number of quasi-Newton iterations allowed in a minimization.
@max_evals	Maximum number of evaluations in the minimiser
Condition	Only used in point estimation.
Type	Integer
Default	1000
Effects	Defines the maximum number of objective function evaluations allowed in a minimization.
@max_iters_intermediate	Maximum number of iterations in early phases
Condition	Only used in multi-phase point estimation.
Type	Integer
Default	max_iters
Effects	Defines the maximum number of quasi-Newton iterations allowed in all but the last phase of a multi-phase minimization.
@max_evals_intermediate	Maximum number of evaluations in early phases
Condition	Only used in multi-phase point estimation.
Type	Integer
Default	max_evals

Effects	Defines the maximum number of objective function evaluations allowed in all but the last phase of a multi-phase minimization.
@grad_tol	Minimiser convergence threshold
Condition	Only used in point estimation.
Type	Constant
Default	0.002
Effects	Defines the convergence criterion for minimization. The minimiser converges successfully if the maximum absolute gradient, of the objective function with regard to the transformed free parameters, divided by the absolute value of the objective function, is less than <code>grad_tol</code> . In other words, make the convergence criterion more severe by decreasing <code>grad_tol</code> .
Notes	The minimiser also converges and claims it is successful if the quasi-Newton stepsize becomes very small. We are not sure what the implications of this result are yet. If in doubt, act as if the minimiser had not converged — do more runs from different starting points.

8.3 Defining likelihood or posterior profiling

@profile	Profile block command
Conditions	Only used in likelihood or posterior profiling.
Effects	Defines any following commands as <code>@profile</code> subcommands
Notes	The <i>i</i> th <code>@profile</code> block relates to the <i>i</i> th parameter to be profiled.
parameter	Name of the parameter to be profiled
Command	<code>profile [i]</code>
Conditions	Only used in likelihood or posterior profiling.
Type	String
Effects	Defines the name of the <i>i</i> th parameter to be profiled
Notes	Only scalar parameters can be profiled.
n	Number of values at which to profile the parameter
Command	<code>profile [i]</code>
Conditions	Only used in likelihood or posterior profiling.
Type	Integer
Default	10
Effects	Defines the number of values at which to profile the parameter. See <code>l</code> and <code>u</code> .
l, u	Range of values at which to profile the parameter
Command	<code>profile [i]</code>
Conditions	Only used in likelihood or posterior profiling.
Type	$2 \times \text{constant}$
Default	$l = \text{lower bound on parameter} + (\text{range of parameter} / 2n)$ $u = \text{upper bound on parameter} - (\text{range of parameter} / 2n)$
Effects	Defines the lower and upper values at which to profile the parameter. See <code>n</code> .

8.4 Defining MCMC

@MCMC	MCMC block command
Conditions	Only used in MCMC.
Effects	Defines any following commands as <code>@MCMC</code> subcommands.
Notes	Some of these are only used when running the chain (<code>casal -m, -a</code>). The rest are only used when creating a sub-sample (<code>casal -C</code>).
start	Covariance multiplier for the starting point of the Markov chain
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m</code>).
Type	Constant

Effects	If 0, defines the starting point of the chain as the point estimate. If >0, defines the starting point as randomly generated, with covariance matrix equal to the approximate covariance (inverse Hessian) times the value of this <code>start</code> parameter.
Notes	This parameter can be overridden by the <code>casal -i</code> switch. If the file specified with <code>-i</code> contains two parameter vectors, the second is used to start the chain (the first is used to start the initial point estimate).
length	Length of the Markov chain
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>).
Type	Integer
Effects	Defines the length of the Markov chain (as a number of iterations)
Notes	(unless you stop the chain first, of course) We recommend an absolute minimum of 100 000 for serious runs. With many parameters, 1 000 000 may not be enough.
keep	Spacing between recorded values in the chain
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>).
Type	Integer
Default	1 (i.e., all values are recorded)
Effects	Defines the spacing between recorded values in the chain. Samples from the posterior are written to file only if their sample number is evenly divisible by <code>keep</code> .
max_cor	Maximum absolute correlation in the covariance matrix of the proposal distribution
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>).
Type	Constant
Default	0.8
Effects	Defines the maximum correlation in the covariance matrix of the proposal distribution. Correlations greater than <code>max_cor</code> are decreased to <code>max_cor</code> , and those less than <code>-max_cor</code> are increased to <code>-max_cor</code> .
stepsize	Initial stepsize (as a multiplier of the approximate covariance matrix)
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>).
Type	Constant
Default	$2.4d^{-0.5}$ where d is the number of free parameters.
Effects	Defines the stepsize in the Markov chain.
Notes	The covariance of the proposal distribution is the approximate covariance (inverse Hessian) times this <code>stepsize</code> parameter. See also <code>adaptive_stepsize</code> , <code>adapt_at</code> .
adaptive_stepsize	Should the MCMC stepsize be altered during the chain?
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>).
Type	Switch
Default	False (do not alter stepsize)
Effects	Defines whether the stepsize should be altered adaptively during the chain.
adapt_at	At which iteration numbers can the MCMC stepsize be altered?
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>). Only used if <code>adaptive_stepsize</code> is set.
Type	Constant vector
Effects	Defines the iteration number(s) at which the stepsize is altered adaptively.
Notes	Make sure that the burn-in period is greater than the largest entry of <code>adapt_at</code> .

proposal_t	Should the proposal distribution be multivariate t?
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>).
Type	Switch
Default	False (use multivariate normal)
Effects	Defines whether the proposal distribution should be multivariate t rather than multivariate normal.
df	Degrees of freedom of the multivariate t proposal distribution.
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>). Only used if <code>proposal_t</code> is set.
Type	Integer
Default	4
Effects	Defines the degrees of freedom of the multivariate t proposal distribution.
burn_in	Number of samples to be discarded for the burn-in period
Command	MCMC
Conditions	Only used when creating a posterior sub-sample (<code>casal -C</code>).
Type	Integer
Effects	Defines the number of samples to be discarded at the start of each chain.
Notes	This is the number of recorded samples to be discarded. So, the length of the burn-in period is effectively <code>burn_in × keep</code> .
subsample_size	Size of the sub-sample to be generated
Command	MCMC
Conditions	Only used when creating a posterior sub-sample (<code>casal -C</code>).
Type	Integer
Default	No random sub-sampling is done.
Effects	Defines the size of the sub-sample to be generated using resampling with replacement.
Notes	This is used to decimate down to a sub-sample of manageable size.
systematic	Should sub-sampling be systematic?
Command	MCMC
Conditions	Only used when creating a posterior sub-sample (<code>casal -C</code>). <code>subsample_size</code> must be provided. Cannot be used with <code>prior_reweighting</code> .
Type	Switch
Default	False (sub-sample randomly)
Effects	Defines the sub-sampling from the posterior as systematic (i.e., keep every n th point) rather than random.
prior_reweighting	Should the sub-sample be generated using prior reweighting?
Command	MCMC
Conditions	Only used when creating a posterior sub-sample (<code>casal -C</code>). <code>subsample_size</code> must be provided. <code>systematic</code> must not be set.
Type	Switch
Default	False (no prior reweighting)
Effects	Defines that prior reweighting should be carried out, using the current prior.
Notes	The prior commands in <code>estimation.csl</code> should have changed since you did the MCMC run. The sub-sampling is weighted by the ratio of the old prior (which is saved in the <code>objective</code> files) to the new prior.
adaptive_covariance	Should the MCMC covariance matrix be altered during the chain?
Command	MCMC
Conditions	Only used when running a Markov chain (<code>casal -m, -a</code>).
Type	Switch
Default	False (do not alter covariance matrix)
Effects	Defines whether the covariance matrix should be altered adaptively during the chain. See also <code>adaptive_stepsize</code> .

adapt_covariance_at At which iteration numbers can the MCMC covariance matrix be altered?

Command MCMC
 Conditions Only used when running a Markov chain (`casal -m, -a`). Only used if `adaptive_covariance` is set.
 Type Constant vector
 Effects Defines the iteration number(s) at which the covariance matrix is altered adaptively.
 Notes Make sure that the burn-in period is greater than the largest entry of `adapt_covariance_at`.

adaptive_covariance_discard If the MCMC covariance matrix is altered during the chain, how many observations should be discarded from the start of the chain when taking a subsample for estimating the new covariance matrix?

Command MCMC
 Conditions Only used when running a Markov chain (`casal -m, -a`).
 Type integer
 Effects Defines the number of observations discarded from the start of the chain. All remaining observations are systematically subsampled, and the (modified) covariance of the subsample is used as the new covariance matrix.
 Notes It is a fatal error if the chain does not move at least once before `adapt_covariance_discard` observations have occurred.

adaptive_covariance_transitions If the MCMC covariance matrix is altered during the chain, how many transitions must occur in the part of the chain used to estimate the new covariance matrix?

Command MCMC
 Conditions Only used when running a Markov chain (`casal -m, -a`).
 Type integer
 Effects It is a fatal error if the chain does not move at least `adapt_covariance_transitions` times between the end of the discard period (`adapt_covariance_discard`) and the point where the covariance matrix is adapted.

adaptive_covariance_stepsize If the MCMC covariance matrix is altered during the chain, what is the new stepsize value?

Command MCMC
 Conditions Only used when running a Markov chain (`casal -m, -a`).
 Type Constant
 Default $2.4d^{-0.5}$ where d is the number of free parameters.
 Effects Defines the stepsize in the Markov chain, after the covariance matrix is modified adaptively.
 Notes After modification, the covariance of the proposal distribution is the modified covariance from a sample of the chain times this `adaptive_covariance_stepsize` parameter. See also `stepsize`, `adaptive_stepsize`, `adapt_at`.

@trivariate_normal_test Test MCMC algorithm with a simple trivariate normal example

Conditions Only used when running a Markov chain (`casal -m`).
 Type Integer
 Default False (program runs normally)
 Effects Replace the usual objective function with a trivariate normal density. This is a test of the MCMC algorithm, which should generate a sample from the trivariate normal distribution.

8.5 Defining the free parameters and priors

@estimate Free parameter block command

Effects Defines any following commands as `estimate` subcommands
 Notes The i th `@estimate` block relates to the i^{th} parameter to be estimated.

parameter	Name of the parameter to be estimated
Command	<code>estimate [i]</code>
Type	String
Effects	Defines the name of the parameter to be estimated.
Notes	See Section 5.2, and Section 2.4 for instructions on generating the parameter names.
Example	<code>annual_growths initialization[stock_name] .B0 growth[2] .g</code>
same	Names of the other parameters which are constrained to have the same value
Command	<code>estimate [i]</code>
Type	String
Default	(no parameters)
Effects	Defines the names of all the other parameters which are constrained to have the same value as this parameter
Notes	Do not give these parameters separate <code>estimate</code> blocks. See Section 5.2.
phase	Phase at which this parameter should be estimated, in point estimation
Command	<code>estimate [i]</code>
Type	Integer
Default	1
Effects	Defines the phase at which this parameter should be freed.
Notes	If no <code>phase</code> commands are provided, then estimation is single-phase.
lower_bound, upper_bound	Bounds on this scalar parameter
Command	<code>estimate [i]</code>
Type	2 x constant
Effects	Defines the lower and upper bounds on this scalar parameter.
Notes	See also the vector versions below.
lower_bound, upper_bound	Bounds on this vector parameter
Command	<code>estimate [i]</code>
Type	2 x constant vector
Effects	Defines the vectors of lower and upper bounds on this vector parameter.
Notes	See also the scalar versions above.
MCMC_fixed	Should this parameter be fixed during MCMC?
Command	<code>estimate [i]</code>
Conditions	<code>estimator=Bayes</code>
Type	Switch
Default	False (do not fix the parameter during MCMC)
Effects	Define this parameter as fixed during MCMC.
prior	What type of prior does this parameter have?
Command	<code>estimate [i]</code>
Conditions	<code>estimator=Bayes</code>
Type	String
Effects	Defines the type of prior on this parameter. For scalar parameters, <code>uniform</code> , <code>uniform-log</code> , <code>normal</code> , <code>normal-by-stdev</code> , <code>lognormal</code> , <code>normal-log</code> , <code>beta</code> . For vector parameters, all the above plus <code>normal-AR</code> , <code>normal-log-AR</code> , <code>normal-log-mean1-AR</code> .
Notes	No default.

(Parameters of the prior follow)

mu, cv	What are the mean and c.v. of this normal or lognormal prior on a scalar parameter?
Command	<code>estimate [i]</code>
Conditions	<code>estimator=Bayes, prior=normal or lognormal.</code>

Type	2 x constant
Effects	Defines the prior mean and c.v. (<i>before</i> bounds are applied).
mu, stdev	What are the mean and standard deviation of this normal-by-standard deviation or beta prior on a scalar parameter?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal-by-stdev or beta.
Type	2 x constant
Effects	Defines the prior mean and standard deviation (<i>before</i> bounds are applied).
m, s	What are the mean and standard deviation of the log of this scalar parameter, under the normal-log prior?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal-log.
Type	2 x constant
Effects	Defines the prior mean and standard deviation of the log-parameter (<i>before</i> bounds are applied).
A, B	What are the lower and upper values for the range parameters of the beta prior?
Command	estimate [i]
Default	A=0, and B=1
Conditions	estimator=Bayes, prior=beta.
Type	2 x constant
Effects	Defines the lower and upper range values of the beta prior (<i>before</i> bounds are applied).
Warning	Note the bounds must lie <i>inside</i> the range parameters.
mu, cv	What are the mean and c.v. of each element of this normal or lognormal prior on a vector parameter?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal or lognormal.
Type	2 x constant vector
Effects	Defines the prior mean and c.v. of each element of the vector (<i>before</i> bounds are applied).
mu, stdev	What are the mean and standard deviation of each element of this normal-by-standard deviation or beta prior on a vector parameter?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal-by-std.dev or beta.
Type	2 x constant vector
Effects	Defines the prior mean and standard deviation of each element of the vector (<i>before</i> bounds are applied).
m, s	What are the mean and standard deviation of each element of the log of this vector parameter, under the normal-log prior?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal-log.
Type	2 x constant vector
Effects	Defines the prior mean and standard deviation of each element of the log-parameter (<i>before</i> bounds are applied).
A, B	What are the lower and upper values for each element of the range parameters of the beta prior?
Command	estimate [i]
Default	A=0, and B=1
Conditions	estimator=Bayes, prior=beta.
Type	2 x constant
Effects	Defines the lower and upper range values of each element of the beta prior (<i>before</i> bounds are applied).
Warning	Note the bounds must lie <i>inside</i> the range parameters.

mu, cv, rho	What are the mean, c.v., and ρ of this normal-AR prior on a vector parameter?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal-AR.
Type	3 x constant
Effects	Defines the prior mean, c.v., and ρ .
Notes	The single value is used for each element of the parameter.
m, s, r	What are the log-scale mean, standard deviation, and ρ of this normal-log-AR prior on a vector parameter?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal-log-AR.
Type	3 x constant
Effects	Defines the prior mean, standard deviation, and ρ of the log-parameter.
Notes	The single value is used for each element of the parameter.
s, r	What are the log-scale standard deviation and ρ of this normal-log-mean1-AR prior on a vector parameter?
Command	estimate [i]
Conditions	estimator=Bayes, prior=normal-log-mean1-AR
Type	2 x constant
Effects	Defines the prior standard deviation and ρ of the log-parameter.
Notes	The single value is used for each element of the parameter.

8.6 Defining the relativity constants q

@q_method	Method used for relativity constants q
Type	String
Default	“nuisance”
Effects	Defines the method used to deal with q ’s as nuisance or free.
@q	Relativity constant q block command
Label	the label of the q , as used in the relevant @observations block(s)
Effects	Defines any following commands as q subcommands
Notes	Only needed if q_method=free or you have a curvature parameter b
q	Value of the q parameter
Command	q[label]
Conditions	Only used if q_method=free.
Type	Estimable
Effects	Defines the starting value of the q parameter.
b	Curvature parameter b associated with the q
Command	q[label]
Type	Estimable
Default	(no curvature)
Effects	Defines the curvature parameter.
Notes	You also need to set curvature=true in the command block for the relative abundance observations.

8.7 Defining the observations

@abundance	Absolute abundance block command
Effects	Defines any following commands as @abundance subcommands for the time series.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop)

@relative_abundance	Relative abundance block command
Effects	Defines any following commands as @relative_abundance subcommands for the time series.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop)
years	Years of the time series
Command	abundance [label], relative_abundance [label]
Type	Constant vector
Effects	Defines the years for which there are observations. Should be one entry per observation.
step	Time step in which the observations occur
Command	abundance [label], relative_abundance [label]
Type	Integer
Effects	Defines the time step in which the observations occur.
proportion_mortality	Proportion of the step's mortality, after which the observations occur
Command	abundance [label], relative_abundance [label]
Type	Constant
Default	0.5
Effects	Defines the proportion of the mortality in the time step after which the observations occur.
area	Area in which the observations occur
Command	abundance [label], relative_abundance [label]
Condition	n_areas > 1
Type	String
Effects	Defines the area in which the observations occur. Use an area label as per area_names.
q	Relativity constant q to use
Command	relative_abundance [label]
Type	String
Effects	Defines the label of the q used by the observations.
curvature	Should a curvature parameter be used?
Command	relative_abundance [label]
Type	Switch
Default	False (no curvature)
Effects	Defines these observations as using a curvature parameter b .
Notes	See Section 5.7.2. Provide the value of the curvature parameter in the q command block.
biomass	Are the observations biomass rather than numbers of fish?
Command	abundance [label], relative_abundance [label]
Type	Switch
Effects	Defines whether the observations are biomass (biomass=true) or numbers of fish (biomass=false).
ogive	Which selectivity ogive should be applied?
Command	abundance [label], relative_abundance [label]
Type	String
Default	No selectivity
Effects	Defines which selectivity ogive should be applied when calculating the fits. Use a selectivity label as per selectivity_names.
[year]	Abundance for [year]
Command	abundance [label], relative_abundance [label]
Type	Constant
Effects	Defines the abundance for [year].
Example	If you have observations for 1992, 1995, and 1998, you would put, e.g., years 1992 1995 1998

1992 100000
1995 3000
1998 12

(See Section 8.8 for parameters relating to the likelihood or least-squares weights.)

mature_only	Do these observations include mature fish only?
Command	abundance [label]
Type	Switch
Default	False (include both immature and mature fish)
Effects	Defines whether the observations are mature fish only (mature_only=true) or both mature and immature fish (mature_only=false).
Notes	You will probably only want to use this subcommand when generating pseudo-fits (Section 6.2), so as to output mature abundance.
stock	Which stock do these observations relate to?
Command	abundance [label]
Condition	n_stocks > 1
Type	String
Default	(All stocks)
Effects	Defines the name of the stock which is observed. Use a stock label from stock_names.
Notes	You will probably only want to use this subcommand when generating pseudo-fits (Section 6.2), so as to output abundance of a particular stock.
all_areas	Do these observations cover all areas in the model?
Command	abundance [label]
Condition	n_areas > 1
Type	Switch
Default	False (not all areas)
Effects	Defines the observations as covering all areas. The area command is superseded by this command.
Notes	You will probably only want to use this subcommand when generating pseudo-fits (Section 6.2), so as to output abundance over all areas.
@numbers_at	Numbers_at block command
Effects	Defines any following commands as @numbers_at subcommands for the time series.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop)
@relative_numbers_at	Relative numbers-at block command
Effects	Defines any following commands as @relative_numbers_at subcommands for the time series.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop)
@proportions_at	Proportions-at block command
Effects	Defines any following commands as @proportions_at subcommands for the time series.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop). Use catch_at for commercial catch data.
years	Years of the time series
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Type	Constant vector
Effects	Defines the years for which there are observations. Should be one entry per row of observations.

step	Time step in which the observations occur
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Type	Integer
Effects	Defines the time step in which the observations occur.
proportion_mortality	Proportion of the step's mortality after which the observations occur
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Type	Constant
Default	0.5
Effects	Defines the proportion of the mortality in the time step after which the observations occur.
at_size	Are the observations by size?
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Type	Switch
Default	size_based (i.e., size-based in a size-based model, age-based in an age-based model)
Effects	Defines the observations as size-based (at_size=true) or age-based (at_size=false).
sexed	Are the observations sexed?
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Condition	sex_partition is set
Type	Switch
Default	True (observations are sexed)
Effects	Defines the observations as sexed (sexed=true) or unsexed (sexed=false).
area	Area in which the observations occur
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Condition	n_areas > 1
Type	String
Effects	Defines the area in which the observations occur. Use an area label as per area_names.
q	Relativity constant q to use
Command	relative_numbers_at [label]
Type	String
Effects	Defines the label of the q used by the observations.
ogive	Which selectivity ogive should be applied?
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Type	String
Default	No selectivity
Effects	Defines which selectivity ogive should be applied when calculating the fits. Use a selectivity label as per selectivity_names.
class_mins	What are the size bins of the observations (in an age-based model)?
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Conditions	Age-based model with at_size set, i.e., a size frequency time series in an age-based model
Type	Constant vector
Effects	Defines the lower limits of each of the size classes. If there is no plus group then an additional value defines the upper limit of the last size class.

Notes	If the observations are sexed, both sexes share the same list of size classes (but can use different subsets of them, see <code>min_class</code> , <code>max_class</code>) — do not put separate size classes for the two sexes here.
class_nums	What are the class numbers for the observations (in a size-based model)?
Command	<code>numbers_at [label], relative_numbers_at [label], proportions_at [label]</code>
Conditions	Size-based model. The <code>sexed</code> switch needs to be set to false.
Type	Constant vector of integers.
Effects	Defines the lower limits of each of the size classes, where size classes are indexed by the size class number (not by size). The size class includes the lower limit. If there is no plus group then the last value defines the upper limit of the last size class (but does not include this value).
class_nums_male, class_nums_female	What are the class numbers for the sexed observations (in a size-based model)?
Command	<code>numbers_at [label], relative_numbers_at [label], proportions_at [label]</code>
Conditions	Size-based model. The <code>sexed</code> switch needs to be set to true.
Type	Constant vector of integers.
Effects	Defines the lower limits of each of the size classes, where size classes are indexed by the size class number (not by size). The size class includes the lower limit. If there is no plus group then the last value defines the upper limit of the last size class (but does not include this value).
Notes	If the observations are not sexed then use <code>class_nums</code> . If the observations are sexed then use <code>class_nums_male</code> and <code>plus_group</code>
plus_group	Is the last age or size class a plus group?
Command	<code>numbers_at [label], relative_numbers_at [label], proportions_at [label]</code>
Type	Switch
Default	True (it is a plus group)
Effects	Defines the last age or size class as a plus group.
min_class, max_class	Which age/size classes are covered by the observations?
Command	<code>numbers_at [label], relative_numbers_at [label], proportions_at [label]</code>
Type	2 x constant vector
Default	All classes are covered
Effects	Defines the first and last age or size classes covered by the observations. If the observations are sexed, each parameter is a 2-vector, for males then females. If unsexed, each parameter is a 1-vector. Age classes are indexed by age. Size classes by size class number (not by size).
Example	Unsexed observations covering 2-year-olds to 6-year-olds would have <code>min_class=2, max_class=6</code> . Sexed observations covering 2 to 12-year-old males and 2 to 14-year-old females would have <code>min_class = 2 2, max_class = 12 14</code> . If the model uses size classes of 20-30, 30-40, 40-50, 50+ cm, and the unsexed observations only cover 20-50 cm, then they would have <code>min_class = 1, max_class = 3</code> .
sum_to_one	Should the proportions sum to 1?
Command	<code>proportions_at [label]</code>
Type	Switch
Default	True (proportions are expected to sum to 1)
Effect	Defines proportions as summing to 1 for each year of observations.
Notes	This switch is provided for compatibility with previous NIWA software. It should only be turned off for least-squares estimation. The default is recommended. See Section 5.6 for details.
ageing_error	Should ageing error be applied to these observations?
Command	<code>numbers_at [label], relative_numbers_at [label], proportions_at [label]</code>

Conditions	Only used in an age-based model in which ageing error has been specified.
Type	Switch
Default	True (use ageing error as specified)
Effects	Define that these observations use ageing error.
Notes	This is used to turn ageing error off for an individual time series. Ageing error is only applied if it is specified in an @ageing_error block (see Section 8.10).
[year]	Numbers or proportions for [year]
Command	numbers_at [label], relative_numbers_at [label], proportions_at [label]
Type	Constant
Effects	Defines the observations for [year]. If the observations are sexed, put male observations first then female observations (on the same row).
Example	If you have sexed observations of 2, 3, and 4-year-olds for 1992, 1995, and 1998, you would put, e.g., <pre>years 1992 1995 1998 sexed 1 min_class 2 2 max_class 4 4 # M2 M3 M4 F2 F3 F4 1992 0.1 0.2 0.3 0.1 0.2 0.1 1995 0.2 0.3 0.1 0.2 0.1 0.1 1998 0.3 0.1 0.1 0.2 0.2 0.1</pre>

(See Section 8.8 for parameters relating to the likelihood or least-squares weights.)

@catch_at	Catch_at block command
Effects	Defines any following commands as @catch_at subcommands for the time series.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop)
years	Years of the time series
Command	catch_at [label]
Type	Constant vector
Effects	Defines the years for which there are observations. Should be one entry per row of observations.
fishery	Fishery or fisheries covered by the observations
Command	catch_at [label]
Type	Vector of strings
Effects	Defines the fisheries included by the observations. Use fishery labels as in annual_cycle.fishery_names.
Notes	Typically a single set of observations will cover only one fishery, but it might be the case that a single administrative fishery might be split into several CASAL fisheries, and the observations are provided for all those fisheries combined.
at_size	Are the observations by size?
Command	catch_at [label]
Type	Switch
Default	True (observations are size-based in a size-based model, age-based in an age-based model)
Effects	Defines the observations as size-based (at_size=true) or age-based (at_size=false).
sexed	Are the observations sexed?
Command	catch_at [label]
Condition	sex_partition is set
Type	Switch
Default	True (observations are sexed)

Effects	Defines the observations as sexed (<code>sexed=true</code>) or unsexed (<code>sexed=false</code>).
class_mins	What are the size bins of the observations (in an age-based model)?
Command	<code>catch_at [label]</code>
Conditions	Age-based model with <code>at_size</code> set, i.e., a size frequency time series in an age-based model
Type	Constant vector
Effects	Defines the lower limits of each of the size classes. If there is no plus group then an additional value defines the upper limit of the last size class.
Notes	If the observations are sexed, both sexes share the same list of size classes (but can use different subsets of them, see <code>min_class</code> , <code>max_class</code>) — do not put separate size classes for the two sexes here.
class_nums	What are the class numbers for the observations (in a size-based model)?
Command	<code>catch_at [label]</code>
Conditions	Size-based model. The <code>sexed</code> switch needs to be set to false.
Type	Constant vector of integers.
Effects	Defines the lower limits of each of the size classes, where size classes are indexed by the size class number (not by size). The size class includes the lower limit. If there is no plus group then the last value defines the upper limit of the last size class (but does not include this value).
Notes	If the observations are sexed then use <code>class_nums_male</code> and <code>class_nums_female</code> .
class_nums_male, class_nums_female	What are the class numbers for the sexed observations (in a size-based model)?
Command	<code>catch_at [label]</code>
Conditions	Size-based model. The <code>sexed</code> switch needs to be set to true.
Type	Constant vector of integers.
Effects	Defines the lower limits of each of the size classes, where size classes are indexed by the size class number (not by size). The size class includes the lower limit. If there is no plus group then the last value defines the upper limit of the last size class (but does not include this value).
Notes	If the observations are not sexed then use <code>class_nums</code> .
plus_group	Is the last age or size class a plus group?
Command	<code>catch_at [label]</code>
Type	Switch
Default	True (it is a plus group)
Effects	Defines the last age or size class as a plus group.
min_class, max_class	Which age/size classes are covered by the observations?
Command	<code>catch_at [label]</code>
Type	2 x constant vector
Default	All classes are covered
Effects	Defines the first and last age or size classes covered by the observations. If the observations are sexed, each parameter is a 2-vector, for males then females. If unsexed, each parameter is a 1-vector. Age classes are indexed by age. Size classes by size class number (not by size).
Example	Unsexed observations covering 2-year-olds to 6-year-olds would have <code>min_class = 2, max_class = 6</code> . Sexed observations covering 2 to 12-year-old males and 2 to 14-year-old females would have <code>min_class = 2 2, max_class = 12 14</code> . If the model uses size classes of 20-30, 30-40, 40-50, 50+ cm, and the unsexed observations only cover 20-50 cm, then they would have <code>min_class = 1, max_class = 3</code> .
sum_to_one	Should the proportions sum to 1?
Command	<code>catch_at [label]</code>
Type	Switch
Default	True (proportions are expected to sum to 1)
Effect	Defines proportions as summing to 1 for each year of observations.

Notes	This switch is provided for compatibility with previous NIWA software. It should only be turned off for least-squares estimation. The default is recommended. See Section 5.6 for details.
ageing_error	Should ageing error be applied to these observations?
Command	catch_at [label]
Conditions	Only used in an age-based model in which ageing error has been specified.
Type	Switch
Default	True (use ageing error as specified)
Effects	Define that these observations use ageing error.
Notes	This is used to turn ageing error off for an individual time series. Ageing error is only applied if it is specified in an @ageing_error block (see Section 8.10).
[year]	Numbers or proportions for [year]
Command	catch_at [label]
Type	Constant
Effects	Defines the observations for [year]. If the observations are sexed, put male observations first then female observations (on the same row).
Example	If you have sexed observations of 2, 3, and 4-year-olds for 1992, 1995, and 1998, you would put, e.g., years 1992 1995 1998 sexed 1 min_class 2 2 max_class 4 4 # M2 M3 M4 F2 F3 F4 1992 0.1 0.2 0.3 0.1 0.2 0.1 1995 0.2 0.3 0.1 0.2 0.1 0.1 1998 0.3 0.1 0.1 0.2 0.2 0.1

(See Section 8.8 for parameters relating to the likelihood or least-squares weights.)

@proportions_mature Proportions_mature block command

Effects	Defines any following commands as @proportions_mature subcommands for the time series.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop)
years	Years of the time series
Command	proportions_mature [label]
Type	Constant vector
Effects	Defines the years for which there are observations. Should be one entry per row of observations.
step	Time step in which the observations occur
Command	proportions_mature [label]
Type	Integer
Effects	Defines the time step in which the observations occur.
proportion_mortality	Proportion of the step's mortality, prior to when the observations occur
Command	proportions_mature [label]
Type	Constant
Default	0.5
Effects	Defines the proportion of the mortality in the time step, prior to when the observations occur.
sexed	Are these observations sexed?
Command	proportions_mature [label]
Type	Switch
Default	True (for a sex-based model) and false otherwise

Effects	Defines these observations as sexed (<code>sexed=true</code>) or unsexed (<code>sexed=false</code>).
females_only	Are these observations for females only?
Command	<code>proportions_mature [label]</code>
Type	Switch
Default	False (both males and females)
Effects	Defines these observations as females only (<code>females_only=true</code>) or both sexes separately (<code>females_only=false</code>).
Notes	Do not use <code>females_only=true</code> for <code>sexed=false</code> observations or in an unsexed model.
at_size	Are the observations by size?
Command	<code>proportions_mature [label]</code>
Type	Switch
Default	<code>size_based</code> (i.e., size-based in a size-based model, age-based in an age-based model)
Effects	Defines the observations as size-based (<code>at_size=true</code>) or age-based (<code>at_size=false</code>).
area	Area in which the observations occur
Command	<code>proportions_mature [label]</code>
Condition	<code>n_areas > 1</code>
Type	String
Effects	Defines the area in which the observations occur. Use an area label as per <code>area_names</code> .
ogive	Which selectivity ogive should be applied?
Command	<code>proportions_mature [label]</code>
Type	String
Default	No selectivity
Effects	Defines which selectivity ogive should be applied when calculating the fits. Use a selectivity label as per <code>selectivity_names</code> .
Notes	The selectivity ogive only affects the results if it is a size-based ogive in an age-based model.
class_mins	What are the size bins of the observations (in an age-based model)?
Command	<code>proportions_mature [label]</code>
Conditions	Age-based model with <code>at_size</code> set, i.e., a size frequency time series in an age-based model
Type	Constant vector
Effects	Defines the lower limits of each of the size classes. If there is no plus group then an additional value defines the upper limit of the last size class.
Notes	If the observations are sexed, both sexes share the same list of size classes (but can use different subsets of them, see <code>min_class</code> , <code>max_class</code>) — do not put separate size classes for the two sexes here.
plus_group	Is the last age or size class a plus group?
Command	<code>proportions_mature [label]</code>
Type	Switch
Default	True (it is a plus group)
Effects	Defines the last age or size class as a plus group.
min_class, max_class	Which age/size classes are covered by the observations?
Command	<code>proportions_mature [label]</code>
Type	2 x constant vector
Default	All classes are covered
Effects	Defines the first and last age or size classes covered by the observations. If the observations are sexed, each parameter is a 2-vector, for males then females. If unsexed, each parameter is a 1-vector. Age classes are indexed by age. Size classes by size class number (not by size).

Example	Unsexed observations covering 2-year-olds to 6-year-olds would have <code>min_class = 2, max_class = 6</code> . Sexed observations covering 2 to 12-year-old males and 2 to 14-year-old females would have <code>min_class = 2 2, max_class = 12 14</code> . If the model uses size classes of 20-30, 30-40, 40-50, 50+ cm, and the unsexed observations only cover 20-50 cm, then they would have <code>min_class = 1, max_class = 3</code> .
ageing_error	Should ageing error be applied to these observations?
Command	<code>proportions_mature [label]</code>
Conditions	Only used in an age-based model in which ageing error has been specified.
Type	Switch
Default	True (use ageing error as specified)
Effects	Define that these observations use ageing error.
Notes	This is used to turn ageing error off for an individual time series. Ageing error is only applied if it is specified in an <code>@ageing_error</code> block (see Section 8.10).
[year]	Proportions mature for [year]
Command	<code>proportions_mature [label]</code>
Type	Constant
Effects	Defines the observations for [year]. If the observations are for both sexes, put male observations first then female observations (on the same row).
Example	If you have sexed observations of 2, 3, and 4-year-olds for 1992 and 1995, you would put, e.g., <pre>years 1992 1995 min_class 2 2 max_class 4 4 # M2 M3 M4 F2 F3 F4 1992 0.1 0.5 0.8 0.1 0.4 0.9 1995 0.2 0.6 1.0 0.2 0.7 0.9</pre>
@proportions_migrating	Proportions_migrating block command
Effects	Defines any following commands as <code>@proportions_migrating</code> subcommands for the time series.
Label	The text label of the time series (should be unique, not <code>Bpre</code> or <code>Bpost</code> , and not contain a full stop)
years	Years of the time series
Command	<code>proportions_migrating [label]</code>
Type	Constant vector
Effects	Defines the years for which there are observations. Should be one entry per row of observations.
migration	Migration to which the observations apply
Command	<code>proportions_migrating [label]</code>
Type	String
Effects	Defines the migration to which the observations apply. Use a label from <code>annual_cycle.migration_names</code> .
sex	Which sex do the observations apply to?
Command	<code>proportions_migrating [label]</code>
Condition	<code>sex_partition</code> is set
Type	Integer
Default	(observations apply to both sexes combined)
Effects	Defines the observations as male (<code>sex=1</code>) or female (<code>sex=2</code>).
Notes	If you have separate observations for males and females, enter them as two different time series.
at_size	Are the observations by size?
Command	<code>proportions_migrating [label]</code>
Type	Switch

Default	size_based (i.e., size-based in a size-based model, age-based in an age-based model)
Effects	Defines the observations as size-based (at_size=true) or age-based (at_size=false).
area	Area in which the observations occur
Command	proportions_migrating[label]
Condition	n_areas > 1
Type	String
Effects	Defines the area in which the observations occur. Use an area label as per area_names.
ogive	Which selectivity ogive should be applied?
Command	proportions_migrating[label]
Type	String
Default	No selectivity
Effects	Defines which selectivity ogive should be applied when calculating the fits. Use a selectivity label as per selectivity_names.
Notes	The selectivity ogive only affects the results if it is a size-based ogive in an age-based model. Probably this parameter is unnecessary.
class_mins	What are the size bins of the observations (in an age-based model)?
Command	proportions_migrating[label]
Conditions	Age-based model with at_size set, i.e., a size frequency time series in an age-based model
Type	Constant vector
Effects	Defines the lower limits of each of the size classes. If there is no plus group then an additional value defines the upper limit of the last size class.
Notes	If the observations are sexed, both sexes share the same list of size classes (but can use different subsets of them, see min_class, max_class) — do not put separate size classes for the two sexes here.
plus_group	Is the last age or size class a plus group?
Command	proportions_migrating[label]
Type	Switch
Default	True (it is a plus group)
Effects	Defines the last age or size class as a plus group.
min_class, max_class	Which age/size classes are covered by the observations?
Command	proportions_migrating[label]
Type	2 x constant vector
Default	All classes are covered
Effects	Defines the first and last age or size classes covered by the observations. If the observations are sexed, each parameter is a 2-vector, for males then females. If unsexed, each parameter is a 1-vector. Age classes are indexed by age. Size classes by size class number (not by size).
Example	Unsexed observations covering 2-year-olds to 6-year-olds would have min_class = 2, max_class = 6. Sexed observations covering 2 to 12-year-old males and 2 to 14-year-old females would have min_class = 2 2, max_class = 12 14. If the model uses size classes of 20-30, 30-40, 40-50, 50+ cm, and the unsexed observations only cover 20-50 cm, then they would have min_class = 1, max_class = 3.
ageing_error	Should ageing error be applied to these observations?
Command	proportions_migrating[label]
Conditions	Only used in an age-based model in which ageing error has been specified.
Type	Switch
Default	True (use ageing error as specified)
Effects	Define that these observations use ageing error.
Notes	This is used to turn ageing error off for an individual time series. Ageing error is only applied if it is specified in an @ageing_error block (see Section 8.10).

[year]	Proportions migrating for [year]
Command	proportions_migrating[label]
Type	Constant
Effects	Defines the observations for [year]. If the observations are for both sexes, put male observations first then female observations (on the same row).
Example	If you have sexed observations of 2, 3, and 4-year-olds for 1992 and 1995, you would put, e.g., <pre>years 1992 1995 min_class 2 2 max_class 4 4 # M2 M3 M4 F2 F3 F4 1992 0.1 0.5 0.8 0.1 0.4 0.9 1995 0.2 0.6 1.0 0.2 0.7 0.9</pre>
@age_size	Age_size block command
Effects	Defines any following commands as @age_size subcommands for the age/size dataset.
Conditions	Only used in an age-based model. Do not use age/size data in a growth-path model, nor one where size-at-age depends on maturity.
Label	The text label of the data (should be unique and not contain a full stop).
year	Year in which the data were collected
Command	age_size[label]
Type	Integer
Effects	Defines the year in which the data were collected.
step	Time step in which the data were collected
Command	age_size[label]
Type	Integer
Effects	Defines the time step in which the data were collected.
proportion_mortality	Proportion of the step's mortality, prior to when the observations occur
Command	age_size[label]
Type	Constant
Default	0.5
Effects	Defines the proportion of the mortality in the time step, prior to when the observations occur.
area	Area in which the observations occur
Command	age_size[label]
Condition	n_areas > 1
Type	String
Effects	Defines the area in which the observations occur. Use an area label as per area_names.
stock	Stock for which the data were collected
Command	age_size[label]
Conditions	Not used in a single-stock model
Type	String
Default	Data were collected from all stocks in the selected area
Effects	Defines the stock which the fish belong to. Use a text label from stock_names.
Notes	This parameter would not normally be necessary.
sample	Sampling method under which the observations were generated
Command	age_size[label]
Type	String
Effects	Defines the sampling method used to generate the age-size observations. Options are random, random_at_sex, random_at_age, random_at_size, random_at_sex_and_size, and random_at_sex_and_age.

Notes	See the main text for an explanation of these options.
ogive	Which selectivity ogive should be applied?
Command	age_size[label]
Type	String
Default	No selectivity
Effects	Defines which selectivity ogive was used when taking this sample. Use a selectivity label as per selectivity_names.
Notes	With some sample structures there is no point in specifying a size-based selectivity because it has no effect (and involves more calculations). This is always true with a random_at_sex_and_size sample, and it is true with a random_at_size sample as long as the size-based selectivity function is not sex dependent. Similarly with some sample structures there is no point in using an age-based selectivity. This is always true with random_at_sex_and_age, and with random_at_age so long as the age-based selectivity is not sex dependent. See the main text for more explanation.
ageing_error	Should ageing error be applied to these observations?
Command	age_size[label]
Conditions	Only if ageing error has been specified.
Type	Switch
Default	True (use ageing error as specified)
Effects	Define that these observations use ageing error.
Notes	This is used to turn ageing error off for an individual set of observations. Ageing error is only applied if it is specified in an @ageing_error block (see Section 8.10).
ages	Age data
Command	age_size[label]
Type	Constant vector
Effects	Defines the age data. For each entry of 'ages', there should be a corresponding entry of 'sizes', and of 'sexes' in a sexed model.
Example	<pre>@age_size sizedata year 1990 step 1 ages 1 1 2 2 2 3 3 3 3 4 ... sizes 12 15 23 26 27 36 38 32 40 48 ... sexes 1 2 1 2 2 1 1 1 2 1 ...</pre>
sizes	Size data
Command	age_size[label]
Type	Constant vector
Effects	Defines the size data. For each entry of 'sizes', there should be a corresponding entry of 'ages', and of 'sexes' in a sexed model.
sexes	Sex data
Command	age_size[label]
Conditions	Only used in a sexed model
Type	Constant vector
Effects	Defines the sex data. 1 denotes male, 2 denotes female. For each entry of 'sexes', there should be a corresponding entry of 'ages' and of 'sizes'.
@age_at_maturation	Age_at_maturation block command
Effects	Defines any following commands as @age_at_maturation commands.
Label	The text label of the time series (should be unique, not Bpre or Bpost, and not contain a full stop)
sexed	Are these observations sexed?
Command	age_at_maturation[label]
Type	Switch
Default	True (for a sex-based model) and false otherwise

Effects	Defines these observations as sexed (<code>sexed=true</code>) or unsexed (<code>sexed=false</code>).
sampled_ages	What were the estimated ages of these fish at sampling?
Command	<code>age_at_maturation[label]</code>
Type	Constant vector
Effects	Defines the age of each fish at sampling.
Notes	For each entry in 'sampled_ages', there should be a matching entry in <code>maturation_ages</code> , and in <code>sexes</code> if <code>sexes=T</code> .
maturation_ages	What were the estimated ages of these fish at maturation?
Command	<code>age_at_maturation[label]</code>
Type	Constant vector
Effects	Defines the age of each fish at maturation.
Notes	For each entry in <code>maturation_ages</code> , there should be a matching entry in <code>sampled_ages</code> , and in <code>sexes</code> if <code>sexes=T</code> .
sexes	What were the sexes of these fish?
Command	<code>age_at_maturation[label]</code>
Conditions	Only use if <code>sexed=True</code>
Type	Constant vector
Effects	Defines the sex of each fish sampled.
Notes	1=male, 2=female. For each entry in 'sexes', there should be a matching entry in <code>maturation_ages</code> and <code>sampled_ages</code> .
stock	Which stock do these observations relate to?
Command	<code>age_at_maturation[label]</code>
Condition	<code>n_stocks > 1</code>
Type	String
Effects	Defines the name of the stock which is observed. Use a stock label from <code>stock_names</code> .
Notes	This parameter must be supplied in a multi-stock model, e.g., there is no default.
ageing_error	Should ageing error be applied to these observations?
Command	<code>age_at_maturation[label]</code>
Conditions	Only used if ageing error has been specified.
Type	Switch
Default	True (use ageing error as specified)
Effects	Define that the ages at maturation (not at capture!) are subject to ageing error.
Notes	This is used to turn ageing error off for an individual time series. Ageing error is only applied if it is specified in an <code>@ageing_error</code> block (see Section 8.10).
k	After how many years can maturation be detected?
Command	<code>age_at_maturation[label]</code>
Type	Constant
Default	$k = 0$
Effects	Defines the k parameter, e.g., the number of years after maturation during which maturation cannot be detected.

(See Section 9.2 for the `@selectivity_at_observation` command)

8.8 Defining the objective function associated with the observations

weight	Weight of this time series, in the Cordue weighted least-squares scheme
Command	<code>abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label]</code>
Conditions	Only used if <code>estimator=least_squares</code> , <code>weighting=Cordue</code> .

Type	Constant
Effects	Define the weight u for the time series.
cv_[year]	C.v.s by year for this time series, in the Cordue weighted least-squares scheme
Command	abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label]
Conditions	Only used if estimator=least_squares, weighting=Cordue. A parameter of the same name is used for some likelihoods.
Type	Constant
Effects	Defines the c.v.s by year for the time series. One c.v. for each year.
Example	years 1992 1993 1994 cv_1992 0.2 cv_1993 0.3 cv_1994 0.4
dist	Likelihood of the observations
Command	abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label], proportions_migrating[label]
Conditions	Only used if estimator=likelihood or Bayes. Can be user-supplied only if you have written a user.likelihood.C which calculates the likelihood function.
Type	String
Effects	Defines the type of likelihood for this time series. Can be: normal, lognormal, normal-log, or normal-by-stdev for abundance or proportions_mature or proportions_migrating; OR binomial for proportions_mature or proportions_migrating; OR normal, lognormal, robustified-lognormal, or normal-log for relative_abundance or relative_numbers_at; OR multinomial, Fournier, Coleraine, lognormal, or robustified-lognormal for proportions_at or catch_at; OR user-supplied for any of the above time series types.
r	Robustifying constant
Command	abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label], proportions_migrating[label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=multinomial or robustified-lognormal or binomial
Type	Constant
Default	0 (i.e., no robustification)
Effects	Defines the robustifying constant r for this time series. Constants r are used for the multinomial, the binomial, and the robustified lognormal — they mean quite different things in the three contexts.
cv	C.v. for all observations in this time series, used with likelihoods
Command	abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal, lognormal, robustified-lognormal, or normal-log.
Type	Constant
Effects	Defines a single c.v. used for every observation in the time series.
Notes	Alternatively supply c.v.s for each year using cv_[year] or for each individual observation using cvs_[year].

cv_[year]	C.v.s by year for this time series, used with likelihoods
Command	abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal, lognormal, robustified-lognormal, or normal-log. A parameter of the same name is used for the Cordue weighting scheme.
Type	Constant (1 per year)
Effects	Defines the c.v.s by year for the time series. One c.v. for each year.
Example	years 1992 1993 1994 cv_1992 0.2 cv_1993 0.3 cv_1994 0.4
Notes	Alternatively supply c.v.s for each individual observation using cvs_[year] or a single c.v. for all years using cv.
cvs_[year]	C.v.s by observation and year for this time series, used with likelihoods
Command	abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal, lognormal, robustified-lognormal, or normal-log.
Type	Constant vector (1 per year)
Effects	Defines the c.v.s by observation and year for the time series. One c.v. for each observation in each year.
Example	years 1992 1993 1994 sexed 1 min_class 2 2 max_class 4 4 # M2 M3 M4 F2 F3 F4 cvs_1992 0.40 0.15 0.25 0.55 0.15 0.20 cvs_1993 0.60 0.17 0.12 0.75 0.18 0.14 cvs_1994 0.80 0.19 0.08 0.85 0.31 0.15
Notes	Alternatively supply a single c.v. for each year using cv_[year] or a single c.v. for all years using cv.
cv_process_error	Process error c.v. for this time series, used with likelihoods parameterised by the c.v.
Command	abundance[label], relative_abundance[label], relative_numbers_at[label], proportions_at[label], catch_at[label], proportions_mature[label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal, lognormal, robustified-lognormal, or normal-log.
Type	Estimable
Effects	Defines a c.v. to be 'added' to all c.v.s by observation and year for the time series.
stdev	Standard deviation for all observations in this time series, used with normal-by-standard deviation likelihoods
Command	abundance[label], proportions_mature[label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal-by-stdev.
Type	Constant
Effects	Defines a single standard deviation used for every observation in the time series.
Notes	Alternatively supply standard deviation for each year using stdev_[year] or for each individual observation using stdevs_[year].

stdev_[year]	Standard deviation by year for this time series, used with normal-by-standard deviation likelihoods
Command	abundance [label], proportions_mature [label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal-by-stdev.
Type	Constant (1 per year)
Effects	Defines the standard deviations by year for the time series. One standard deviation for each year.
Example	years 1992 1993 1994 stdev_1992 0.2 stdev_1993 0.3 stdev_1994 0.4
Notes	Alternatively supply standard deviation for each individual observation using stdevs_[year] or a single std. dev. for all years using stdev.
stdevs_[year]	Standard deviation by observation and year for this time series, used with normal-by-standard deviation likelihoods
Command	abundance [label], proportions_mature [label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal-by-stdev.
Type	Constant vector (1 per year)
Effects	Defines the standard deviations by observation and year for the time series. One std. dev. for each observation in each year.
Example	years 1992 1993 1994 sexed 1 min_class 2 2 max_class 4 4 # M2 M3 M4 F2 F3 F4 stdev_1992 0.40 0.15 0.25 0.55 0.15 0.20 stdev_1993 0.60 0.17 0.12 0.75 0.18 0.14 stdev_1994 0.80 0.19 0.08 0.85 0.31 0.15
Notes	Alternatively supply a single standard deviation for each year using stdev_[year] or a single standard deviation for all years using stdev.
stdev_process_error	Process error standard deviation for this time series, used with likelihoods parameterised by the standard deviation
Command	abundance [label], proportions_mature [label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=normal-by-stdev.
Type	Estimable
Effects	Defines a std. dev. to be 'added' to all standard deviation by observation and year for the time series.
N	N for all years in this time series, used with proportions likelihoods
Command	proportions_at [label], catch_at [label], proportions_mature [label], proportions_migrating [label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=multinomial, Fournier, or Coleraine.
Type	Constant
Effects	Defines a single N used for every year in the time series.
Notes	Alternatively supply N's for each year using N_[year] or for each individual observation using Ns_[year].
N_[year]	N's by year for this time series, used with proportions likelihoods
Command	proportions_at [label], catch_at [label], proportions_mature [label], proportions_migrating [label]

Conditions	Only used if estimator=likelihood or Bayes AND dist=multinomial, Fournier, or Coleraine.
Type	Constant (1 per year)
Effects	Defines the N 's by year for the time series. One N for each year.
Example	years 1992 1993 1994 N_1992 100 N_1993 120 N_1994 150
Notes	Alternatively supply N 's for each individual observation using $Ns_ [year]$ or a single N for all years using N .
Ns_[year]	N's by observation and year for this time series, used with proportions likelihoods
Command	proportions_at [label], catch_at [label], proportions_mature [label], proportions_migrating [label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=multinomial, Fournier, or Coleraine.
Type	Constant vector (1 per year)
Effects	Defines the N 's by observation and year for the time series. One N for each observation in each year.
Example	years 1992 1993 1994 sexed 1 min_class 2 2 max_class 4 4 # M2 M3 M4 F2 F3 F4 Ns_1992 20 100 150 30 90 120 Ns_1993 30 122 156 40 80 125 Ns_1994 40 140 180 50 70 123
Notes	Alternatively supply a single N for each year using $N_ [year]$ or a single N for all years using N .
N_process_error	Process error N for this time series, used with likelihoods parameterised by the effective sample size.
Command	proportions_at [label], catch_at [label], proportions_mature [label], proportions_migrating [label]
Conditions	Only used if estimator=likelihood or Bayes AND dist=multinomial, Fournier, or Coleraine.
Type	Estimable
Effects	Defines an N to be 'added' to all N 's by observation and year for the time series.

8.9 Defining the penalties

@ogive_smoothing_penalty Ogive smoothing penalty block command

Effects	Defines any following commands as @ogive_smoothing_penalty subcommands
Notes	The i th @ogive_smoothing_penalty block relates to the i th penalty. The penalty is on the sum of squares of r th differences in the ogive values, encouraging the ogive to be like a polynomial of degree $r-1$.
label	The name of the penalty
Command	ogive_smoothing_penalty [i]
Type	String
Effects	Defines a text label for the penalty
ogive	The name of the ogive parameter to which the penalty is applied
Command	ogive_smoothing_penalty [i]

Type	String
Effects	Defines the name of the ogive parameter to which the penalty is applied.
Notes	Should be an <code>allvalues</code> or <code>allvalues_bounded</code> ogive. See Section 2.4 for instructions on generating parameter names.
r	Penalty is applied to <i>r</i>th differences
Command	<code>ogive_smoothing_penalty[i]</code>
Type	Integer
Effects	Defines the effect of the penalty, which acts on the sum of squares of <i>r</i> th differences (and hence encourages the ogive to be like a polynomial of degree <i>r</i> -1).
lower_bound, upper_bound	Penalty is applied for age or size classes lower_bound to upper_bound
Command	<code>ogive_smoothing_penalty[i]</code>
Type	2 x integer
Default	<code>min_age, max_age</code>
Effects	Defines the effect of the penalty, which acts on ogive elements for age or size classes of <code>lower_bound</code> to <code>upper_bound</code> . Everything else is dropped off before differencing.
multiplier	Multiply the penalty by this factor
Command	<code>ogive_smoothing_penalty[i]</code>
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.
@catch_limit_penalty	Catch limit penalty block command
Effects	Defines any following commands as <code>@catch_limit_penalty</code> subcommands
Notes	The <i>i</i> th <code>@catch_limit_penalty</code> block relates to the <i>i</i> th such penalty. The penalty is on the sum of squares of (actual catch - specified catch), optionally on a log scale, for a single fishery.
label	The name of the penalty
Command	<code>catch_limit_penalty[i]</code>
Type	String
Effects	Defines a text label for the penalty
fishery	The label of the fishery to which the penalty is applied
Command	<code>catch_limit_penalty[i]</code>
Type	String
Effects	Defines the label of the fishery to which the penalty is applied.
log_scale	Should sums of squares be calculated on the log scale?
Command	<code>catch_limit_penalty[i]</code>
Type	Switch
Default	False
Effects	Defines sum of squares to be calculated on the log scale (with <code>log_scale=true</code>) or the linear scale (with <code>log_scale=false</code>).
multiplier	Multiply the penalty by this factor
Command	<code>catch_limit_penalty[i]</code>
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.
@vector_average_penalty	Vector average penalty block command
Effects	Defines any following commands as <code>@vector_average_penalty</code> subcommands
Notes	The <i>i</i> th <code>@vector_average_penalty</code> block relates to the <i>i</i> th such penalty. The penalty is on the square of (mean(vector) - <i>k</i>), or of (mean(log(vector)) -

		l), or of $(\log(\text{mean}(\text{vector}) / m))$. This encourages the vector to average arithmetically to k or m , or geometrically to $\exp(l)$.
label	The name of the penalty	
Command	Vector_average_penalty[i]	
Type	String	
Effects	Defines a text label for the penalty	
vector	The name of the vector parameter to which the penalty is applied	
Command	Vector_average_penalty[i]	
Type	String	
Effects	Defines the name of the vector parameter to which the penalty is applied.	
Notes	See Section 2.4 for instructions on generating parameter names.	
k, l, m	Vector should average arithmetically to k or m, or geometrically to l.	
Command	Vector_average_penalty[i]	
Conditions	Supply one of k , l , and m .	
Type	Constant	
Effects	Define the number which the vector should average to. The penalty is on the square of $(\text{mean}(\text{vector}) - k)$, or of $(\text{mean}(\log(\text{vector})) - l)$, or of $(\log(\text{mean}(\text{vector}) / m))$.	
multiplier	Multiply the penalty by this factor	
Command	Vector_average_penalty[i]	
Type	Constant	
Effects	Defines the factor by which the penalty is multiplied.	
Notes	The larger this number, the more severe the penalty.	
@vector_smoothing_penalty	Vector smoothing penalty block command	
Effects	Defines any following commands as @vector_smoothing_penalty subcommands	
Notes	The i th @vector_smoothing_penalty block relates to the i th penalty. The penalty is on the sum of squares of r th differences in the vector values, encouraging the vector to be like a polynomial of degree $r - 1$.	
label	The name of the penalty	
Command	vector_smoothing_penalty[i]	
Type	String	
Effects	Defines a text label for the penalty	
vector	The name of the vector parameter to which the penalty is applied	
Command	vector_smoothing_penalty[i]	
Type	String	
Effects	Defines the name of the vector parameter to which the penalty is applied.	
Notes	See Section 2.4 for instructions on generating parameter names.	
r	Penalty is applied to rth differences	
Command	vector_smoothing_penalty[i]	
Type	Integer	
Effects	Defines the effect of the penalty, which acts on the sum of squares of r th differences (and hence encourages the vector to be like a polynomial of degree $r - 1$).	
lower_bound, upper_bound	Penalty is applied for vector index classes lower_bound to upper_bound	
Command	vector_smoothing_penalty[i]	
Type	2 x integer	
Default	1, and the length of vector	
Effects	Defines the effect of the penalty, which acts on vector elements of lower_bound to upper_bound. Everything else is dropped off before differencing.	

multiplier	Multiply the penalty by this factor
Command	<code>vector_smoothing_penalty[i]</code>
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.
@element_difference_penalty Element difference penalty block command	
Effects	Defines any following commands as <code>@element_difference_penalty</code> subcommands
Notes	The <i>i</i> th <code>@element_difference_penalty</code> block relates to the <i>i</i> th such penalty. The penalty is on the square of $(\text{vector}_1[i] - \text{vector}_2[i])$, and is intended to encourage the <i>i</i> th elements of the two vectors to be similar.
label	The name of the penalty
Command	<code>element_difference_penalty[i]</code>
Type	String
Effects	Defines a text label for the penalty
vector1, vector2	The name of the vector parameters to which the penalty is applied
Command	<code>element_difference_penalty[i]</code>
Type	2 x string
Effects	Defines the names of the vector parameters to which the penalty is applied.
Notes	See Section 2.4 for instructions on generating parameter names.
i	Penalise differences in the <i>i</i>th elements of the two vectors
Command	<code>element_difference_penalty[i]</code>
Type	Integer
Effects	Define which element of the two vectors is to be penalised. The penalty is on the squared difference between the <i>i</i> th elements of the vectors.
multiplier	Multiply the penalty by this factor
Command	<code>element_difference_penalty[i]</code>
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.
@YCS_difference_penalty YCS difference penalty block command	
Effects	Defines any following commands as <code>@YCS_difference_penalty</code> subcommands
Notes	The <i>i</i> th <code>@element_difference_penalty</code> block relates to the <i>i</i> th such penalty. The penalty is on the square of the difference in YCS for a given year for two stocks, and is intended to encourage the two stocks to have the same YCS for that year.
label	The name of the penalty
Command	<code>YCS_difference_penalty[i]</code>
Type	String
Effects	Defines a text label for the penalty
stock1, stock2	The names of the stocks to which the penalty is applied
Command	<code>YCS_difference_penalty[i]</code>
Type	2 x string
Effects	Defines the names of the stocks to whose YCS the penalty is applied.
year	Year for which the penalty is to be applied
Command	<code>YCS_difference_penalty[i]</code>
Type	Integer
Effects	Define which year the penalty is applied to. The penalty is on the squared difference between the YCS of the two stocks for this year.
Notes	This refers to the year in which the year class is spawned.
multiplier	Multiply the penalty by this factor

Command	YCS_difference_penalty[i]
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.
@similar_qs_penalty	Similar q's penalty block command
Effects	Defines any following commands as @similar_qs_penalty subcommands
Notes	The i th @similar_qs_penalty block relates to the i th such penalty. The penalty is on the square of $(\log(q_1) - \log(q_2))$, and is intended to encourage q_1 and q_2 to be similar.
label	The name of the penalty
Command	similar_qs_penalty[i]
Type	String
Effects	Defines a text label for the penalty
q1, q2	The names of the two q's to which the penalty is applied
Command	similar_qs_penalty[i]
Type	2 x string
Effects	Defines the names of the q 's to which the penalty is applied.
Notes	Use the labels of the q 's as specified in the relative observations commands.
multiplier	Multiply the penalty by this factor
Command	similar_qs_penalty[i]
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.
@ogive_comparison_penalty	Ogive comparison penalty block command
Effects	Defines any following commands as @ogive_comparison_penalty subcommands
Notes	The i th @ogive_comparison_penalty block relates to the i th such penalty. The penalty is on the sum of squares of $\max(\text{ogive}_1 - \text{ogive}_2, 0)$, and is intended to encourage ogive_1 to be at or below ogive_2 . Cannot be used on size-based ogives in an age-based model.
label	The name of the penalty
Command	ogive_comparison_penalty[i]
Type	String
Effects	Defines a text label for the penalty
ogive1, ogive2	The name of the ogive parameters to which the penalty is applied
Command	ogive_comparison_penalty[i]
Type	2 x string
Effects	Defines the names of the vector parameters to which the penalty is applied.
Notes	See Section 2.4 for instructions on generating parameter names.
lower_bound, upper_bound	Penalty is applied for age or size classes lower_bound to upper_bound
Command	ogive_comparison_penalty[i]
Type	2 x integer
Default	min_age, max_age
Effects	Defines the effect of the penalty, which acts on ogive elements for age or size classes of lower_bound to upper_bound. Everything else is dropped off before comparing.
multiplier	Multiply the penalty by this factor
Command	ogive_comparison_penalty[i]
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.

@ogive_difference_penalty Ogive difference penalty block command

Effects	Defines any following commands as @ogive_difference_penalty subcommands
Notes	The <i>ith</i> @ogive_difference_penalty block relates to the <i>ith</i> such penalty. The penalty is on the square of (ogive ₁ - ogive ₂) for a given size or age class, and is intended to encourage the two ogives to have similar values for that size or age class. Cannot be used on size-based ogives in an age-based model.
label	The name of the penalty
Command	ogive_difference_penalty[i]
Type	String
Effects	Defines a text label for the penalty
ogive1, ogive2	The name of the ogive parameters to which the penalty is applied
Command	ogive_difference_penalty[i]
Type	2 x string
Effects	Defines the names of the vector parameters to which the penalty is applied.
Notes	See Section 2.4 for instructions on generating parameter names.
class	The age, or size class number, to which the penalty is applied
Command	ogive_difference_penalty[i]
Type	Integer
Effects	Defines the effect of the penalty, which acts on the specified age, or size class. If an age class, give the age in years. If a size class, give a number from 1, 2, ... where 1 is the smallest size class.
multiplier	Multiply the penalty by this factor
Command	ogive_difference_penalty[i]
Type	Constant
Effects	Defines the factor by which the penalty is multiplied.
Notes	The larger this number, the more severe the penalty.

8.10 Defining the ageing error

@ageing_error	Ageing error block command
Conditions	Only used in an age-based model.
Effects	Defines any following commands as @ageing_error subcommands.
Notes	Note spelling.
type	Type of ageing error model
Command	ageing_error
Conditions	Only used in an age-based model.
Type	String
Default	none (i.e., no ageing error)
Effects	Defines the ageing error model as none, off_by_one, normal, or misclassification_matrix.
Notes	Note punctuation.
p1, p2, k	Parameters of the off_by_one ageing error model
Command	ageing_error
Conditions	Only used in an age-based model. Only used if type=off_by_one.
Type	2 x estimable, 1 x integer
Default	The default of k is 0, i.e., fish of all ages can be misclassified. There are no defaults for p1 and p2.
Effects	p1 and p2 define the proportions of misclassifications down and up by 1 year respectively. k defines the minimum age of fish which can be misclassified — fish under age k have no ageing error.

c	Parameter of the normal ageing error model
Command	<code>ageing_error</code>
Conditions	Only used in an age-based model. Only used if <code>type=normal</code> .
Type	Estimable
Effects	Define the c.v. of misclassification.
[age]	Row of the misclassification matrix for the <code>misclassification_matrix</code> ageing error model.
Command	<code>ageing_error</code>
Conditions	Only used in an age-based model. Only used if <code>type=misclassification_matrix</code> .
Type	Constant vector
Effects	Define a row of the misclassification matrix. Enter one row per age class. The subcommand is the age class, from <code>min_age</code> to <code>max_age</code> .
Notes	Each row corresponds to a true age, each column to a reported age.
Example	With <code>min_age=2</code> , <code>max_age=4</code> , you might enter commands as follows, <pre>2 0.9 0.1 0.0 3 0.1 0.7 0.2 4 0.0 0.2 0.8</pre> Then, for example, 20% of 4-year-olds are misclassified as 3-year-olds.

8.11 CASAL extensions

@user_parameterisation	Is there a reparameterisation of the population section using <code>user.parameterisation.C</code>?
Type	Switch
Default	False (no reparameterisation)
Effects	Tells CASAL whether the population parameters are transformed using the user-supplied function <code>user.parameterisation.C</code> .
Notes	Only available if you have access to the CASAL source code.
@user_components	Lists the names of the user-defined priors or penalties calculated in <code>user.prior_penalty.C</code>
Conditions	Only used if you have written code in <code>user.prior_penalty.C</code>
Type	Vector of strings
Default	No user-defined objective function components
Effects	Gives the labels of each of the objective function components calculated in <code>user.prior_penalty.C</code> .
Notes	Only available if you have access to the CASAL source code.

9. THE OUTPUT.CSL FILE

The output parameters are specified in the `output.csl` file. See Section 6 for information about the output section, and Section 2.4 for instructions on writing a CASAL data file.

9.1 Defining the printouts

@print	Printouts block command
Effects	Defines any following commands as @print subcommands.
parameters	Print the population, estimation, and output parameters?
Command	print
Type	Switch
Default	False (do not print)
Effects	Print the population, estimation, and output parameters to the standard output as they are read in from <code>population.csl</code> , <code>estimation.csl</code> , <code>output.csl</code>
unused_parameters	Print a list of the parameters that were never used?
Command	print
Conditions	Only usable with <code>-r</code> , <code>-e</code> , <code>-E</code> .
Type	Switch
Default	False (do not print)
Effects	Print a list of the names of parameters which were never accessed by CASAL.
Notes	There are several reasons why a parameter might occur on this list. <ol style="list-style-type: none"> 1. Because its name was spelt incorrectly 2. Because it was an unnecessary or nonexistent parameter 3. Because the task you were doing does not use that parameter, e.g., <code>casal -r</code> does not use <code>max_iters</code> 4. Something unexpected happened.
population_section	Print a description of the population section?
Command	print
Conditions	Unless suppressed by <code>-q</code>
Type	Switch
Default	False (do not print)
Effects	Print a text description of the population section to the standard output.
requests, results	Print a description of the requests sent to the population section and the corresponding results?
Command	print
Conditions	Unless suppressed by <code>-q</code>
Type	2 x switch
Default	False (do not print)
Effects	Print the requests sent to the population section by the estimation and output sections, and the corresponding results, to the standard output.
Notes	Can produce a lot of output, only intended for use when debugging.
initial_state, state_annually, state_every_step, final_state	Print the state of the population?
Command	print
Conditions	Unless suppressed by <code>-q</code>
Type	4 x switch
Default	False (do not print)
Effects	Print the initial state, the state after every year, after every step, or the final state, to the standard output.
Notes	Can produce a lot of output, only intended for use when debugging.
estimation_section	Print a description of the estimation section?
Command	print
Type	Switch

Default	False (do not print)
Effects	Print a text description of the estimation section to the standard output.
fits, resids, pearson_resids, normalised_resids	Print the fits, residuals, and standardised residuals?
Command	print
Type	4 x switch
Default	False (do not print)
Effects	At the end of a model run (-r) or a parameter estimation (-e, -E), print the fits, residuals, Pearson residuals, and/or normalised residuals to the standard output.
covariance	Print the approximate covariance matrix of the free parameters
Command	print
Type	Switch
Default	False (do not print)
Effects	At the end of a parameter estimation (-e, -E), print the approximate covariance matrix of the free parameters, i.e., the inverse of the approximation to the Hessian. In addition, print out the sorted list of eigenvalues of the Hessian matrix.
Notes	The minimiser needs a number of iterations to build up an accurate Hessian approximation. If few iterations were done, the approximation will be poor. In any case the inverse Hessian is not a good approximation to the covariance of the free parameters, and should not be used for, for example, constructing confidence bounds on parameters.
yields	Print a description of the yield calculations?
Command	print
Conditions	Only used during yield calculations.
Type	Switch
Default	True (print)
Effects	Print a text description of the yield calculations to the standard output.
fits_every_eval, objective_every_eval, parameters_every_eval, parameter_vector_every_eval	Print the fits, objective function, or parameters at every function evaluation?
Command	print
Type	4 x switch
Default	False (do not print)
Effects	During each evaluation of the objective function, print the fits and residuals, objective function and components, free parameters in a verbose format, and/or free parameters as a vector, to the standard output.
Notes	Can produce a lot of output, only intended for use when debugging.
@print_sizebased_ogives_at	Sizes for which size-based ogives should be printed in an age-based model
Conditions	If printing of size-based ogives in an age-based model is required then the command must be specified. Ignored in a size-based model.
Type	Constant vector
Effects	Defines the sizes for which each size-based ogive will be output, (i) if you ask for a printout of a population section which has size-based selectivity ogives, and (ii) if your output quantities include any of the size-based ogive parameters.

9.2 Defining the output quantities

@quantities	Output quantities block command
Effects	Defines any following commands as @quantities subcommands.
all_free_parameters	Output quantities include all free parameters?
Command	quantities

Type	Switch
Default	False (do not print)
Effects	Include all free parameters in the output quantities.
Notes	If a parameter is an ogive, then the values of the ogive are printed, not the ogive arguments.
scalar_parameters, vector_parameters, ogive_parameters Output quantities include these parameters	
Command	quantities
Type	3 x vector of strings
Default	False (do not print)
Effects	Include some parameters in the output quantities. For each of these three subcommands, give a list of parameter names (which can be constant or free).
Notes	Generate parameter names as described in Section 2.4. If a parameter is an ogive, then the values of the ogive are printed, not the ogive arguments.
ogive_arguments Output quantities include these ogive arguments	
Command	quantities
Type	Vector of strings
Default	False (do not print)
Effects	Include the arguments of some ogives in the output quantities. Give a list of ogive parameter names.
Notes	Generate parameter names as described in Section 2.4.
nuisance_qs Output quantities include nuisance q's?	
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include nuisance q 's in the output quantities.
B0, R0, Bmean, Rmean Binitial, Rinitial Output quantities include $B_0, R_0, B_{mean}, R_{mean}, B_{initial}, R_{initial}$?	
Command	quantities
Type	6 x switch
Default	False (do not print)
Effects	Include $B_0, R_0, B_{mean}, R_{mean}, B_{initial}, R_{initial}$ for each stock in the output quantities.
SSBs Output quantities include SSBs?	
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include SSBs for each stock in the output quantities.
actual_catches, actual_catches_by_stock Output quantities include actual catches (by stock)?	
Command	quantities
Type	2 x switch
Default	False (do not print)
Effects	Include actual catches, and actual catches of each stock, in the output quantities.
recruitments Output quantities include recruitments?	
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include recruitments (as absolute numbers of fish, indexed by the year in which they recruit) in the output quantities.
YCS Output quantities include YCS?	
Command	quantities
Type	Switch
Default	False (do not print)

Effects	Include YCS (as deviates, indexed by the year in which they are spawned) in the output quantities.
Notes	The YCS can be different from the YCS parameter you input for several reasons. First, because if $n_rinitial > 0$, the YCS are prefixed by one or more entries of $R_{initial}$. Second, if you use the Haist YCS parameterisation. Third, if you are randomising some YCS in projection, or all YCS in yield simulation.
true_YCS	Output quantities include ‘true YCS’ = $YCS \times SR \times CR$?
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include <code>true_YCS</code> (defined as $YCS \times SR \times CR$, indexed by the year in which the fish are spawned) in the output quantities.
Notes	See <code>YCS</code> above. The ‘true YCS’ are arguably more informative than the <code>YCS</code> when there is a stock- or climate-recruitment relationship. Multiplying the <code>true_YCS</code> by R_0 gives the absolute recruitments.
Ts	Output quantities include climate variable T?
Command	quantities
Conditions	A climate-recruitment relationship is used
Type	Switch
Default	False (do not print)
Effects	Include the climate variable T (for each stock) in the output quantities.
fishing_pressures	Output quantities include fishing pressures?
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include fishing pressures for each fishery in the output quantities.
stock_crash	Output quantities include ‘stock crash’?
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include the ‘stock crash’ quantity for each stock in the output quantities. This is the indicator of the event that SSB falls below 20% B_0 in the projection period.
Notes	For projections only. The stock risk is the projected expectation of this quantity.
fits	Output quantities include the ‘fits’ or the expected value for each observation?
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include the ‘fits’ quantity for each stock in the output quantities.
Warning	The use of this option can generate a large amount of output, depending on the number of observations in your input data files.
resids	Output quantities include the ‘residuals’ or the observed less expected value for each observation?
Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include the ‘resids’ quantity for each stock in the output quantities.
Warning	The use of this option can generate a large amount of output, depending on the number of observations in your input data files.
pearson_resids	Output quantities include the ‘Pearson residuals’ for each observation?
Command	quantities
Type	Switch
Default	False (do not print)

Effects	Include the 'pearson_resids' quantity for each stock in the output quantities.
Warning	The use of this option can generate a large amount of output, depending on the number of observations in your input data files.

normalised_resids Output quantities include the 'normalised residuals' for each observation?

Command	quantities
Type	Switch
Default	False (do not print)
Effects	Include the 'normalised_resids' quantity for each stock in the output quantities.
Warning	The use of this option can generate a large amount of output, depending on the number of observations in your input data files.

Pseudo-fits are a special case. Specify the pseudo-observations using the observations parameters listed below or in Section 8.7. Do not use relative observation types. Do not provide the observations values, or any of the commands in Section 8.8, which relate to the objective function.

@selectivity_at Selectivity-at block command

Effects	Defines any following commands as @selectivity_at subcommands for the time series.
Label	The text label of the pseudo-observation time series (should be unique, not Bpre or Bpost, and not contain a full stop).
Note	Selectivity_at will probably only ever be used as a pseudo-observation. Its purpose is to extract the values of a selectivity ogive for a particular year, time step, area, stock.

ogive Which selectivity should be applied?

Command	selectivity_at [label]
Type	String
Default	None
Effects	Defines which selectivity ogive is being queried. Use a selectivity label as per selectivity_names.

years Years of the time series

Command	selectivity_at [label]
Type	Constant vector
Effects	Defines the years for which there are observations. Should be one entry for each year in which you want to query the selectivity.

step Time step in which the observations occur

Command	selectivity_at [label]
Type	Integer
Effects	Defines the time step in which the observations occur.

proportion_mortality Proportion of the step's mortality after which the observations occur

Command	selectivity_at [label]
Type	Constant
Default	0.5
Effects	Defines the proportion of the mortality in the time step after which the observations occur.

sexed Are the observations sexed?

Command	selectivity_at [label]
Condition	sex_partition is set
Type	Switch
Default	True (observations are sexed)
Effects	Defines the observations as sexed (sexed=true) or unsexed (sexed=false).

area	Area in which the observations occur
Command	<code>selectivity_at [label]</code>
Condition	<code>n_areas > 1</code>
Type	String
Effects	Defines the area in which the observations occur. Use an area label as per <code>area_names</code> .
mature_only	Do these observations include mature fish only?
Command	<code>selectivity_at [label]</code>
Type	Switch
Default	False (include both immature and mature fish)
Effects	Defines whether the observations are biomass (<code>biomass=true</code>) or numbers of fish (<code>biomass=false</code>).
Notes	You would only use this subcommand when generating pseudo-fits (Section 6.2), so as to output mature abundance.
stock	Which stock do these observations relate to?
Command	<code>selectivity_at [label]</code>
Condition	<code>n_stocks > 1</code>
Type	String
Default	(All stocks)
Effects	Defines the name of the stock which is observed. Use a stock label from <code>stock_names</code> .
Notes	You would only use this subcommand when generating pseudo-fits (Section 6.2), so as to output abundance of a particular stock.
[year]	Numbers or proportions for [year]
Command	<code>selectivity_at [label]</code>
Type	Constant
Effects	Defines the observations for [year].
Notes	This probably should never be supplied, assuming that this observation type is only ever used for pseudo-observations.

9.3 Defining projections

Most of the projections parameters are in the `population.csl` file. See *final*, the recruitment variability parameters in Section 7.4, and the future catch parameters in Section 7.11. There is one projections parameter in the `output.csl` file.

@n_projections	Number of projections to be done (from a point estimate)
Conditions	Only used when you are projecting from a point estimate (as opposed to a sample from the posterior)
Type	Integer
Default	300
Effects	Defines the number of projections to be done.
Notes	When projecting from a sample from the posterior, 1 projection is done for each sample point.

See also the output quantity commands in Section 9.2.

9.4 Defining yield calculations

@catch_split	Catch split used in yield simulations
Conditions	Only used in yield simulations, if there is more than one fishery.
Type	Constant vector
Effects	Defines a catch split used in yield simulations, as proportions of the total catch by fishery.

Notes	Order the proportions in the same order as the <code>annual.cycle.fishery_names</code> parameter in <code>population.csl</code> . If they do not sum to 1, they will be rescaled to sum to 1. (Check the output to make sure you specified this parameter correctly.)
@deterministic_yields_mortality_rate, @MCY_CAY_mortality_rate	Definition of the mortality rate F used in deterministic yield calculations and in CAY calculations
Conditions	Only used in yield simulations, if there is more than one fishery, and the Baranov catch equation is used.
Type	String
Effects	Defines the mortality rate F as either <code>exploitation_rate</code> or <code>instantaneous_mortality</code> .
Notes	If the above conditions are not met, F is always an exploitation rate.
@B_pre	Pre-fishery biomass B_{pre} block command
Conditions	Only used if the mortality rate in simulations is defined as an exploitation rate (either because there is more than one fishery, or because the Baranov catch equation is not used, or because <code>deterministic_yields_mortality_rate</code> or <code>MCY_CAY_mortality_rate</code> is defined as <code>exploitation_rate</code>).
Effects	Defines any following commands as <code>B_pre</code> subcommands.
Notes	B_{pre} is the pre-fishery biomass which is multiplied by the mortality rate to yield the catch, in yield simulations. Note spelling (<code>B_pre</code> and not <code>Bpre</code>).
mature_only	Pre-fishery biomass B_{pre} is mature fish only
Command	<code>B_pre</code>
Conditions	Only used if the mortality rate <code>deterministic_yields_mortality_rate</code> or <code>MCY_CAY_mortality_rate</code> is defined as <code>exploitation_rate</code> .
Type	Switch
Default	True (B_{pre} includes mature fish only)
Effects	Defines B_{pre} to include mature fish only.
area	Area for which pre-fishery biomass B_{pre} is calculated
Command	<code>B_pre</code>
Conditions	Only used if the mortality rate <code>deterministic_yields_mortality_rate</code> or <code>MCY_CAY_mortality_rate</code> is defined as <code>exploitation_rate</code> .
Type	String
Default	all areas combined
Effects	Defines B_{pre} to be calculated for a single area. Use an area label as in <code>area_names</code> .
step	Time step in which pre-fishery biomass B_{pre} is calculated
Command	<code>B_pre</code>
Conditions	Only used if the mortality rate <code>deterministic_yields_mortality_rate</code> or <code>MCY_CAY_mortality_rate</code> is defined as <code>exploitation_rate</code> .
Type	Integer
Default	time step of first fishery in the year
Effects	Defines the time step in which B_{pre} is calculated.
Notes	This time step must be earlier than the time step of any fishery, or it can be the time step of the first fishery if <code>proportion_mortality=0</code> .
proportion_mortality	Proportion of the time step's mortality after which pre-fishery biomass B_{pre} is calculated
Command	<code>B_pre</code>
Conditions	Only used if the mortality rate <code>deterministic_yields_mortality_rate</code> or <code>MCY_CAY_mortality_rate</code> is defined as <code>exploitation_rate</code> .

Type	Constant
Default	0
Effects	Defines the proportion of the mortality in the time step after which B_{pre} is calculated.
Notes	This has to be 0 unless <code>step</code> is earlier than the time step of any fishery.
selectivity	Selectivity ogive with which pre-fishery biomass B_{pre} is calculated
Command	<code>B_pre</code>
Conditions	Only used if the mortality rate <code>deterministic_yields_mortality_rate</code> or <code>MCY_CAY_mortality_rate</code> is defined as <code>exploitation_rate</code> .
Type	String
Default	(no selectivity)
Effects	Defines the selectivity ogive with which B_{pre} is calculated. Use a selectivity label from <code>selectivity_names</code> .

9.5 Defining deterministic yields

@per_recruit	Per-recruitment block command
Conditions	Only used in yield calculations.
Effects	Defines any following commands as <code>@per_recruit</code> subcommands.
do_YPR_SPR	Supply data to plot a YPR or SPR curve?
Command	<code>per_recruit</code>
Type	Switch
Default	False (do not supply data)
Effects	Print out F , YPR, and SPR for a list of mortality rates you supply through the <code>F</code> subcommand.
F	F's at which to calculate YPR & SPR
Command	<code>per_recruit</code>
Conditions	Only used if <code>do_YPR_SPR</code> is set.
Type	Constant vector
Effects	Defines the mortality rates F for which F , YPR, and SPR should be printed.
do_Fmax	Calculate F_{max}?
Command	<code>per_recruit</code>
Type	Switch
Default	False (do not calculate F_{max})
Effects	Print out F_{max}
do_F0.1	Calculate $F_{0.1}$?
Command	<code>per_recruit</code>
Type	Switch
Default	False (do not calculate $F_{0.1}$)
Effects	Print out $F_{0.1}$
Notes	Note spelling
do_Fx	Calculate $F_{x\%}$?
Command	<code>per_recruit</code>
Type	Switch
Default	False (do not calculate $F_{x\%}$)
Effects	Print out $F_{x\%}$
Notes	The value of x is supplied through the <code>x</code> subcommand.
x	x at which to calculate $F_{x\%}$
Command	<code>per_recruit</code>
Conditions	Only used if <code>do_Fx</code> is set.
Type	Constant vector
Effects	Defines the percentage x for which $F_{x\%}$ should be calculated.
Notes	Supply a percentage not a proportion.

guess	First guess at F_{max}, $F_{0.1}$, $F_{x\%}$
Command	per_recruit
Conditions	Only used if do_Fmax, do_F0_1, or do_Fx is set.
Type	Constant
Effects	Defines a guesstimate of F used by the minimiser in the calculation of F_{max} , $F_{0.1}$ and $F_{x\%}$.
@deterministic_MSY	Deterministic MSY block command
Conditions	Only used in yield calculations.
Effects	Defines any following commands as @deterministic_MSY subcommands.
do_MSY	Calculate deterministic MSY?
Command	deterministic_MSY
Type	Switch
Default	False (do not calculate MSY)
Effects	Print out deterministic MSY
do_yield_vs_SSB	Supply data to plot a yield versus SSB curve?
Command	deterministic_MSY
Type	Switch
Default	False (do not supply data)
Effects	Print out F , yield, and SSB for a list of mortality rates you supply through the F subcommand.
F	F's at which to calculate yield and SSB
Command	deterministic_MSY
Conditions	Only used if do_yield_vs_SSB is set.
Type	Constant vector
Effects	Defines the mortality rates F for which F, yield, and SSB should be printed.
guess	First guess at F_{MSYdet}
Command	deterministic_MSY
Conditions	Only used if do_MSY is set.
Type	Constant
Effects	Defines a guesstimate of F used by the minimiser in the calculation of F_{MSYdet} .

9.6 Defining stochastic yields

@MCY_CAY	MCY/CAY block command
Conditions	Only used in yield calculations.
Effects	Defines any following commands as @MCY_CAY subcommands.
do_MCY	Calculate MCY?
Command	MCY_CAY
Type	Switch
Default	False (do not calculate MCY)
Effects	Print out MCY
do_CAY	Calculate CAY?
Command	MCY_CAY
Type	Switch
Default	False (do not calculate CAY)
Effects	Print out CAY
interactive	Should MCY/CAY be calculated interactively?
Command	MCY_CAY
Conditions	Only used if do_MCY or do_CAY is set.
Type	Switch
Default	False (calculate MCY/CAY using the automated minimiser)

Effects	Calculate MCY and CAY interactively, i.e., you will be prompted for values of the harvest rate H .
p, q	Risk constraints in MCY/CAY analysis
Command	MCY_CAY
Conditions	Only used if do_MCY or do_CAY is set.
Type	2 x constant
Default	0.2, 0.1
Effects	Define the risk constraint parameters p and q . The stock is at risk if $SSB < pB_0$, but this can safely happen up to proportion q of the time.
Notes	p and q are proportions not percentages.
MCY_CAY_risk_year	Specify the year whose SSB will be used in the MCY/CAY risk constraint
Command	MCY_CAY
Type	Integer
Default	No year: instead use the unfished equilibrium SSB, B_0 , as the reference point
Effects	Specify the year whose SSBs will be used in the risk constraint for all MCY/CAY analysis. The constraint requires that the spawning stock biomass falls below $p \times SSB_{MCY_CAY_risk_year}$ less than $q \times 100\%$ of the time.
n_simulations	Number of simulations for each harvest rate H in MCY/CAY analyses (point-based only)
Command	MCY_CAY
Conditions	Only used if do_MCY or do_CAY is set, and you have supplied exactly one vector of free parameters using -i.
Type	Integer
Default	100
Effects	Define the number of point-based simulations to be done for each harvest rate H .
Notes	For sample-based simulations, one simulation is done for each sample from the posterior.
n_discard, n_keep	Number of years to discard and to keep in each MCY/CAY simulation.
Command	MCY_CAY
Conditions	Only used if do_MCY or do_CAY is set.
Type	2 x integer
Effects	Define the numbers of years to be kept and discarded in each MCY/CAY simulation.
Notes	Try several different values and see if it makes a difference.
max_upper_iter	Maximum number of times upper bound on H can be increased when searching for an optimal yield
Command	MCY_CAY
Conditions	Only used if do_MCY or do_CAY is set.
Type	Integer
Default	40
Effects	Defines the maximum number of times that the optimisation routine can consecutively increase the upper bound on H when searching for an optimum yield.
Notes	Normally the upper bound will only be increased repeatedly when the yield exhibits asymptotic behaviour.
MCY_uncertainty_dist	Distribution of uncertainty in virgin biomass (point-based only)
Command	MCY_CAY
Conditions	Only used if do_MCY is set, and you have supplied exactly one vector of free parameters using -i.
Type	String
Default	none
Effects	Define the distribution of uncertainty in virgin abundance in MCY simulations as none, normal, or lognormal.

Notes	Not used for sample-based simulations, where the posterior distribution should reflect the uncertainty in virgin biomass.
MCY_uncertainty_cv C.v. of uncertainty in virgin biomass (point-based only)	
Command	MCY_CAY
Conditions	Only used if uncertainty_dist is set to normal or lognormal.
Type	Constant
Default	0.2
Effects	Define the c.v. of uncertainty in virgin abundance in MCY simulations.
MCY_guess First guess at MCY	
Command	MCY_CAY
Conditions	Only used if do_MCY is set.
Type	Constant
Effects	Defines a guesstimate of MCY used by the minimiser.
F_CAY_guess First guess at F_{CAY}	
Command	MCY_CAY
Conditions	Only used if do_CAY is set.
Type	Constant
Effects	Defines a guesstimate of F _{CAY} used by the minimiser in the calculation of CAY.
@CSP CSP block command	
Conditions	Only used in yield calculations.
Effects	Defines any following commands as @CSP subcommands.
do_CSP Calculate CSP?	
Command	CSP
Type	Switch
Default	False (do not calculate CSP)
Effects	Print out CSP
individual_stocks Calculate CSP by individual stock?	
Command	CSP
Conditions	Only used if do_CSP is set, in a multi-stock model
Type	Switch
Default	False (calculate a single overall CSP)
Effects	Print out CSP by individual stock, rather than overall
CSP_guess First guess at CSP	
Command	CSP
Conditions	Only used if do_CSP is set.
Type	Constant
Effects	Defines a guesstimate of CSP used by the minimiser.
@B_post Post-fishery biomass B_{post} block command	
Conditions	Only used if do_CSP is set.
Effects	Defines any following commands as @B_post subcommands.
Notes	B _{post} is the post-fishery biomass whose expectation is the same in year current as in year current+1, if the catch in year current+1 is the CSP. Note spelling (B_post and not Bpost).
mature_only Post-fishery biomass B_{post} is mature fish only	
Command	B_post
Conditions	Only used if do_CSP is set.
Type	Switch
Default	True (B _{post} includes mature fish only)
Effects	Defines B _{post} to include mature fish only.
area Area for which post-fishery biomass B_{post} is calculated	
Command	B_post

Conditions	Only used if do_CSP is set.
Type	String
Default	All areas combined
Effects	Defines B_{post} to be calculated for a single area. Use an area label as in <i>area_names</i> .
step	Time step in which post-fishery biomass B_{post} is calculated
Command	B_post
Conditions	Only used if do_CSP is set.
Type	Integer
Default	Time step of first fishery in the year
Effects	Defines the time step in which B_{post} is calculated.
Notes	This time step must be earlier than the time step of any fishery, or it can be the time step of the first fishery if <i>proportion_mortality</i> =0.
proportion_mortality	Proportion of the time step's mortality after which post-fishery biomass B_{post} is calculated
Command	B_post
Conditions	Only used if do_CSP is set.
Type	Constant
Default	0
Effects	Defines the proportion of the mortality in the time step after which B_{post} is calculated.
Notes	This has to be 0 unless <i>step</i> is earlier than the time step of any fishery.
selectivity	Selectivity ogive with which post-fishery biomass B_{post} is calculated
Command	B_post
Conditions	Only used if do_CSP is set.
Type	String
Default	(No selectivity.)
Effects	Defines the selectivity ogive with which B_{post} is calculated. Use a selectivity label from <i>selectivity_names</i> .

10. TROUBLESHOOTING

10.1 Typical errors

CASAL is a complex system providing many opportunities for error — either because your parameter files do not correctly specify your model, or because the model you tried to specify does not work. When in doubt, ask an experienced user. Debugging versions of CASAL can be compiled that help to track down cryptic errors.

A short list of common errors that have been seen include;

Misspelt commands

If you misspell a command name in a parameter file, CASAL will not read it. You might just as well have commented out the line. There are two possible consequences. If the command is compulsory, CASAL will error out, complaining that it failed to find a parameter. If it is optional, CASAL may run but do the wrong thing. One way to diagnose misspelt commands is using `print.unused_parameters` command in the input data file `output.csl` (see Section 9.1). Note that commands are case sensitive.

No carriage return at the end of a parameter file

If you don't put a carriage return after your last command, then that last command will not be read by CASAL. You might just as well have commented out the last line. See above.

Misspelt arguments

Probably more likely to be picked up by CASAL than misspelt commands, but still a potential problem. Note that most arguments are case sensitive.

Commented out command

If you have the following lines in an input data file, for example,

```
@size_based 0
@fishery trawl
selectivity trawl_sel
...
```

and you comment out the `@fishery` line, then this will effect the `@size_based` command. CASAL will think `@size_based` is the beginning of a command block, because it now appears to have arguments (`selectivity`, etc.). Whereas in fact it is a single standalone command and should be treated as such. So the correct value of `size_based` will not be read. The fix for this one is to comment out or remove the entire `@fishery` block, and not just the first line.

No penalty on exceeding exploitation rate limits

Unless you use penalties, there is nothing to stop CASAL from coming up with a parameter estimate under which there is not enough fish for the catches to be taken. You probably want to use a catch limit penalty for each fishery (see Section 5.7.6).

No penalty on year class strengths not averaging to 1

If you estimate recruitment with the standard or Haist YCS parameterisation, you may want to force year class strengths to average to 1. You need to use a vector average penalty (see Section 5.7.6) to do this. If you don't, you may find that all your year class strengths are much more than 1 or much less than 1.

Proportions maturing specified as proportions mature

Be clear that the maturation rates supplied to CASAL (when maturity is a partition character) are proportions of fish which *become* mature at each age, not proportions which *are* mature at each age,

10.2 Other errors

When CASAL generates an error and the error message makes no sense to you, please let the CASAL development team know. Even if you manage to fix the problem yourself, the development team may be able to implement a more helpful error message and make life easier for the next person to encounter the problem. Especially let them know if the message says 'Betadiff error'.

Some parameter values of functions or ogives can result in either very large or very small numbers. These can, on occasion, generate internal numeric overflow errors within CASAL. The most common cause of a "Betadiff error" (especially when accompanied by the message "Betadiff returned a NaN gradient for parameters ...") is a calculation that has resulted in an overflow error.

CASAL does do some range checking of parameter values before attempting calculations, but there will be some instances where these range checks are not adequate. For example, in the logistic ogive, a very small value of a_{1095} (say, less than 0.2) can result in a very large value of $19^{(a_{50-x})/a_{1095}}$ for some values of x . Without range checking, this would cause an overflow error. The solution to this type of error is to impose bounds on parameters that exclude the possibility of an overflow error.

10.3 Reporting errors

If you wish to report a bug or problem with CASAL, then please send a bug report to CASAL@niwa.co.nz — after reading the guidelines below. Use the text "CASAL bug report" as the subject line in the email. Following the guidelines will assist the CASAL development team identify, reproduce, and hopefully solve any reported bugs. It is helpful to be as specific as possible when describing the problem. But before submitting a bug report, please check that you are using the most recent version of CASAL (see <http://www.niwa.co.nz/ncfa/tools/casal/> for details about the current release of CASAL).

Note that CASAL is distributed as unsupported software. NIWA does not usually provide help for users of CASAL outside of NIWA. However, the CASAL development team may provide assistance to reasonable requests, but while the development team would appreciate being notified of any problems or errors in CASAL they may not be able to provide timely solutions.

Guidelines for reporting a bug in CASAL

1. Detail the version of CASAL are you using? e.g., “CASAL v2.01-2003/08/01 Microsoft Windows executable”
2. What operating system/environment are you using? e.g., “IBM-PC running Microsoft Windows 2000 Release 3”.
3. Give a brief one-line description of the problem, e.g., “a segmentation fault was reported” or “a betadiff error was reported”.
4. If the problem is reproducible, please list the *exact* steps required to cause it, remembering to include the relevant CASAL input parameter files. Also specify the exact command line arguments that were used, e.g., “Using the command `casal -e -q -f my- > logfile.out` reports a segmentation fault. The files `my-population.csl`, `my-estimation.csl`, and `my-output.csl` are attached.”
5. If the problem is not reproducible (only happened once, or occasionally for no apparent reason), please describe the circumstances in which it occurred and the symptoms observed (but note it is much harder to reproduce and hence fix non-reproducible bugs, but if several reports are made over time that relate to the same thing, then this may help to track down the problem), e.g., “CASAL crashed, but I cannot reproduce how I did it. It seemed to be related to a local network crash but I cannot be sure.”
6. If the problem causes any error messages to appear, please give the exact text displayed, e.g., “Segmentation fault (core dumped)”
7. Remember to attach all relevant *input* and *output* files so that the problem can be reproduced (it can helpful to compress these into a single file). Without these, it may not be possible to determine the cause of the problem.

11. CASAL EXTENSIONS

Note: this section only applies if you have access to the casal source code. Source code, and hence CASAL extensions, are not available outside NIWA.

It may happen that you want to make a small change to CASAL, but don't want to have to delve into the source code. This section describes relatively simple ways of making the following common changes:

1. Changing the parameterisation used by the CASAL population section.
2. Applying new priors, or new penalties which depend only on the free parameters.
3. Applying new likelihoods.

These changes are implemented by writing snippets of C++ code, using templates supplied with CASAL (files `user.parameterisation.C`, `user.prior_penalty.C`, `user.likelihood.C` respectively). The source code is set up to make this as straightforward as possible. The CASAL source code is required to add such snippets, so that when you recompile CASAL, the new version will incorporate your changes.

To do this, you will need:

1. Access to the CASAL source code.
2. To be able to compile CASAL.
3. The ability to write a (simple) C++ program, with some basic use of `Betadiff` and of the Standard Template Library (STL) container classes.

If you use these CASAL extensions, remember that for someone else to replicate your results, they need your C++ code as well as your CASAL data files. Also, if you find yourself using a particular CASAL extension frequently, we recommend that you have it implemented and documented in CASAL (or at least document it separately).

11.1 New parameterisations (via `user.parameterisation.C`)

This section explains how to write your `user.parameterisation.C` to implement a different parameterisation in CASAL.

Why use a different parameterisation? It may turn out that the population parameters implemented in CASAL are not suited to the objective function minimisation or MCMC algorithms, and lead to poor convergence. For example, when fitting a von Bertalanffy curve to data, the standard parameterisation of k , t_0 , and L_{inf} is not ideal, because of the high correlations between the parameters. You can get better results by changing the parameterisation to t_0 , L_1 , and L_2 , where L_1 and L_2 are the lengths at two reference ages (see Smith et al. 2002). Similarly, the 2002 hoki assessment parameterised initial abundance in terms of \log_{BO_total} (the log of the sum of the B_0 values for the two stocks) and BO_prop_stock1 (the ratio of B_0 for the first of the two stocks, to the sum of the B_0 values for the two stocks), on the basis that this may lead to better convergence in the minimiser and MCMC.

To deal with this situation, we need to make a distinction between the original parameterisation used by the population section, and the new parameterisation which you want to use in the estimation section (and hence in the minimisation and MCMC algorithms). You can do this by writing your data files in terms of the new parameterisation and providing a snippet of code which converts the new parameters back to the original parameters. The CASAL estimation section works with the new parameters, but calls your code whenever it

invokes the population section, to translate your parameters back to the original parameterisation. In the first example above, we want the estimation section to work with the new parameterisation of t_0 , L_1 , and L_2 , while leaving the population section in terms of the familiar k , t_0 , and L_{inf} . We can do this by writing our data files in terms of t_0 , L_1 , and L_2 and providing a snippet of code which converts these parameters back to k , t_0 , and L_{inf} .

Unfortunately, you cannot use this method to reparameterise ogives. In CASAL, ogives are quite complicated objects (once you take ogive shifts and size-based ogives in age-based models into account): they are too complex to manipulate with this simple technique.

So, your `population.csl` file should include your new parameters, using the standard CASAL parameter syntax. For example,

```
@initialization
log_B0 11
```

The `@estimate` blocks in your `estimation.csl` file should also be in terms of the new parameters,

```
@estimate
parameter initialization.log_B0
lower_bound 8
upper_bound 16
prior uniform
```

You then need to write a `user.parameterisation.C` file which converts the new parameters to the old parameters. So, in the example above, you need to convert `initialization.log_B0` to `initialization.B0`. Write the program using the template provided.

1. The inputs are the names and values of the free parameters (under the new parameterisation). Scalar and vector parameters are supplied separately: ogive parameters are not supplied. Bear in mind that parameters which are not free, i.e., not currently being estimated, are not supplied.
2. The outputs are the names and values of the free population parameters (under the old parameterisation). You need to return scalars and vectors separately: you cannot return ogives. Be careful to check the type of each parameter (a 1-vector is different from a scalar).
3. Everything in between is up to you.

More detail is provided in the template file `user.parameterisation.C`.

Finally, you need to set `@user_parameterisation T` in your `estimation.csl` file. This lets CASAL know that it should call your `user_parameterisation()` function.

Here is a version of the code to do the $\log(B_0) - B_0$ conversion for a single-stock model (see `user.parameterisation.example.C`).

```
#####
template<CDVM>
void user_parameterisation(
    std::vector<std::string>& new_scalar_names,
    std::vector<DOUBLE>& new_scalar_vals,
    std::vector<std::string>& new_vector_names,
    std::vector<VECTOR>& new_vector_vals,
    std::vector<std::string>& old_scalar_names,
    std::vector<DOUBLE>& old_scalar_vals,
    std::vector<std::string>& old_vector_names,
    std::vector<VECTOR>& old_vector_vals){
/*
*** TEMPLATE - BRIAN BULL, 23/5/02 ***

Use this function to implement a new parameterisation in the CASAL
population section.

The function should convert the values of the 'new' parameters (i.e.,
those in your *.csl files) back to the 'old' parameters (i.e., those
in the CASAL manual).

The inputs to this function are the names and values of the new
scalar and vector parameters, the outputs are the names and values of
the old scalar and vector parameters.

You cannot reparameterise ogives.

You need to use the std::vector and std::string classes (these are in
the STL, see any good C++ book e.g., Stroustrup (2000) The C++
Programming Language, 3rd edition) and DOUBLE and VECTOR, which are
templates either for double and dvector or for dvariable and dvv
(apart from double, these are Betadiff classes: see betadiff.h)

Note that the indices of these VECTORS should always start at 1 (in
this particular function) but the indices of the std::vectors always
start at 0. So, the name of the first new vector parameter is
new_vector_names[0], and its values are new_vector_vals[0][1],
new_vector_vals[0][2]. The output arguments are passed as vectors of
length 0: you need to grow them and fill them in.
*/

/*
*** Brian Bull, 24/5/02. Converts log(B0) to B0 ***

Converts initialization.log_B0 to initialization.B0.

Need to insert the following in estimation.csl:
@user_parameterisation T
*/

DEBUG1("user_parameterisation");

if(!in<std::string>(new_scalar_names,"initialization.log_B0")
    {fatal("You need to estimate initialization.log_B0");}

DOUBLE log_B0 = new_scalar_vals[pos<std::string>
    (new_scalar_names,"initialization.log_B0")];
old_scalar_names.push_back("initialization.B0");
old_scalar_vals.push_back(exp(log_B0));
}
#####
```

11.2 New priors and penalties (via `user.prior_penalty.C`)

This section explains how to write your `user.prior_penalty.C` to add new priors, and new penalty functions on the free parameters, to CASAL.

Why would you want to add new priors and penalties on the free parameters? There are several possible reasons:

- You want to use some kind of penalty which is not allowed for in CASAL. For example, you want to use a penalty to encourage the value of q for one set of observations to be twice the value of q for another set of observations.
- You want to use some kind of prior which is not allowed for in CASAL. For example, you want a Student's t prior on M .
- You are using a user-defined parameterisation as per Section 11.1, and you want to apply a prior or penalty to the parameters using the old parameterisation. For example, you have reparameterised from B_0 to $\log(B_0)$, but you still want to apply the prior to B_0 .

The first two situations could be dealt with by changing the CASAL code, but you may not want to do this, for various reasons, because for example, (i) it would be difficult, (ii) because you think your changes would not be useful to other users, or (iii) because you just want to do a quick fix rather than a proper documented change. (The third would, in general, be genuinely hard to do by changing the main CASAL code.)

To add these objective function components, you need to do two things. First, add the `@user_components` command to your `estimation.csl`: the arguments are the text labels of the priors and penalties that you are going to supply. Second, write a `user.prior_penalty.C` which takes the free parameter values and calculates and returns the components. Write the program using the template provided.

- The inputs are:
 - (a) the names and values of the free parameters (i.e., those listed in your `estimation.csl`). Scalar, vector, and ogive parameters are supplied separately. (For ogives, only the estimable parameters are supplied, not the non-estimable parameters, like L and H in an `allvalues_bounded` ogive, and not the ogive values.)
 - (b) the names and values of the scalar and vector parameters calculated by `user.parameterisation.C` (if any).
- The outputs are the values of the new objective function components, and their labels, which should match those given in the `@user_components` command.
- Everything in between is up to you.

More detail is provided in the template file `user.prior_penalty.C`.

Here is a version of the code to implement the three examples above (see `user.prior_penalty.example.C`):

```
#####  
template<CDVM>  
void user_prior_penalty(  
    std::vector<std::string>& free_scalar_names,  
    std::vector<DOUBLE>& free_scalar_vals,
```

```

    std::vector<std::string>& free_vector_names,
    std::vector<VECTOR>& free_vector_vals,
    std::vector<std::string>& free_ogive_names,
    std::vector<VECTOR>& free_ogive_arguments,
    std::vector<std::string>& user_scalar_names,
    std::vector<DOUBLE>& user_scalar_vals,
    std::vector<std::string>& user_vector_names,
    std::vector<VECTOR>& user_vector_vals,
    std::vector<std::string>& objective_component_names,
    std::vector<DOUBLE>& objective_component_vals){
/*
*** TEMPLATE - BRIAN BULL, 23/5/02 ***

```

Use this function to implement new priors, and new penalties on the free parameters.

This function takes the names and values of the free parameters, and returns user-defined objective components (i.e., priors and penalties).

The inputs to this function are (a) the names and values of the free parameters (i.e., those in @estimate blocks in your estimation.csl). If there are free ogive parameters, only their estimable arguments are supplied - like a1, sL and sR for a double-normal ogive - not the actual values of the ogives, and (b) the names and values of the user-defined parameters calculated by user.parameterisation.C (if any) The outputs are the names and values of the user-defined objective components.

You need to use the std::vector and std::string classes (these are in the STL, see any good C++ book) and DOUBLE and VECTOR, which are templates either for double and dvector or for dvariable and dv (apart from double, these are Betadiff classes: see betadiff.h)

Note that the indices of these VECTORS should always start at 1 (in this particular function) but the indices of the std::vectors always start at 0. So, the name of the first free vector parameter is free_vector_names[0], and its values are free_vector_vals[0][1], free_vector_vals[0][2] ...

The output arguments are passed as vectors of length 0: you need to grow them and fill them in.

```

*/
/*
*** Brian Bull, 24/5/02. ***

```

- (1) Applies a penalty to encourage q[AEX].q to be approximately 2 x q[TAN].q.
- (2) Applies a t prior to natural_mortality.all
- (3) Applies a Cauchy prior to initialization.B0 (which has been reparameterised as initialization.log_B0).

Notes:

- Both the above q's must be ordinary free parameters (not nuisance parameters)
- natural_mortality.all must be free
- user.parameterisation.C converts initialization.log_B0 to initialization.B0
- Need to insert the following in estimation.csl:
- @user_components AEX_TAN_penalty prior_on_M prior_on_B0

```
*/  
  
DEBUG1("user_prior_penalty");  
  
if (!in<std::string>(free_scalar_names, "q[AEX].q")) {  
    fatal("You need to estimate q[AEX].q");  
}  
if (!in<std::string>(free_scalar_names, "q[TAN].q")) {  
    fatal("You need to estimate q[TAN].q");  
}  
DOUBLE qAEX = free_scalar_vals[pos<std::string>  
    (free_scalar_names, "q[AEX].q")];  
DOUBLE qTAN = free_scalar_vals[pos<std::string>  
    (free_scalar_names, "q[TAN].q")];  
objective_component_vals.push_back(5 * pow(qAEX-2*qTAN, 2));  
objective_component_names.push_back("AEX_TAN_penalty");  
  
if (!in<std::string>(free_scalar_names,  
    "natural_mortality.all")) {  
    fatal("You need to estimate natural_mortality.all");  
}  
DOUBLE M = free_scalar_vals[pos<std::string>  
    (free_scalar_names, "natural_mortality.all")];  
objective_component_vals.push_back(2.5 * log(1 + 0.25 *  
    pow((M - 0.25)/0.05, 2)));  
objective_component_names.push_back("prior_on_M");  
  
if (!in<std::string>(user_scalar_names, "initialization.B0")) {  
    fatal("You need to estimate initialization.B0 through  
    reparameterisation from initialization.log_B0");  
}  
DOUBLE B0 = user_scalar_vals[pos<std::string>  
    (user_scalar_names, "initialization.B0")];  
objective_component_vals.push_back(log(1+B0*B0));  
objective_component_names.push_back("prior_on_B0");  
}  
#####
```

11.3 New likelihoods (via user.likelihood.C)

This section explains how to write your user.likelihood.C to add new likelihood functions to CASAL.

Why add new likelihoods? You may want to use some kind of likelihood which is not allowed for in CASAL. For example, you might want a Student's t likelihood for an abundance series. This could be done by changing the CASAL code, but you may not want to do this, for various reasons (because it would be difficult, because you think your change would not be useful to other users, because you just want to do a quick fix rather than a proper documented change).

Note that you cannot use nuisance q 's for relative time series with user-defined likelihoods. Also, process error is ignored for user-defined likelihoods.

CASAL will not be able to calculate Pearson or normalised residuals for time series for which the likelihood is user-defined (they will be reported as zeros).

To add likelihoods, you need to do two things. First, put user-supplied as the argument to the dist subcommand in each observations block in estimation.csl for which you want to add a new likelihood. This means that CASAL will expect your user.likelihood.C file to provide a likelihood for that set of observations. Second,

write a `user.likelihood.C` file which calculates and returns the components. Use the template provided.

- The inputs are the observations and fits (each as a matrix) and the text label of the time series for which the likelihood is to be calculated.
- The output is a negative-log-likelihood for the time series.
- Everything in between is up to you.

More detail is provided in the template file `user.likelihood.C`. In particular, the format of the observations and fits matrices is described.

Here is a version of the code to calculate a Student's t likelihood for a relative abundance series (see `user.likelihood.example.C`).

```
#####
template<CDVM>
DOUBLE user_likelihood(std::string& label, const MATRIX& fits, const
MATRIX& observations){

/*
*** TEMPLATE - BRIAN BULL, 23/5/02 ***
```

Use this function to implement new likelihood functions.

This function takes the observations and fits, and returns the negative-log-likelihoods.

You need to set the 'dist' argument for each set of observations for which you are supplying a likelihood to 'user_supplied'.

You cannot use nuisance q 's for relative time series with user-defined likelihoods (use free q 's instead).

The inputs to this function are the text label of a time series and the observations and fits.

The obs and fits are stored as matrices. There is one row per year. For abundance data, there is one column. For all other data, there is one column per age or size class, per sex if they are sexed observations, with males before females. This is exactly the same as the format used by CASAL to print the data. The output is the corresponding negative-log-likelihood. If you want to use user-defined likelihoods for more than one time series, then you may need to put in an if-statement based on the 'label' argument, which tells you which time series is currently being dealt with.

You need to use the DOUBLE, VECTOR, and MATRIX classes which are templates either for double, dvector and dmatrix or for dvariable, dvv, and dvm (apart from double, these are Betadiff classes: see betadiff.h) Note that the indices of these VECTORS and MATRIXs should always start at 1 (in this particular function)

```
*/

/*
*** Brian Bull, 24/5/02. Calculates a Student's T likelihood for a
relative abundance index. ***
```

Notes: We use a t -likelihood with 4 degrees of freedom and a scale parameter of 20% of the mean (which means the c.v. is over 20%). The

q for this abundance index must be an ordinary free parameter, not a nuisance parameter.

*/

```
DEBUG1("user_likelihood");
```

```
if (label! = "acoustics"){  
    fatal("user_likelihood.C doesn't know how to calculate  
    a likelihood for " + label);}
```

```
DOUBLE result = 0;
```

```
for (int i = 1; i <= observations.rowmax(); i++){  
    result += log(fits[i][1])  
    + 2.5 * log(1 + 0.25*pow((observations[i][1]-  
    fits[i][1])/(0.2*fits[i][1]),2));
```

```
}
```

```
return result;
```

```
}
```

```
#####
```


12. POST-PROCESSING OF CASAL OUTPUT

12.1 Introduction

A set of S/S-Plus/R functions are available for reading CASAL output. These are contained in the file `extract_CASAL.v2.01.s`

Key functions include `extract.header`, `extract.objective.function`, `extract.fits`, `extract.free.parameters`, `extract.quantities`, and `read.files.for.BOA`. These functions allow the importation of standard output from `casal -r`, `-e` or `-E` into S-Plus for post-processing. When running CASAL, redirect the standard output into a text file and pass the name of the text file to the S/S-Plus/R functions.

Two other functions may be useful. These are `extract.free.parameters.from.table`, and `extract.quantities.from.table`. The former is used to extract free parameters from files in the flat format described in Section 2.3 (this is the format used to pass free parameters to CASAL with `-i`, and is also the output format for samples from a posterior using `-m` or `-C`). The latter is used to extract output quantities from files in tabular format (this is the output format of `-v` and `-P`).

All these functions are documented below.

12.2 `extract.header`

Extracts the header file information from a CASAL standard output file.

DESCRIPTION

Creates an S/S-Plus/R object representing header file information contained in a CASAL standard output file.

USAGE

```
extract.header(file)
```

REQUIRED ARGUMENTS

`file` character string giving the file name where the data is to be retrieved.

VALUE

A list object.

`extract.header` returns a partial list of the contents of the file that describe the header information. This includes CASAL function call, run date and time, CASAL version information, user and machine names.

<code>call</code>	CASAL function call and command line parameters
<code>date</code>	Creation date time of the output file
<code>sources</code>	Version identification information
<code>user</code>	User (login) name
<code>machine</code>	Machine (computer) name

WARNING

The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

```
extract.objective.function,  
extract.free.parameters, extract.fits,  
extract.quantities,  
extract.free.parameters.from.table,  
extract.quantities.from.table
```

EXAMPLES

```
a<-extract.header("abc")  
# store header information from the CASAL output file "abc"  
# as the S object a
```

12.3 extract.objective.function

Extracts the objective function information from a CASAL standard output file

DESCRIPTION

Creates an S/S-Plus/R object representing objective function information contained in a CASAL standard output file.

USAGE

```
extract.objective.function(file)
```

REQUIRED ARGUMENTS

`file` character string giving the file name where the data is to be retrieved.

VALUE

A list object.

`extract.objective.function` returns a partial list of the contents of the file that describe the objective function information. This includes objective function value and components.

`value` Objective function value

`components` data.frame of the objective function component labels and the component values

WARNING

The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

```
extract.header, extract.free.parameters,  
extract.fits, extract.quantities,  
extract.free.parameters.from.table,  
extract.quantities.from.table
```

EXAMPLES

```
a<-extract.objective.function("abc")  
# store objective function information from the CASAL output file "abc"  
# as the S object a
```

12.4 `extract.free.parameters`

Extracts free parameters information from a CASAL standard output file

DESCRIPTION

Creates an S/S-Plus/R object representing free parameter information contained in a CASAL standard output file.

USAGE

```
extract.free.parameters(file)
```

REQUIRED ARGUMENTS

`file` character string giving the file name where the data is to be retrieved.

VALUE

A list object.

`extract.free.parameters` returns a partial list of the contents of the file that describe the free parameter information. This would typically include the initialization.B0 value, selectivity estimates, and year class strength estimates. List elements are named by the text string used to define each free parameter in CASAL.

WARNING

Illegal S names can be used to name free parameters in CASAL. Such names are converted to legal S names by enclosing the string in quotes. The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

```
extract.header, extract.objective.function,
extract.fits, extract.quantities,
extract.free.parameters.from.table,
extract.quantities.from.table
```

EXAMPLES

```
a<-extract.free.parameters("abc")
# store free parameter information from the CASAL output file "abc"
# as the S object a
```

12.5 `extract.fits`

Extracts the observation and associated fit & residual information from a CASAL standard output file

DESCRIPTION

Creates an S/S-Plus/R object representing the observations, fits, and residual information contained in a CASAL standard output file.

USAGE

```
extract.fits(file)
```

REQUIRED ARGUMENTS

`file` character string giving the file name where the data is to be retrieved.

VALUE

A list object.

`extract.fits` returns a partial list of the contents of the file that describe the observation, fit, and residual information. Each element of the object represents a single time series of observations from the CASAL input files. List elements are named by the text string used to define time series in CASAL. Each element has sub-elements as follows:

`year` Vector of years that identify the years of each of the observations

`obs` Vector or matrix of observation data

`fits` Vector or matrix of associated fitted value for each observation

`resids` Vector or matrix of associated residuals for each observation

`"pearson_resids"` Vector or matrix of associated Pearson residuals for each observation

`"normalised_resids"` Vector or matrix of associated normalised residuals for each observation

WARNING

Illegal S names can be used to name quantities in CASAL. Such names are converted to legal S names by enclosing the string in quotes. The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

`extract.header`, `extract.objective.function`,
`extract.free.parameters`, `extract.quantities`,
`extract.free.parameters.from.table`,
`extract.quantities.from.table`

EXAMPLES

```
a<-extract.fits("abc")
# store time series observations, fits & residuals from the CASAL
# output file "abc" as the S object a
```

12.6 `extract.quantities`

Extracts the estimated quantities from a CASAL standard output file

DESCRIPTION

Creates an S/S-Plus/R object representing the estimated quantities contained in a CASAL standard output file.

USAGE

```
extract.quantities(file)
```

REQUIRED ARGUMENTS

`file` character string giving the file name where the data is to be retrieved.

VALUE

A list object.

`extract.quantities` returns a partial list of the contents of the file that describe the estimated free quantities. This would typically include scalar and vector parameter values, ogives, q 's, B_0 , R_0 , and spawning stock biomass estimates. Each element of the object represents a single quantity from the CASAL output files. These elements can include, but are not limited to, the following:

- "Scalar parameter values" List object describing the values of each of the scalar parameter values.
- "Vector parameter values" List object describing the values of each of the vector parameter values.
- "Vector ogive arguments" List object describing the values of each of the vector ogive values.
- "Nuisance q 's" List object describing the values of each of the nuisance q values.

B_0 Estimated value of B_0

R_0 Estimated value of R_0

$SSBs$ List object with vector elements

$SSBs\$SSB$ Vector of spawning stock biomass estimates

$SSBs\$year$ Vector of years associated with the spawning stock biomass estimates

"Actual_catches" $\$catch$ Vector of recorded catch values

"Actual_catches" $\$year$ Vector of years associated with the catch

WARNING

Illegal S names can be used to name free parameters in CASAL. Such names are converted to legal S names by enclosing the string in quotes. The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

`extract.header`, `extract.objective.function`,
`extract.free.parameters`, `extract.fits`,
`extract.free.parameters.from.table`,
`extract.quantities.from.table`

EXAMPLES

```
a<-extract.fits("abc")
# store time series observations, fits, and residuals from the CASAL
# output file "abc" as the S object a
```

12.7 extract.free.parameters.from.table

Extracts the free parameters from a CASAL free parameter flat file

DESCRIPTION

Creates an S/S -Plus/ R object representing the output quantities contained in a CASAL free parameter flat file.

USAGE

```
extract.free.parameters.from.table(file)
```

REQUIRED ARGUMENTS

`file` character string giving the file name where the data is to be retrieved.

VALUE

A data.frame object, with one column per free parameter.

WARNING

Illegal S names can be used to name output quantities in CASAL. Such names are converted to legal S names by enclosing the string in quotes. The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

```
extract.header, extract.objective.function,  
extract.free.parameters, extract.fits,  
extract.quantities.from.table
```

EXAMPLES

```
a<-extract.free.parameters.from.table("abc")  
# store free parameters from the CASAL output file "abc" as the S  
# object a
```

12.8 extract.quantities.from.table

Extracts the output quantities from a CASAL output file in tabular format.

DESCRIPTION

Creates an S/S-Plus/R object representing the output quantities contained in a CASAL tabular format file.

USAGE

```
extract.quantities.from.table(file)
```

REQUIRED ARGUMENTS

`file` character string giving the file name where the data is to be retrieved.

VALUE

A data.frame object, with one column per output quantity.

WARNING

Illegal S names can be used to name output quantities in CASAL. Such names are converted to legal S names by enclosing the string in quotes. The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

```
extract.header, extract.objective.function,
extract.free.parameters, extract.fits,
extract.free.parameters.from.table
```

EXAMPLES

```
a<-extract.quantities.from.table("abc")
# store output quantities from the CASAL output file "abc" as the S
# object a
```

12.9 read.files.for.BOA

Extracts the output quantities from a CASAL MCMC output file in a format suitable for use with the S-Plus/R library BOA (Bayesian Output Analysis program, version 1.0)

DESCRIPTION

Creates an S/S-Plus/R object representing the output quantities contained in the CASAL MCMC output file. Use `casal -m` to generate the MCMC output file suitable for use with `read.files.for.BOA`

USAGE

```
read.files.for.BOA(samples.file,objectives.file,
burn.in = 0)
```

REQUIRED ARGUMENTS

`samples.file` character string giving the file name where the MCMC chain data is to be retrieved.
`objectives.file` character string giving the file name where the objective function data is to be retrieved.
`burn.in` number specifying the number of samples at the beginning of the `samples.file` that should be considered to be samples from the "burn-in" phase of the MCMC run.

VALUE

A data.frame object, with one row per output quantity.

WARNING

Illegal S names can be used to name output quantities in CASAL. Such names are converted to legal S names by enclosing the string in quotes. The exact format of the object on the file is subject to change. No error checking is undertaken to ensure that the data in the S object accurately represents the data in the CASAL file. Note also that exact equality of read-in numeric data is machine specific.

SEE ALSO

The add-on library/package BOA

EXAMPLES

```
a<-read.files.for.BOA("samples.1",
"objectives.1", burn.in = 0)
# store MCMC chain quantities from the CASAL output file "samples.1"
# as the S object a
```


13. AN EXAMPLE OF AN AGE-BASED MODEL

13.1 Introduction

The example provided here is a simple age-based model, loosely based on the Chatham Rise hake fishery in New Zealand. (Note that we do not suggest that this is a good model, but rather it is provided as an example of how CASAL can be used.) The biomass observations consist of a CPUE series and a trawl survey biomass series. In addition, the trawl survey series has a series of associated proportions-at-age data. Proportions-at-age data are also available from the commercial catch.

The partition is defined to allow for age-classes in the model from 2–25. Note that there is no plus group, so the model assumes that there are no fish older than 25 in the population. The partition is defined for males and females, but does not include maturity.

The annual cycle is described in the `population.csl` file, and consists of a year round fishery (time step 1), followed by an instantaneous spawning period (time step 2), and then an age increment (time step 3).

An example of the partition at time step 1 (i.e., before the arrival of new recruits) as printed by CASAL is,

```
Partition:
Sex      2      3      4      5  ...      23      24      25
Male    0  915.9  676.2  246.53  ...  0.5768  0.4575  0.3662
Female  0  915.9  700.7  250.03  ...  1.981   1.618   1.331
```

The remainder of the `population.csl` file defines the recruitment processes, maturation, natural mortality, fishing mortality (and catches), selectivities to be applied to the fishing and trawl surveys, size-at-age, and size-weight relationships. Note that your results may look slightly different, depending on the computer and platform used to run CASAL.

13.2 The input parameter files

The input parameter files (named with a prefix of R1- to identify them as the first ‘run’) are;

R1-population.csl

```
#INITIALSATION (the starting value for B0 is set as 45000 t)
@initialization
B0 45000

# PARTITION
@size_based false # Define the model as age-based
@min_age 2
@max_age 25 # The partition keeps account of fish aged 2-25
@plus_group false # and excludes all fish over the age of 25
@sex_partition true # The model is sex-based
@mature_partition false # Maturity is excluded from the partition
@n_areas 1 # Only a single fishing area is defined
@area_names chat # with the (optional in a single area model)
# label "chat"
@n_stocks 1 # This is a single stock model
@stock_names HAK4 # and the stock has the (optional in a single
# stock model) name "HAK4"
```

```

# TIME SEQUENCE
@initial 1975      # The model is defined to run from 1975
@current 2002     # to the current year, 2002
@final 2007       # Projections are run up to the year 2007
@annual_cycle
time_steps 3      # There are three time steps: Oct-Aug, Sep, Sep
recruitment_time 2 # Recruitment occurs in time step 2
recruitment_areas chat # in the area "chat" (the only area defined)
spawning_time 2   # Spawning occurs in time step 2
spawning_part_mort 0.5 # and SSBs are calculated after spawning fish have
                    # undergone 0.5 of the mortality assigned to this time
                    # step
spawning_areas chat # Spawning occurs in the area "chat"
spawning_ps 1     # and all mature fish spawn
aging_time 3      # Age incrementation occurs in time step 3
growth_props 1.00 1.00 0.00 # All fish growth occurs in time step 1
M_props 1.00 0.00 0.00 # All natural mortality occurs in time step 1
baranov false     # The baranov equation is not used
midmortality_partition weighted_sum
fishery_names chatFishery # The fishery has the label "chatFishery"
fishery_times 1     # and occurs in the first time step
fishery_areas chat # in the area labelled "chat"
n_migrations 0     # No migrations are defined

# RECRUITMENT
@y_enter 2         # Recruits enter at age 2
@standardise_YCS true # Use the "Haist" parameterisation of YCS
@recruitment      # the two following lines define the starting values for
                  # recruitment for the years 1973-2000
YCS_years 1973 1974 1975 1976 1977 1978 1979 1980 1981 1982 1983 1984 1985
          1986 1987 1988 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999
          2000
YCS 1.00 1.00 1.00 1.00 1.00 1.00 1.00 1.04 1.21 0.85 1.38 1.03 1.25 0.89
     1.47 1.40 1.84 2.66 2.28 2.17 1.70 1.59 0.71 1.00 1.00 1.00 1.00 1.00
p_male 0.5        # 50% of 'recruits' are males
first_free 1979  # with standardisation occurring over the years
last_free 1998   # 1979-1998
sigma_r 0.6      # Standard deviation of YCS for projections
SR BH           # Use the Beverton-Holt stock-recruit relationship
steepness 0.9    # with a steepness parameter of 0.9

# RECRUITMENT VARIABILITY
@randomisation_method lognormal # Use the lognormal distribution when
                                # assigning YCS to unknown years during projections
@first_random_year 1999 # Defines the first unknown YCS as 1999

#MATURATION
@maturity_props # Define the maturity ogive for males and females
male allvalues_bounded 2 10 0.00 0.00 0.02 0.07 0.31 0.78 1.00 1.00 1.00
female allvalues_bounded 2 10 0.00 0.00 0.02 0.04 0.07 0.45 0.86 1.00 1.00

# NATURAL MORTALITY
@natural_mortality
avg 0.20 # Define the average natural mortality of males & females as 0.20
diff 0.02 # and define the difference (female-male) as 0.02,
          # i.e., male_M=0.21 and female_M=0.19

# FISHING
@fishery chatFishery # Define the catch from the chatFishery for years 1975-
                    2002
years 1975 1976 1977 1978 1979 1980 1981 1982 1983 1984 1985 1986 1987 1988
      1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002
catches 191 488 1288 34 609 750 997 596 302 344 544 362 509 574 804 977 991
        2454 2775 2898 4094 3760 3761 3673 3524 3700 3700 3700
selectivity chatFsel # Defines that the catch is removed from the population
                    # using the selectivity defined by the label "chatFsel"
U_max 0.4          # with a maximum possible exploitation rate of 0.4
future_years 2003 2004 2005 2006 2007 # Defines the future years and

```

```

future_catches 3700 3700 3700 3700 3700 # catches for use in projections

# SELECTIVITIES
@selectivity_names chatTANsel chatFsel # Define the two selectivities used
@selectivity chatFsel # Fishing selectivity is a logistic
male logistic 9 4 # curve for males and females, with
female logistic_capped 9 4 0.7 # the female rates relative to males
@selectivity chatTANsel # And similarly for trawl survey
male logistic 9 4 # selectivity
female logistic_capped 9 4 0.7

# SIZE AT AGE
@size_at_age_type von_Bert # Defines that the age-length relationship is
@size_at_age_dist normal # von-Bertalanffy (defined separately for males
@size_at_age # and females), with a distribution defined
k_male 0.277 # as normal with c.v.=0.1
t0_male -0.11
Linf_male 90.3
cv_male 0.1
k_female 0.202
t0_female -0.20
Linf_female 113.4
cv_female 0.1

# SIZE-WEIGHT
@size_weight # Defines the length-weight relationship
a_male 2.49e-9
b_male 3.234
a_female 1.70e-9
b_female 3.328
verify_size_weight 50 0.5 1.5 # Verify that a 50 cm fish has a
# weight between 0.5 and 1.5 kgs

```

R1-estimation.csl

```

# ESTIMATION
@estimator Bayes # Use the Bayes estimation method
@max_iters 300 # With maximum of 300 iterations for the point estimates
@max_evals 1000 # and 1000 function evaluations
@grad_tol 0.002 # Set the tolerance for the convergence test at 0.002
@MCMC
start 0 # Start the MCMC at 0
length 1000000 # and evaluate for 1 million steps
keep 1000 # keeping every 1000th sample
stepsize 0.02 # with the stepsize for the MCMC set at 0.02
adaptive_stepsize true # but adapt the stepsize during the evaluation
adapt_at 50000 100000 # after the 50000 and 100000 step
burn_in 100 # The MCMC has a burn-in period of 100*1000=100000 steps

# OBSERVATIONS Chatham Rise
@relative_abundance chatCPUE # Define a relative abundance series "chatCPUE"
biomass true # This time series is an abundance index
q chatCPUEq # and has a relativity constant called "chatCPUEq"
years 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 # index years
step 1 # Occurs in time step 1
proportion_mortality 0.5 # after 0.5 of mortality has been recorded
area chat # Occurs in the area called "chat"
ogive chatFsel # and is applied with the selectivity "chatFsel"
1992 1.50 # The values of the index are ...
1993 1.10
1994 0.93
1995 1.33
1996 1.53
1997 0.90
1998 0.68
1999 0.75
2000 0.57

```

```

2001 1.23
cv_1992 0.35 # and each point has the c.v.s ...
cv_1993 0.35
cv_1994 0.35
cv_1995 0.35
cv_1996 0.35
cv_1997 0.35
cv_1998 0.35
cv_1999 0.35
cv_2000 0.35
cv_2001 0.35
dist lognormal # where the c.v.s have lognormal distribution
cv_process_error 0.0 # and there is no process error applied

@relative_abundance chatTANbiomass # Define a relative abundance series
# "chatTANbiomass"
biomass true # This time series is an abundance index
q chatTANq # and has a relativity constant called "chatTANq"
years 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002 # index years
step 1 # Occurs in time step 1
proportion_mortality 1 # after all mortality has been recorded
area chat # Occurs in the area "chat"
ogive chatTANsel # and is applied with the selectivity "chatTANsel"
1992 4180 # the values of the index are ...
1993 2950
1994 3353
1995 3303
1996 2457
1997 2811
1998 2873
1999 2302
2000 2090
2001 1589
2002 1567
cv_1992 0.15 # and each point has the c.v.s ...
cv_1993 0.17
cv_1994 0.10
cv_1995 0.23
cv_1996 0.13
cv_1997 0.17
cv_1998 0.18
cv_1999 0.12
cv_2000 0.09
cv_2001 0.13
cv_2002 0.15
dist lognormal # where the c.v.s have lognormal distribution
cv_process_error 0.0 # and there is no process error applied

@proportions_at chatTANage # Define a series of relative proportions-at-age
years 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002 # index years
step 1 # Occurs in time step 1
proportion_mortality 1 # after all mortality has been recorded
area chat # in area "chat"
sexed true # The observations are recorded by sex
plus_group true # The oldest age group is a 'plus-group'
sum_to_one true # and the proportions sum to one over each year
min_class 3 3 # The minimum age class is 3 for males and 3 for females
max_class 15 15 # The maximum/plus-group age class is 15 for males and 15
# for females
ogive chatTANsel # And is applied with the selectivity "chatTANsel"
# M3 M4 M5 M6 M7 M8 M9 M10 M11 M12 M13 M14 M15+ F3 F4 F5 F6 F7 F8 F9 F10 F11
# F12 F13 F14 F15+ are the proportions-at-age in each column
1992 0.0186 0.0219 0.0249 0.0390 0.0512 0.0646 0.0422 0.0677 0.0523 0.0687
0.0299 0.0132 0.0515 0.0055 0.0254 0.0199 0.0320 0.0268 0.0394 0.0250
0.0536 0.0346 0.0423 0.0489 0.0304 0.0705
1993 0.0449 0.0346 0.0372 0.0112 0.0286 0.0220 0.0279 0.0156 0.0310 0.0300
0.0690 0.0283 0.0614 0.0552 0.0231 0.0457 0.0301 0.0296 0.0426 0.0122
0.0415 0.0277 0.0541 0.0697 0.0127 0.1142

```

```

1994 0.1063 0.0405 0.0431 0.0274 0.0179 0.0182 0.0170 0.0198 0.0317 0.0252
      0.0127 0.0195 0.0969 0.1176 0.0483 0.0421 0.0513 0.0271 0.0350 0.0095
      0.0154 0.0060 0.0256 0.0083 0.0303 0.1072
1995 0.0933 0.0730 0.0365 0.0123 0.0379 0.0321 0.0195 0.0149 0.0158 0.0314
      0.0322 0.0544 0.0329 0.0788 0.0609 0.0501 0.0368 0.0382 0.0288 0.0034
      0.0476 0.0218 0.0231 0.0197 0.0390 0.0659
1996 0.0528 0.0505 0.0904 0.0485 0.0226 0.0323 0.0064 0.0149 0.0037 0.0039
      0.0046 0.0111 0.0265 0.0689 0.1471 0.1374 0.0786 0.0450 0.0260 0.0135
      0.0282 0.0032 0.0168 0.0286 0.0133 0.0251
1997 0.0942 0.0797 0.0590 0.0498 0.0318 0.0528 0.0045 0.0187 0.0091 0.0152
      0.0093 0.0230 0.0344 0.1124 0.0682 0.0927 0.0814 0.0225 0.0233 0.0180
      0.0073 0.0026 0.0027 0.0297 0.0121 0.0457
1998 0.0397 0.0678 0.0862 0.0457 0.0676 0.0354 0.0201 0.0225 0.0092 0.0176
      0.0066 0.0260 0.0422 0.0195 0.0606 0.0660 0.0831 0.0711 0.0527 0.0291
      0.0170 0.0362 0.0225 0.0095 0.0049 0.0411
1999 0.0683 0.0771 0.0408 0.0364 0.0228 0.0380 0.0148 0.0226 0.0138 0.0109
      0.0045 0.0050 0.0585 0.0628 0.0307 0.0711 0.0411 0.0372 0.0740 0.0521
      0.0465 0.0232 0.0270 0.0180 0.0152 0.0876
2000 0.0623 0.0466 0.0521 0.0292 0.0369 0.0524 0.0508 0.0414 0.0385 0.0138
      0.0120 0.0227 0.0234 0.0131 0.0358 0.0336 0.0433 0.0445 0.0699 0.0413
      0.0265 0.0298 0.0368 0.0187 0.0370 0.0878
2001 0.0033 0.0274 0.0554 0.0259 0.0455 0.0611 0.0413 0.0404 0.0337 0.0204
      0.0124 0.0034 0.0195 0.0064 0.0314 0.0278 0.0364 0.0983 0.0549 0.0798
      0.0681 0.0728 0.0488 0.0076 0.0210 0.0567
2002 0.0173 0.0193 0.0241 0.0346 0.0365 0.0657 0.0427 0.0667 0.0326 0.0307
      0.0272 0.0141 0.0319 0.0353 0.0249 0.0146 0.0133 0.0547 0.0488 0.0745
      0.0660 0.0750 0.0646 0.0304 0.0147 0.0399
# with c.v.s for each observation defined as
cvs_1992 0.710 0.469 0.370 0.350 0.419 0.413 0.373 0.383 0.365 0.301 0.393
          0.518 0.302 0.689 0.326 0.394 0.313 0.305 0.322 0.336 0.301 0.308 0.289
          0.276 0.360 0.228
cvs_1993 0.321 0.336 0.361 0.861 0.421 0.532 0.408 0.555 0.409 0.489 0.361
          0.527 0.369 0.346 0.451 0.370 0.428 0.437 0.376 0.780 0.336 0.494 0.366
          0.342 0.628 0.307
cvs_1994 0.268 0.379 0.315 0.492 0.555 0.507 0.587 0.546 0.468 0.510 0.818
          0.547 0.268 0.238 0.290 0.315 0.278 0.405 0.382 0.782 0.640 0.838 0.457
          0.742 0.374 0.188
cvs_1995 0.250 0.367 0.372 0.621 0.655 0.810 0.924 0.772 0.865 0.713 0.632
          0.608 0.612 0.254 0.329 0.340 0.400 0.379 0.587 1.028 0.432 0.542 0.499
          0.506 0.475 0.286
cvs_1996 0.497 0.445 0.353 0.409 0.729 0.560 1.161 0.929 1.481 1.128 1.335
          1.228 0.687 0.392 0.227 0.256 0.317 0.381 0.566 0.716 0.542 3.000 0.881
          0.568 0.930 0.558
cvs_1997 0.298 0.315 0.430 0.441 0.612 0.521 1.504 0.736 0.994 0.867 1.033
          0.668 0.607 0.289 0.313 0.279 0.301 0.631 0.518 0.695 1.059 3.000 1.751
          0.556 1.051 0.458
cvs_1998 0.464 0.403 0.343 0.440 0.423 0.575 0.710 0.703 1.227 0.774 1.503
          0.691 0.445 0.533 0.342 0.341 0.324 0.358 0.357 0.548 0.689 0.511 0.689
          1.097 1.381 0.493
cvs_1999 0.572 0.417 0.454 0.510 0.577 0.521 0.834 0.782 0.808 1.004 1.373
          1.170 0.457 0.458 0.550 0.314 0.440 0.448 0.324 0.335 0.450 0.690 0.588
          0.813 0.749 0.426
cvs_2000 0.414 0.522 0.328 0.399 0.351 0.319 0.320 0.328 0.374 0.631 0.613
          0.589 0.419 1.334 0.569 0.452 0.403 0.371 0.294 0.343 0.427 0.488 0.389
          0.593 0.416 0.229
cvs_2001 1.726 0.527 0.446 0.510 0.510 0.392 0.462 0.442 0.551 0.643 0.761
          1.439 0.637 1.180 0.434 0.552 0.445 0.301 0.429 0.352 0.353 0.368 0.441
          0.767 0.610 0.362
cvs_2002 1.091 0.770 0.539 0.421 0.412 0.297 0.367 0.322 0.391 0.510 0.523
          0.734 0.481 0.612 0.643 0.756 0.772 0.399 0.369 0.331 0.306 0.304 0.309
          0.461 0.752 0.423
dist lognormal # where the c.v.s have lognormal distribution
cv_process_error 0.0 # and there is no process error applied
ageing_error true # apply an ageing error model to the observations

@catch_at chatOBS # Define a series of proportions-at-age from the catch
years 1998 1999 2000 2001 # for the years 1998-2001
fishery chatFishery # that occur in the fishery "chatFishery"

```

```

sexed true          # The observations are recorded by sex
plus_group true    # The oldest age group is a 'plus-group'
sum_to_one true    # and the proportions sum to one over each year
min_class 3 3      # The minimum age class is 3 for males and 3 for females
max_class 15 15    # The maximum/plus-group age class is 15 for males and 15
                  # for females
# M3 M4 M5 M6 M7 M8 M9 M10 M11 M12 M13 M14 M15+ F3 F4 F5 F6 F7 F8 F9 F10 F11
  F12 F13 F14 F15+ are the proportions-at-age in each column
1998 0.1079 0.0696 0.0580 0.0607 0.0865 0.0706 0.0288 0.0247 0.0062 0.0077
     0.0076 0.0070 0.0115 0.0905 0.0581 0.0608 0.0373 0.0427 0.0548 0.0241
     0.0247 0.0245 0.0105 0.0063 0.0036 0.0152
1999 0.0264 0.0641 0.0445 0.0714 0.0413 0.0516 0.0329 0.0271 0.0270 0.0117
     0.0023 0.0021 0.0209 0.0229 0.0690 0.0485 0.0913 0.0563 0.0537 0.0594
     0.0517 0.0412 0.0133 0.0157 0.0137 0.0401
2000 0.0161 0.0441 0.0605 0.0509 0.0658 0.0590 0.0715 0.0432 0.0291 0.0154
     0.0116 0.0051 0.0180 0.0129 0.0405 0.0315 0.0428 0.0766 0.1011 0.0573
     0.0309 0.0436 0.0248 0.0071 0.0060 0.0346
2001 0.0087 0.0280 0.0422 0.0427 0.0849 0.0887 0.0788 0.0711 0.0566 0.0275
     0.0162 0.0166 0.0507 0.0019 0.0383 0.0246 0.0332 0.0786 0.0594 0.0345
     0.0295 0.0240 0.0219 0.0120 0.0062 0.0233
# with c.v.s for each observation defined as
cvs_1998 0.175 0.232 0.227 0.207 0.173 0.172 0.278 0.299 0.515 0.487 0.440
         0.494 0.307 0.183 0.209 0.228 0.269 0.234 0.184 0.296 0.282 0.295 0.459
         0.402 0.746 0.275
cvs_1999 0.328 0.221 0.313 0.250 0.269 0.283 0.270 0.361 0.380 0.481 1.162
         1.189 0.409 0.332 0.217 0.329 0.196 0.257 0.235 0.222 0.244 0.250 0.421
         0.583 0.512 0.259
cvs_2000 0.495 0.264 0.237 0.233 0.196 0.219 0.203 0.238 0.349 0.426 0.458
         0.714 0.343 0.424 0.299 0.404 0.252 0.188 0.146 0.208 0.275 0.217 0.294
         0.424 0.514 0.184
cvs_2001 0.383 0.332 0.311 0.300 0.192 0.200 0.227 0.245 0.268 0.372 0.551
         0.602 0.257 1.445 0.257 0.333 0.305 0.191 0.195 0.211 0.255 0.265 0.295
         0.356 0.481 0.239
dist lognormal      # where the c.v.s have lognormal distribution
cv_process_error 0.2 # and there is a c.v.=0.2 process error applied
ageing_error true   # Apply an ageing error model to the observations

# RELATIVITY CONSTANTS
@q_method nuisance # Use the "nuisance" method for estimating q

@estimate
parameter q[chatCPUEq].q # Estimate the parameter q[chatCPUEq].q when
                        # fitting the model
lower_bound 1e-6        # with a lower bound
upper_bound 10          # and upper bound
prior uniform-log       # and use a uniform-log prior

@estimate
parameter q[chatTANq].q # Estimate the parameter q[chatTANq].q when
                        # fitting the model
lower_bound 1e-8
upper_bound 1
prior uniform-log

#FREE PARAMETERS
@estimate
parameter initialization.B0 # Estimate B0 when fitting the model
phase 2                    # but use two-phase estimation, and
                          # only try to "fit" this parameter after
                          # fitting all other parameters first
lower_bound 2500          # Define the lower bound
upper_bound 150000        # Define the upper bound
prior uniform-log        # and use a uniform-log prior

@estimate
parameter recruitment.YCS # Estimate YCS when fitting the model

```

```

#the YCSyears are 1973 1974 1975 1976 1977 1978 1979 1980 1981 1982 1983
    1984 1985 1986 1987 1988 1989 1990 1991 1992 1993 1994 1995 1996 1997
    1998 1999 2000
lower_bound 1 1 1 1 1 1 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01
    0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 0.01 1 1 # lower bounds
upper_bound 1 1 1 1 1 1 100 100 100 100 100 100 100 100 100 100 100 100 100
    100 100 100 100 100 100 1 1 # upper bounds. Note that some YCS are
    # constrained to be equal to one

prior lognormal # Use a lognormal prior
mu 1 # with  $\mu=1$ , and c.v.=1.1
cv 1.1

@estimate
parameter selectivity[chatTANsel].male # Estimate the "chatTANsel.male"
    # ogive.
lower_bound 0 0 # the two logistic parameters have lower and upper
upper_bound 30 30 # bounds
prior uniform # and they have uniform priors

@estimate
parameter selectivity[chatTANsel].female # Estimate the "chatTANsel.female"
    # ogive.
lower_bound 0 0 0.2 # The three logistic parameters have lower and upper
upper_bound 30 30 5 # bounds
prior uniform # And they have uniform priors

@estimate
parameter selectivity[chatFsel].male # Estimate the "chatFsel.male"
    # ogive.
lower_bound 0 0 # The two logistic parameters have lower and upper
upper_bound 30 30 # bounds
prior uniform # And they have uniform priors

@estimate
parameter selectivity[chatFsel].female # Estimate the "chatFsel.female"
    # ogive.
lower_bound 0 0 0.2 # The three logistic parameters have lower and upper
upper_bound 30 30 5 # bounds
prior uniform # And they have uniform priors

{
# This is a comment block, commenting out the request to estimate the two
# parameters natural_mortality.avg and natural_mortality.diff
@estimate
parameter natural_mortality.avg
phase 1
prior lognormal
mu 0.20
cv 0.20
lower_bound 0.10
upper_bound 0.30

@estimate
parameter natural_mortality.diff
phase 1
prior normal-by-stdev
mu 0
stdev 0.05
lower_bound -0.20
upper_bound 0.20
}

# PENALTIES
@catch_limit_penalty # This specifies that the model must attempt to have a
    # biomass large enough so that the catch is takable
    # from the population
label chatCatchMustBeTaken
fishery chatFishery

```

```
log_scale true
multiplier 1000      # The penalty has a high "multiplier"

@ageing_error        # Specify the ageing error model used
type normal          # Ageing error is of type "normal"
c 0.08               # with a c.v.=0.08

@vector_average_penalty # Defines that there is a penalty applied
label YCS_average_1_CR # to encourage the YCS to have mean=1
vector recruitment.YCS
k 1
multiplier 50        # but this has a moderate "multiplier"
```

R1-output.csl

```
@print # Specifies the outputs that CASAL should generate
# estimation section
parameters false
fits_every_eval false
objective_every_eval false
parameters_every_eval false
parameter_vector_every_eval false
fits true
resids false
pearson_resids false
normalised_resids true
estimation_section false
# population section
requests true
initial_state false
state_annually false
state_every_step false
final_state true
results false
#output section
yields true
unused_parameters true

@quantities
all_free_parameters true
fishing_pressures true
nuisance_qs true
true_YCS true
B0 true
R0 true
SSBs true
YCS true
actual_catches false
ogive_parameters selectivity[chatTANsel].male selectivity[chatTANsel].female
                selectivity[chatFsel].male selectivity[chatFsel].female

@MCY_CAY
do_MCY true
MCY_guess 10000
n_discard 100
n_keep 100
n_simulations 100
do_CAY true
F_CAY_guess 0.2
interactive false
```


13.3 CASAL output

The call `casal -e -q -g 0 -f R1- > R1-estimate.log` with the above parameter files generates an output file. The first few lines of the output file, `R1-estimate.log`, are;

```
CASAL (C++ algorithmic stock assessment laboratory)
Call: c:\WINNT\casal.exe -e -q -g 0 -O MPD.dat -f R1-
Date: Fri Aug 01 12:56:37 2003
Key C++ source files used (all times are UTC):
$Id: model.C,v 1.109 2003/08/01 00:47:00 adunn Exp $
$Id: estimation.C,v 1.149 2003/07/28 08:32:35 bull Exp $
$Id: population.C,v 1.152 2003/07/28 08:32:36 bull Exp $
$Id: output.C,v 1.54 2003/07/28 08:32:36 bull Exp $
$Id: input.C,v 1.65 2003/07/28 08:32:36 bull Exp $
$Id: dictionary.C,v 1.30 2003/07/10 02:00:50 bull Exp $
$Id: development.h,v 1.60 2003/07/17 07:54:34 bull Exp $
$Id: betadiff.C,v 1.59 2003/07/17 06:15:56 bull Exp $
$Id: betadiff.h,v 1.60 2003/06/26 06:35:35 bull Exp $
v2.01-2003/08/01 (c) Copyright NIWA 2002, 2003
User name: dunn
Machine name: Alistair

Prefix for the input parameter filenames : R1-

Random number seed : 0

A male fish of size 50 cm has a weight of 0.777443 kg in your model. This
size-weight scale check assumes that the growth curve is in
centimetres, and that the catch is in tonnes.

A female fish of size 50 cm has a weight of 0.766692 kg in your model. This
size-weight scale check assumes that the growth curve is in
centimetres, and that the catch is in tonnes.

A male fish of size 50 cm has a weight of 0.777443 kg in your model. This
size-weight scale check assumes that the growth curve is in
centimetres, and that the catch is in tonnes.

A female fish of size 50 cm has a weight of 0.766692 kg in your model. This
size-weight scale check assumes that the growth curve is in
centimetres, and that the catch is in tonnes.

In phase 1 Minimiser achieved convergence after 107 quasi-Newton iterations
using 158 objective function evaluations

In phase 2 Minimiser achieved convergence after 40 quasi-Newton iterations
using 95 objective function evaluations

An entry of recruitment.YCS is at or near the lower bound

These lines give information on the CASAL version that was run, the date, and machine, file
prefix, random number generator seed, length weight validation output, and the convergence
summary from the minimiser. Note that the estimation carried out here is two phase
estimation, with the first phase holding  $B_0$  constant.

The output also warns that an estimated YCS is close to a bound.

The remainder of the output file gives fits, residuals, and other requested output information,
as defined in the output.csl file. Some of this output is reproduced here.
```

Start extracting output from here

```

Point estimate:
initialization.B0
current value: 26563.8

recruitment.YCS
current values:
1 1 1 1 1 1 0.0100004 3.47924 0.01 1.87201 0.278127 1.15573 0.607087
    0.915851 0.527776 1.15652 0.908426 1.33845 1.94922 1.53675 1.22428
    1.1685 0.882493 0.50438 0.37749 0.102835 1 1

selectivity[chatTANsel].male (parameters)
current values:
11.976 11.4088

selectivity[chatTANsel].female (parameters)
current values:
10.0316 11.0072 0.685787

selectivity[chatFsel].male (parameters)
current values:
8.52579 12.4255

selectivity[chatFsel].female (parameters)
current values:
6.32959 11.1406 0.618114
    
```

In a format suitable for -i:

```

initialization.B0 recruitment.YCS 28 selectivity[chatTANsel].male 2
    selectivity[chatTANsel].female 3 selectivity[chatFsel].male 2
    selectivity[chatFsel].female 3
26563.8 1 1 1 1 1 1 0.0100004 3.47924 0.01 1.87201 0.278127 1.15573 0.607087
    0.915851 0.527776 1.15652 0.908426 1.33845 1.94922 1.53675 1.22428
    1.1685 0.882493 0.50438 0.37749 0.102835 1 1 11.976 11.4088 10.0316
    11.0072 0.685787 8.52579 12.4255 6.32959 11.1406 0.618114
    
```

Objective function : -141.843

```

Components :
                                chatCPUE -6.99776
                                chatTANbiomass -18.72
                                chatTANage -75.3991
                                chatOBS -40.4331
    prior_on_initialization.B0 10.1873
    prior_on_recruitment.YCS 0.15125
    prior_on_selectivity[chatTANsel].male 0
    prior_on_selectivity[chatTANsel].female 0
    prior_on_selectivity[chatFsel].male 0
    prior_on_selectivity[chatFsel].female 0
    prior_on_q_chatCPUEq -9.50212
    prior_on_q_chatTANq -1.12913
    chatCatchMustBeTaken 0
    YCS_average_1_CR 1.69823e-006
    ...
    
```

The output then continues with fits (if requested), other output quantities, unused parameters, and other requested quantities.

These include the following lines ...

```

* Nuisance q's
chatCPUEq 7.4693e-005
chatTANq 0.323314
    
```

```

* B0
26563.8

* R0
2.04757e+006

* SSBs
SSB 26429.7 26095.8 25224.6 25290 24962.6 24571.7 24055.6 23863.2 23858.8
      23953.7 23729 23338.2 25623.2 26843.4 26862.5 25793.4 24546.7 22226.9
      19935.2 17840.5 15907.7 14876.5 14577.2 15129.6 15512.5 14873.1 13485.8
      11545.1
year 1975 1976 1977 1978 1979 1980 1981 1982 1983 1984 1985 1986 1987 1988
      1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002

* YCS
YCS 1 1 1 1 1 1 0.00999783 3.47835 0.00999742 1.87153 0.278056 1.15543
      0.60693 0.915615 0.52764 1.15622 0.908191 1.3381 1.94872 1.53635
      1.22397 1.1682 0.882265 0.50425 0.377392 0.102808 1 1
year 1973 1974 1975 1976 1977 1978 1979 1980 1981 1982 1983 1984 1985 1986
      1987 1988 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 2000
...

```

The file `R1-estimate.log` gives the estimated point estimates, in this case, for the MPD fit using Bayesian estimation. MCMC output can be generated by the call (directing the output to `R1-mcmc.log`). In this example, the random number seed has been specified to be zero (using the switch `-g 0`), rather than allowing CASAL to generate it from the local computer time.

```
>casal -m -q -g 0 -f R1- > R1-mcmc.log.
```

This generates two additional output files `objectives.1` and `samples.1` (assuming that the directory contains no other objectives or samples output). The `objectives.1` file lists the model call and header, the estimated covariance matrix from the MPD fitting, and the MCMC sampling diagnostics for each requested step (in this case, every 1000th step), i.e.,

```

...
Main table:
sample posterior prior likelihood penalties stepsize acceptance_rate
      stepsize_changes
1000 -121.808 10.5515 -136.774 4.41441 0.02 0.637 0
2000 -126.865 10.4032 -137.394 0.125419 0.02 0.6455 0
3000 -121.505 10.4738 -136.299 4.31987 0.02 0.658667 0
...

```

`samples.1` contains the MCMC output (in a format compatible with `casal -i [filename]`) i.e.,

```

initialization.B0 recruitment.YCS 28 selectivity[chatTANsel].male 2
      selectivity[chatTANsel].female 3 selectivity[chatFsel].male 2
      selectivity[chatFsel].female 3
32867.4 1 1 1 1 1 1 0.161592 5.17398 0.310461 2.76911 0.79027 1.16298
      0.834752 1.62442 0.34993 1.63625 1.3293 1.69915 2.69013 2.13613 1.56574
      1.64539 1.10604 0.737461 0.514854 0.0817992 1 1 8.04818 10.2657 8.09448
      12.0459 0.954801 5.53225 9.50579 3.71554 13.5413 0.790136
28337.7 1 1 1 1 1 1 0.0875087 3.80083 0.746385 0.423447 1.50323 0.741527
      0.569092 1.20748 0.420993 1.04968 1.15698 1.52884 1.72183 1.74998
      1.24671 1.39139 0.926057 0.469134 0.544376 0.116864 1 1 11.5451 11.7334
      8.69597 9.49665 0.613545 9.87319 17.0835 6.01954 8.93415 0.478103
30409.7 1 1 1 1 1 1 0.154603 3.03452 1.79086 1.68867 1.12878 0.827664
      1.16567 1.27732 0.460741 1.92121 1.28814 1.69889 3.07818 2.19876 1.6665

```

```
1.94298 1.39616 0.793912 0.577871 0.138751 1 1 10.2948 9.60845 14.3265
15.1816 1.30681 8.76698 11.5262 8.98394 15.1572 0.739428
```

...

The MCMC output contains estimates of the parameters defined in the `estimation.csl` file. To generate a file of quantities, CASAL must be run with the command

```
>casal -i samples.1 -v quantities.1 -f R1- > R1-quantities.log
```

`quantities.1` contains the MCMC quantities output. The file contains all of the estimated quantities at each point in the `samples.1` output file. For example,

```
Quantity values :
initialization.B0 recruitment.YCS[1] recruitment.YCS[2] recruitment.YCS[3]
  recruitment.YCS[4] recruitment.YCS[5] recruitment.YCS[6]
  recruitment.YCS[7] recruitment.YCS[8] recruitment.YCS[9] ...
31744.1 1 1 1 1 1 1 0.245979 12.735 0.547051 5.72196 2.55924 3.97424
```

...

etc.,

The output, `quantities.1`, is then available, in tabular format, for importing into another package (e.g., Excel, S-Plus, or R) for plotting and summarising.

13.4 Generating simulated output

The call `casal -e -O MPD.dat -q -g 0 -f R1-` with the above parameter files generates an additional output file, `MPD.dat`. This file contains only two lines, a header, and the free parameter values at the fitted MPD, e.g.,

```
initialization.B0 recruitment.YCS 28 selectivity[chatTANsel].male 2
  selectivity[chatTANsel].female 3 selectivity[chatFsel].male 2
  selectivity[chatFsel].female 3
26569.3 1 1 1 1 1 1 0.0102062 3.53837 0.0100016 1.90027 0.285513 1.17286
  0.617232 0.930923 0.536297 1.17543 0.923212 1.36033 1.98089 1.5619
  1.24429 1.18763 0.897004 0.512692 0.383785 0.10455 1 1 11.9662 11.4039
  10.0264 11.0038 0.686283 8.52119 12.4223 6.32529 11.1334 0.618281
```

The call `casal -s 10 simulated -i MPD.dat -g 0 -q -f R1-` generates 10 sets of files with simulated observations, derived from the fit specified by `MPD.dat`. These files are named `simulated.par1.sim[n]`, where n is a number from 1 to 10.

These can be appended (one at a time) to another `estimation.csl` file (that has no observations) called, say, `R1-estimation.stub`, the result renamed to `Sim-estimation.csl`, and CASAL called sequentially on each of these to parametric bootstrap around the MPD, i.e., assuming that the files `Sim-population.csl` and `Sim-ouput.csl` are copies of the corresponding `R1-` files, then repeating the call

```
casal -e -i MPD.dat -o simulations.dat -f Sim-
```

for each set of simulated observations will generate the file `simulations.dat` that contains 10 bootstrap estimates around the MPD, i.e.,

```
initialization.B0 recruitment.YCS 28 selectivity[chatTANsel].male 2
  selectivity[chatTANsel].female 3 selectivity[chatFsel].male 2
  selectivity[chatFsel].female 3
```

```
27665.4 1 1 1 1 1 1 1.36606 2.91338 0.0100011 2.13207 0.457089 1.20367
0.37985 1.06572 0.310713 1.01199 0.811997 1.51831 1.79606 1.58328
0.912118 1.18974 0.663791 0.465778 0.277822 0.0725319 1 1 14.5971
16.6122 7.44284 8.55728 0.48304 7.0429 7.5742 4.4964 8.36031 0.643517
27127.9 1 1 1 1 1 1 0.0100007 4.44302 0.0100011 0.415619 2.24848 0.0627381
0.878195 0.65192 0.408248 1.22999 0.876439 1.17551 2.07134 1.46113
1.20296 1.05986 0.861373 0.457994 0.443146 0.133978 1 1 7.92103 7.5092
10.2896 11.9435 1.10429 7.57573 11.6446 6.86198 9.46643 0.6284
28644.4 1 1 1 1 1 1 0.346082 2.42542 2.1835 0.0395063 1.41114 0.614499
0.214362 0.989823 0.359256 0.953136 1.06654 1.31012 1.85754 1.61956
1.12444 1.15749 0.988708 0.629174 0.511974 0.164525 1 1 10.1979 10.3475
10.3097 10.3891 1.00083 10.5809 13.7244 6.47676 10.1475 0.481152
...
```

This file can then be used to generate bootstrap quantities (using `casal -i simulations.dat`, as earlier with the MCMC output)

14. ENHANCEMENTS AND OTHER CHANGES FROM CASAL v1.02-2002/10/21

14.1 Issues fixed

1. A wide range of additional error checking has been implemented. CASAL now does much more extensive checks to ensure that the commands, sub-commands, and parameters for the input parameter files are valid.
2. In certain circumstances, simulations (i.e., MCY/CAY calculations) produced erroneous results. These calculations have been repaired.
3. The burn-in period specified for Bayesian models in the `estimation.csl` file has been repaired so that it now acts as specified in the manual.
4. A bug relating to the application of size-based fishery selectivity ogives in age-based model with no ogive shift, has been repaired.
5. The application of mature and immature selectivities has been swapped, so it is now acting as it should.
6. Some minor bugs relating to the implementation of length based models have been repaired.
7. Natural mortality can now be estimated for independently for different stocks in a multi-stock model.
8. A bug in `@profile` has been repaired so the output when profiling more than one parameter is as expected.
9. If a logical (switch) argument is supplied that is not of the form true/false, CASAL will now report an error.
10. A bug in the output of quantities when `Cinitial` was specified has been corrected.
11. A problem with determining the initial state (`Cinitial`) that occurred if ageing happened before recruitment in the annual cycle has been identified. CASAL now reports an error message if it cannot determine the initial state under these circumstances.
12. Overflow errors that resulted from some calculations are identified before being attempted.. CASAL now handles most situations where parameter values would lead to an overflow using improved algorithms.
13. An error in the `logistic_producing` ogive has been repaired, so it is now acting as it should.
14. The `@estimate.same` subcommand has been repaired so that it now works as it should with estimable estimation parameters (i.e., catchability constants, q).
15. A small number of other problems in lesser used functions and print functions have also been corrected.

14.2 Changes and enhancements

1. A new command line switch `casal -l` has been added to display the end-user licence.
2. New command line switches (`-o` and `-O`) have been added that dump the free parameters (from `casal -r`, `-e`, or `-E`) to a text file, in a form suitable for use with `casal -i`.
3. An additional transition process, disease mortality, has been added to the annual cycle. See Section 4 for details.
4. Additional output for length based selectivities in and age-based model have been added. These allow the user to request that CASAL output length based selectivities as age based selectivities, given the partition, year, and time step.
5. An additional penalty, `vector_smoothing_penalty` has been introduced. This behaves in a manner similar to the `ogive_smoothing_penalty`.
6. Additional ogives have been added; `double_normal_plateau`, `double_logistic`, and `logistic_product`.
7. Fits, residuals, Pearson residuals, and normalised residuals can now be output as quantities.
8. The `pseudo`, `vector`, and `ogive` sub-commands have been classified as obsolete. These are no longer necessary.
9. The constraint that did not allow migrations to be both a source and a destination in the same time step has been removed. Migrations are now applied in the order specified in the parameter files.
10. The beta prior has been enhanced to include scale and shift parameters. It is now specified in terms of a standard deviation rather than a c.v.
11. Multiple copies of single warning messages are now (mostly) suppressed, so that a warning message is only reported once.
12. The behaviour of the random number seed has been changed so that the default behaviour (when no seed is supplied) is to create a seed from the computer clock. Use the command `casal -g 0` option to replicate previous behaviour.
13. Two new command line switches allow the user to modify the names of the parameter input files (`casal -f` to add a prefix and `casal -F` to modify the `cs1` suffix)
14. A new forms of parameterising YCS has been added (the Francis parameterisation). See Section 4.4.2.
15. The `@recruitment` subcommand, `standardise_YCS`, has been made a command. Now the form of YCS parameterisation must be specified for the entire model with a single command, rather than separately by stock (i.e., by using the command `@standardise_YCS true`).
16. The `proportions_mature` observation type has been enhanced to allow for (a) size-based models and (b) unsexed models.

17. The `age_size` observation type has been rewritten and considerably enhanced. See Section 5.6 for details.
18. The default method for generating recruitments in MCY/CAY simulations has been changed so that there is no default. This must now be specified by the user.
19. Alternative reference years for defining risk in the MCY/CAY simulations can now be specified.
20. CASAL has been ported from `gcc` version 2.95/2.96 to `gcc` version 3.2.3. CASAL has also been implemented as a native Microsoft Windows application.
21. Simulations with a growth curve of type `size_at_age = data` are now implemented for multi-stock models. The `@size_at_age.simulation_male`, `simulation_female`, and `simulation_all` subcommands are now obsolete. CASAL uses the mean size of all supplied observations in simulations instead.
22. An additional error check on values of user supplied units of the size-weight parameters has been implemented. CASAL will now either warn if the values of the size-weight parameters supplied are not within the default range, or will error out if the size-weight parameters are outside a user supplied range.
23. A new observation type, age-at-maturation has been added. See Section 5.6.3 for details.
24. Size based observations in a size-based model can now be integer combinations (i.e., pooled combinations) of the size-classes defined by the partition. This allows size observations with measurements made at a coarser scale than that defined by the partition to be used in the model.
25. CASAL can now generate simulated observations, i.e., generate observed values with random error that are based on a supplied “fit”. See Section 5.9.
26. If the `@print.covariance` is requested for a `-e` or `-E` run, CASAL will also print the eigenvalues of the Hessian.
27. CASAL can now, optionally, update the Hessian matrix during the burn-in phase of an MCMC.
28. The binomial likelihood has been added for use with proportions mature and proportions migrating observations. See Section 5.7.2.

ACKNOWLEDGMENTS

We thank Andreas Griewank and the other creators of the ADOL-C package — the basis of the `Betadiff` software underlying CASAL — for freely supplying this product. Thanks also to Robert Davies for providing the `newran` random number generation library.

We thank Ian Doonan and Allan Hicks for help in the planning stages of the development of CASAL. Neville Phillips, Peter Horn, and Stuart Hanchet provided helpful comments on early drafts of this manual. We also thank Susan Kim and Paul Breen for helpful discussions and comments. Gavin Macaulay provided invaluable advice on the technical aspects of the C++ implementation and in assisting in porting CASAL to MinGW. The development of CASAL was funded by NIWA.

REFERENCES

- Bull, B.; Livingston, M.E. (2001). Links between climate variation and the year class strength of New Zealand hoki (*Macruronus novaezelandiae*): an update. *New Zealand Journal of Marine and Freshwater Research* 35(5): 871–880.
- Clark, W.G. (1991). Groundfish exploitation rates based on life history parameters. *Canadian Journal of Fisheries and Aquatic Sciences* 48: 734–750.
- Cordue, P.L. (2000). MIAEL estimation of biomass and fishery indicators for the 1999 assessment of hoki stocks. *New Zealand Fisheries Assessment Report 2000/10*. 69 p.
- Davies, R.B. (1998). Newran02A. A random number generator library <http://www.robertnz.net/index.html>.
- Dennis Jr, J.E.; Schnabel, R.B. (1996). Numerical methods for unconstrained optimisation and nonlinear equations. Prentice Hall. 378 p.
- Fournier, D.A. (1994). AUTODIF. A C++ array language extension with automatic differentiation for use in nonlinear modeling and statistics. Otter Research Ltd., Sidney, Canada. 123 p.
- Fournier, D.A.; Sibert, J.R.; Majkowski, J.; Hampton, J. (1990). MULTIFAN: a likelihood-based method for estimating growth parameters and age composition from multiple length frequency data sets illustrated using data for southern bluefin tuna. *Canadian Journal of Fisheries and Aquatic Sciences* 47: 301–317.
- Francis, R.I.C.C. (1988). Maximum likelihood estimation of growth and growth variability from tagging data. *New Zealand Journal of Marine and Freshwater Research* 22(1): 43–51.
- Francis, R.I.C.C. (1992). Recommendations concerning the calculation of Maximum Constant Yield (MCY) and Current Annual Yield (CAY). New Zealand Fisheries Assessment Research Document 92/8. 23 p. MAF (Fisheries). (Unpublished report held in NIWA library, Wellington.)
- Francis, R.I.C.C.; Haist, V.; Bull, B. (2003). Assessment of hoki (*Macruronus novaezelandiae*) in 2002 using a new model. *New Zealand Fisheries Assessment Report 2003/6*. 69 p.
- Francis, R.I.C.C.; Horn, P.L. (1997). Transition zone in otoliths of orange roughy (*Hoplostethus atlanticus*) and its relationship to the onset of maturity. *Marine Biology* 129: 681–687.
- Gelman, A.B.; Carlin, J.S.; Stern, H.S.; Rubin, D.B. (1995). Bayesian data analysis. Chapman and Hall, London. 526 p.
- Gilks, W.R.; Richardson, S.; Spiegelhalter, D.J. (eds.) (1998). Markov chain Monte Carlo in practice. *Interdisciplinary statistics*. 399 p. Chapman and Hall/CRC Press, Boca Raton, Florida.
- Griewank, A.; Juedes, D.; Mitev, H.; Utke, J.; Vogel, O.; Walther, A. (1996). ADOL-C: A package for the automatic differentiation of algorithms written in C/C++. *ACM Transactions on Mathematical Software* 22(2): 131–167. Algorithm 755.

Gulland, J.A.; Boerema, L.K. (1973). Scientific advice on catch levels. *Fishery Bulletin* 71: 325–335.

Harley, S.J.; Myers, R.A.; Dunn, A. (2001). Is catch-per-unit-effort proportional to abundance? *Canadian Journal of Fisheries and Aquatic Sciences* 58(9): 1760–1772.

Hilborn, R.; Maunder, M.; Parma, A.; Ernst, B.; Payne, J.; Starr, P.J. (2001). Coleraine: A generalised age-structured stock assessment model users manual. Version 2.0. School of Aquatic & Fishery Sciences, University of Washington. FRI-UW Report Series 0116. 58 p. University of Washington

Ihaka, R.; Gentleman, R. (1996). R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics* 5(3): 299–314.

Otter Research Limited (2000). An introduction to AD model builder, version 4: For use in nonlinear modelling and statistics. 127 p. Otter Research Limited. Sydney, B.C., Canada.

Punt, A.E.; Hilborn, R. (2001). BAYES-SA. Bayesian stock assessment methods in fisheries. User's manual. *FAO Computerized information series (fisheries) 12*. Food and Agriculture Organisation of the United Nations, Rome (Italy). 56 p.

Schnute, J. (1981). A versatile growth model with statistically stable parameters. *Canadian Journal of Fisheries and Aquatic Sciences* 38(9): 1128–1140.

Smith, A.D.M.; Punt, A.E.; Wayte, S.E.; Starr, P.J.; Francis, R.I.C.C.; Stokes, T.K.; Hilborn, R.; Langley, A. (2002). Stock assessment of the northeast Chatham Rise orange roughy for 2001. *New Zealand Fisheries Assessment Report 2002/25*. 30 p.

Smith, B.J. (2001). Bayesian output analysis program. Version 1.00 user's manual. Unpublished manuscript. 45 p. University of Iowa College of Public Health. (see <http://www.public-health.uiowa.edu/boa>)

Starr, P.J.; Bentley, N.; Maunder, M. (1999). Assessment of the NSN and NSS stocks of red rock lobster (*Jasus edwardsii*) for 1998. New Zealand Fisheries Assessment Research Document 99/34. 45 p. Ministry of Fisheries. (Unpublished report held by the Ministry of Fisheries, Wellington.)

Walters, C.J.; Ludwig, D. (1994). Calculation of Bayes posterior probability distributions for key population parameters. *Canadian Journal of Fisheries and Aquatic Sciences* 51: 713–722.

INDEX

- ADOL-C 11
- Age-based model 25
- Ageing error 75
- Misclassification matrix 75
- Misclassification rates 75
- Normal 75
- Off-by-one 75
- Annual cycle 23
- Errors in 27
- AUTODIF 11
- Automatic differentiation 51
- Speed 51
- B_0 28, 38
- Baranov catch equation 34, 35
- Bash shell 13
- Bayesian analysis in CASAL 15
- Bayesian estimation
- Acceptance probability 55
- adaptive proposal distribution 54
- Default proposal distribution 53
- Fixed free parameters 55
- Interruptions in the MCMC chain 52
- Investigating alternative priors from an
 MCMC chain 55
- max_cor 54
- MCMC chain 52
- MCMC output 52
- MCMC starting point 53
- Metropolis algorithm 53
- Multiple MCMC chains 52
- Multivariate t degrees of freedom 55
- Multivariate t proposal distribution 55
- Post-processing of output 53
- Projections 55
- Proposal distribution 53
- Random sub-sampling 55
- stepsize 53, 54
- Stochastic yields 55
- Sub-sampling MCMC chains 52, 55
- Systematic sub-sampling 55
- Bayesian output analysis package
- See BOA 10
- Betadiff 11
- Beverton Holt
- See stock recruitment relationship 28
- $B_{initial}$ 38
- B_{mean} 38
- BOA 10, 52
- C++ code extensions 167
- Likelihoods 172
- new parameterisations 167
- Penalties 170
- Priors 170
- casal.exe 10
- Catch limit penalties
- See Penalties 36
- Catchability constants – q 's 66, 68
- $C_{initial}$ 38
- Citing CASAL 10
- Citing this document 10
- Climate-recruitment relationship 29
- Arctan 29
- Exponential 29
- Identity 29
- Linear combinations 29
- Logistic 29
- Predictions 29
- Command line arguments 13
- Append MCMC results 13
- Calculate projections 14
- Calculate yields 14
- Concatenate MCMC results 14
- Filename prefix 14
- Filename suffix 14
- Finite difference estimation 13
- Generate simulate observations 14
- Input additional data files 14, 15, 210
- Machine name 15
- MCMC 13
- Output values 14
- Point estimates 13
- Posterior profiles 13
- Quiet mode 14
- Random number seed 15
- Run 13
- Save posterior sub-sample 15
- Command-block format 18
- Convergence
- Convergence evaluations 50
- Convergence failure 50
- Convergence iterations 50
- Convergence unclear 50
- Failure to converge 50
- Successful convergence 50
- Correlation matrix 49
- Covariance matrix 49, 51
- CSP
- See Current surplus production 91
- Current surplus production 91
- B_{post} 91
- Multi-stock models 91
- Current year 24
- Density dependent migration 32
- Deterministic MSY 88
- F_{MSYdet} 88
- Simulations 88
- SR_simulation 88
- steepness_simulation 88
- Yield versus SSB curve 88
- Deterministic yields 87
- Constant mortality rate F 87
- Deterministic MSY 87

Multi-stock models.....	87	Initial number of fish.....	38
Per-recruit analysis.....	87	Initial recruitment.....	29
Disease mortality.....	37	Initial state.....	37, 38
Equilibrium abundance.....	37	Initial year.....	24
Equilibrium distribution.....	38	Input data files.....	13
Errors.....	163	estimation.csl.....	13, 17, 18, 82, 83, 119
Betadiff error.....	164	output.csl.....	13, 18, 81, 82, 83, 86, 151
Catch limit penalty.....	163	population.csl.....	13, 18, 36, 37, 47, 53, 82, 93
Comments.....	163	Instantaneous fishing rate.....	36
EOF carriage return.....	163	Instantaneous mortality.....	35
Exceeded maximum exploitation penalty.....	163	Interrupting the MCMC chain.....	15
Maturation rates.....	164	Known parameters.....	21
Misspelt arguments.....	163	Least-squares.....	
Misspelt commands.....	163	Sources.....	63
Penalties.....	163	Weights.....	63
print_unused_parameters.....	163	Likelihood profiles.....	51
YCS average one.....	164	Multi-phase estimation.....	52
Estimable parameters.....	20, 21, 49	Scalar parameters.....	52
Bounds.....	49	Likelihoods.....	
Command block.....	49	Binomial.....	65
Ogives.....	49	Coleraine.....	64
Prior.....	49	Curvature parameters.....	68
same subcommand.....	50	Fournier.....	64
Estimation parameters.....	20, 21	Lognormal.....	66, 67
Estimation section.....	21	Multinomial.....	64
Exit codes.....		Normal.....	65, 67
Convergence unclear.....	50	Normal-log.....	66, 67
Failure to converge.....	50	Nuisance q 's.....	66
Successful convergence.....	50	Process error.....	68
Exploitation rate.....	35	Profiles.....	51
extract_CASAL.v2.01.s.....	10	q 's as free parameters.....	66
False minima.....	51	Robustified lognormal.....	64, 67
Final year.....	24	Robustified multivariate normal.....	64
Finite difference estimation.....	51	Linux.....	10
Speed.....	51	Local optimisation.....	50
Fishery.....	26	Maturation.....	31
Definition.....	34	Maturity.....	47
Fishing mortality.....	34	Maturity in the partition.....	47
Free parameters.....	20, 21	Migrations of mature or immature fish.....	47
gcc.....	11	Proportions mature.....	47
Generate simulated observations.....	76	Maximum.....	
Global optimisation.....	50	Fishing pressure.....	35
Growth.....		Instantaneous fishing mortality rate.....	36
Size based model.....	33	MCMC posterior samples.....	49
Growth curves.....	45	Microsoft Windows.....	10
Annual growth variation.....	45	Migration.....	32
Distribution of size-at-age.....	45	MinGW.....	11
Empirical.....	45	Minimiser starting value.....	50
growth_props parameter.....	45	Modelling numbers of individuals.....	47
Mean size-at-age.....	45	Mortality.....	34
Schnute.....	45	Multi-phase estimation.....	51
von Bertalanffy.....	45	Natural mortality.....	34
Growth-paths.....	24	Necessary files.....	10
Haist parameterisation of YCS.....	29	newran.....	11
Hessian.....	51	$n_{initial}$	29
Approximation to.....	51	Non-estimable parameters.....	20
Eigenvalues.....	51	Nuisance q 's.....	49, 66, 68
Inverse.....	51, 53	Bayesian analysis.....	69
Matrix.....	49	Least-squares.....	69

Maximum likelihood	69	increasing	44
Priors	70	increasing_capped.....	44
Numeric overflow error	40, 164	knife_edge	41
Objective function	21, 49	Length-based model.....	39
Bayesian estimation.....	62	logistic	41
least-squares	49	logistic_bounded.....	42
Least-squares	62	logistic_capped	42
Likelihoods.....	49, 63	logistic_producing	44
Likelihoods for absolute indices.....	65	logistic_product	42
Likelihoods for proportions.....	64, 65	Overview	39
Likelihoods q 's.....	66	Size-based in an age-based model	40
Log-posterior.....	49	Usage	39
Maximum likelihood	62	Optimiser.....	11
Penalties	63	Output	81
Observations	21, 56	Estimation section.....	82
Abundance indices	56	File names	81
Age/size data	58	Header.....	81
Age-at-maturation	56	Output quantities.....	82
Age-at-maturation likelihood	61	Population section.....	81
Age-at-maturation observations	60	Projections	81
Ageing error	58	Results	81
Age-size.....	56	User requests.....	81
Age-size data likelihoods	59	Output parameters	21
Area	57	Output quantities	82
Catch-at-age	56	B_0	83
Catch-at-size.....	56	$B_{initial}$	83
Curvature parameter on q 's	57	B_{mean}	83
Error distributions in a Bayesian analysis	58	Catches.....	83
Fishery.....	57	Climate relationship.....	83
Label.....	57	Fishing pressure	83
Likelihoods		Free parameters.....	82
Age at maturation.....	61	Ogive arguments.....	82
Age-size observations	59	Ogives	82
Numbers-at-age	56	Parameters	82
Numbers-at-size	56	Pseudo fits.....	83
Observation years	57	q 's	83
Plus group.....	56, 57	R_0	83
Proportions mature	57	Recruitment	83
Proportions migrating.....	56	$R_{initial}$	83
Proportions-at-age	56	R_{mean}	83
Proportions-at-size	56	SSBs	82
q 's.....	57	Stock risk	83
Selectivities	57	True YCS	83
Simulating	76	YCS deviates	83
Size classes in an age-based model	57	Output section	21
Specifying observations in		Parameter argument types	19
estimation.csl.....	58	Parameter file format.....	17
Sum to 1	57	Comments	18
Weights in least-squares estimation	58	Constructing a CASAL input parameter file	
Ogives		18
Age-based models	39	Illegal characters.....	18
allvalues	41	Parameter names	20
allvalues_bounded.....	41	Parameter types.....	19, 21
constant.....	41	Ogives	19
Discrete approximation	40	Switches.....	19
double_logistic.....	42	Parametric bootstrap.....	76
double_normal.....	43	Partition.....	23, 24
double_normal_capped.....	43	Area	24
double_normal_coleraine.....	43	Growth-paths	24
double_normal_plateau.....	43	Maturity	24

Sex.....	24	No randomisation of YCS.....	85
Stock.....	24	$P(B_{\text{current}+k} > B_{\text{current}})$	86
Partition matrix	23	Point estimates	84, 86
Penalising estimation parameters.....	20	Posterior distribution	84
Penalties		Post-processing of output.....	86
Catch limit.....	36, 74	Projected expectation.....	85
Element difference	74	Randomised recruitment.....	84
Multiplier.....	73	Randomised recruitment for cohorts.....	84
Ogive comparison	75	Simulations	84
Ogive difference.....	75	Stock risk	86
Ogive smoothing	74	Pseudo fits.....	83
Similar q 's	74	Pseudo observations	83
Vector average.....	74	q 's.....	66, 68
Vector smoothing	74	q 's as free parameters.....	66
YCS difference.....	74	Quasi-Newton minimiser	50
Penalties.....	73	R_0	28, 38
Point estimates.....	21, 49, 50	Random number generator.....	11
Least-squares	49	Recruitment.....	28
MLE	49	Climate recruit relationship	28
MPD.....	49	Stock recruit relationship	28
Population parameters	21	Recruitment for individual stocks	26
Population section.....	21	Residuals	76
Post-processing of CASAL output	10	Normalised.....	76
Priors.....	71	Pearson.....	76
Beta	72	Undefined normalised residuals.....	76
Lognormal	72	Ricker	
Multivariate normal from a stationary		<i>See</i> stock recruitment relationship	28
AR(1) process	73	R_{initial}	29, 38
Normal.....	72	R_{mean}	38
Normal-by-stdev.....	72	S/S-Plus/R	
Normal-log	72	extract.fits.....	175, 177
Uniform.....	72	extract.free.parameters	175, 177
Uniform-log.....	72	extract.free.parameters.from.	
Process error	68	table.....	179
c.v.....	68	extract.header.....	175
Effective sample size.....	68	extract.objective.function	175,
Estimating	68	176	
Standard deviation.....	68	extract.quantities	175, 178
Profiles of parameters	49	extract.quantities.from.table	
Projections	24, 84	180
Adaptive harvest strategies.....	85	read.files.for.BOA	175, 181
Autocorrelation in randomisation of YCS	85	Selectivity ogive	
CAY	84	Shift parameter.....	37
Climate recruit relationship	84	Sex partitions.....	23
Climate recruit relationship and		Simulations	
randomisation of YCS.....	85	Observations	76
CSP.....	84	Simultaneous MCMC chains	16
Current Surplus Production	84	Size data in an age based model.....	25
$E(B_{\text{current}+k}/B_0)$	86	Size-at-age.....	45
$E(B_{\text{current}+k}/B_{\text{current}})$	86	<i>In an age-based model</i>	45
$E(B_{\text{current}+k}/B_{\text{initial}})$	86	Size-weight relationship	46
Empirical randomisation of YCS	85	Size distribution	47
Expectation.....	86	Spawning stock biomass	25
First year for randomisation of YCS	86	Spawning stock biomass per recruit.....	87
Fishery Performance Indicators.....	84, 86	S-Plus/S/R functions	10
FPIs	84, 86	SSB	25
Future catches.....	86	Standard error.....	13
Lognormal randomisation of YCS	85	Standard output	13
Lognormal-empirical randomisation of YCS		State object.....	23
.....	85	Initial.....	37, 38

Stochastic Yields	89	System requirements	10
B_{MAY}	90	Time sequence.....	25
B_{MCY}	90	Time steps	23
Catch split.....	89	Transitions.....	23, 25, 27
CAY	89	Ageing	27
CSP.....	89	Growth	33
Deterministic recruitment.....	89	Maturation	31
F 's	89	Migration	32
F_{CAY}	90	Mortality	34
Harvest rate	89	Recruitment	28
MCY.....	89	$U_{max}(f)$	35
Multiple stocks	90	Version number.....	9
Point based	89	Weightless model.....	47
Recruitment variability.....	90	YCS	
Sample based.....	89	Haist parameterisation	29
Single stock	89	y_{enter}	28
Stochastic recruitment	89	Yield per recruit	87
Uncertainty in B_0	90	Curve	87
Uncertainty in free parameters	90	F at user specified mortality rate	88
Stock	24	$F_{0.1}$	88
Stock recruitment relationship	25	F_{max}	87
Steepness	28		

QUICK REFERENCE

Command line arguments

```
casal [-l] [-r] [-e] [-E] [-p] [-m] [-a number] [-C filelist  
-S outfile] [-s number prefix] [-v outfile] [-P outfile]  
[-Y] [-f prefix] [-F suffix] [-q] [-i infile]  
[-O outfile] [-o outfile] [-g RNG_seed] [-n name]
```

-l	Display the CASAL end user licence.
-r	Run the population section and calculate the objective function.
-e	Calculate the point estimate of the parameters.
-E	Using finite differences instead of automatic differentiation.
-p	Calculate likelihood or posterior profiles.
-m	Use MCMC.
-a [<i>number</i>]	Recover MCMC run; continue; and append further results.
-C [<i>filelist</i>]	Concatenate the MCMC results files for the specified files
-s [<i>no. prefix</i>]	Generate simulated observations, with number of simulations and outfile prefix
-v [<i>outfile</i>]	Output the values of the output quantities.
-P [<i>outfile</i>]	Calculate projected outputs (use with -i).
-Y	Calculate yields.

Optional command line arguments

-f [<i>prefix</i>]	Use a prefix on the names of the three input parameter files.
-F [<i>suffix</i>]	Replace the standard "csl" suffix used on the input parameter filenames with a user defined suffix.
-q	Run quietly.
-i [<i>file</i>]	Input free parameter values (use with -r, -e, -p, -m, -s, -v).
-O [<i>outfile</i>]	Output a set of free parameter values to a text file (use with -r, -e, or -E).
-o [<i>outfile</i>]	Output (with append) a set of free parameter values to a text file (use with -r, -e, or -E).
-S [<i>outfile</i>]	(With -C) dump the posterior sub-sample into <i>outfile</i> .
-g [<i>RNG_seed</i>]	With -m, -s, or -Y seed the random number generator.
-n [<i>name</i>]	(With -m or -a) specify the name of the current machine.

Order of transitions within a time step

1. Ageing (in an age-based model)
2. Recruitment
3. Maturation
4. Migration
5. Growth
6. Mortality (natural and fishing)
7. Disease mortality

Available ogives

`constant`: Has the estimable parameter C ; can be shifted; and can be used as a size-based ogive in an age-based model.

`knife_edge`: Has the non-estimable parameter E ; cannot be shifted; and can be used as a size-based ogive in an age-based model.

`allvalues`: Has estimable parameters $V_{low} V_{low+1} \dots V_{high}$; cannot be shifted; and cannot be used as a size-based ogive in an age-based model.

`allvalues_bounded`: Has non-estimable parameters L and H , and estimable parameters are $V_L V_{L+1} \dots V_H$; cannot be shifted; and cannot be used as a size-based ogive in an age-based model.

`logistic`: Has estimable parameters a_{50} and a_{1095} ; can be shifted; and can be used as a size-based ogive in an age-based model.

`logistic_capped`: Has estimable parameters a_{50} , a_{1095} , and a_{max} ; can be shifted; and can be used as a size-based ogive in an age-based model.

`logistic_bounded`: Has estimable parameters a_{50} and a_{1095} ; can be shifted; and can be used as a size-based ogive in an age-based model.

`double_logistic`: Has estimable parameters a_{50} , a_{1095} , b_{50} , b_{1095} , and a_{max} ; can be shifted; and can be used as a size-based ogive in an age-based model.

`logistic_product`: Has estimable parameters a_{50} , a_{1095} , b_{50} , b_{1095} , and a_{max} ; can be shifted; and can be used as a size-based ogive in an age-based model.

`double_normal`: Has estimable parameters a_1 , s_L , and s_R ; can be shifted; and can be used as a size-based ogive in an age-based model.

`double_normal_capped`: Has estimable parameters a_1 , s_L , s_R , and a_{max} ; can be shifted; and can be used as a size-based ogive in an age-based model.

`double_normal_plateau`: Has estimable parameters a_1 , a_2 , s_L , s_R , and a_{max} ; can be shifted; and can be used as a size-based ogive in an age-based model.

`double_normal_coleraine`: Has estimable parameters a_1 , σ_L^2 , and σ_R^2 ; can be shifted; and can be used as a size-based ogive in an age-based model.

`logistic_producing`: Has the non-estimable parameters L and H , and has estimable parameters a_{50} and a_{1095} ; cannot be shifted; and cannot be used as a size-based ogive in an age-based model.

`increasing`: Has non-estimable parameters L and H , and estimable parameters $\pi_L \pi_{L+1} \dots \pi_H$; cannot be shifted; and cannot be used as a size-based ogive in an age-based model.

`increasing_capped`: Has non-estimable parameters L , H , and C , and estimable parameters $\pi_L \pi_{L+1} \dots \pi_{H-1}$; cannot be shifted; and cannot be used as a size-based ogive in an age-based model.

Available penalties

`ogive_smoothing_penalty`: Sum of squares of r th differences applied to an `allvalues` or `allvalues_bounded` ogive parameter.

`catch_limit_penalty`: Sum of squares of (actual catch less specified catch), optionally on a log scale, for a single fishery.

`vector_average_penalty`: Square of $(\text{mean}(\text{vector}) - k)$, or of $(\text{mean}(\log(\text{vector})) - l)$, or of $(\log(\text{mean}(\text{vector}) / m))$ applied to a vector parameter

`vector_smoothing_penalty`: Sum of squares of r th differences applied to elements of a vector.

`element_difference_penalty`: Square of $(\text{vector}_1[i] - \text{vector}_2[i])$ applied to two vector parameters.

`YCS_difference_penalty`: Squared difference between the YCS values for a given year in the two stocks (for a two stock model).

`similar_qs_penalty`: Square of $(\log(q_i) - \log(q_j))$ applied to two relativity constants q

`ogive_comparison_penalty`: Sum of squares of $\max(\text{ogive}_1 - \text{ogive}_2, 0)$ applied to two ogive parameters.

`ogive_difference_penalty`: Square of $(\text{ogive}_1 - \text{ogive}_2)$ for a single size or age class, applied to two ogive parameters.

List of commands and sub-commands in the `population.cs1` data input file

Defining the partition

<code>@size_based</code>	Should the model be size-based rather than age-based?
<code>@n_classes</code>	Number of size classes
<code>@class_mins</code>	Size class lower limits (plus the upper limit of the last class if it is not a plus group)
<code>@min_age, @max_age</code>	Minimum and maximum age limits
<code>@plus_group</code>	Should a plus age or size group be used?
<code>@plus_group_size</code>	Mean size of plus group
<code>@sex_partition</code>	Is the partition sex-structured?
<code>@mature_partition</code>	Is the partition structured by maturity?
<code>@n_areas</code>	Number of areas in the partition
<code>@area_names</code>	Area names
<code>@n_stocks</code>	Number of stocks in the partition
<code>@stock_names</code>	Stock names
<code>@n_growthpaths</code>	Number of growth paths in the partition
<code>@exclusions_char1</code>	Partition exclusion term 1
<code>@exclusions_val1</code>	Partition exclusion value 1
<code>@exclusions_char2</code>	Partition exclusion term 2
<code>@exclusions_val2</code>	Partition exclusion value 2

Defining the annual cycle and the time sequence

<code>@initial</code>	Initial assessment year
<code>@current</code>	Current assessment year
<code>@final</code>	Final projection year
<code>@annual_cycle</code>	Annual cycle block command

time_steps	Number of time steps
recruitment_time	Time step in which recruitment occurs
recruitment_areas	Area in which where recruitment occurs, for each stock
spawning_time	Time step for recording SSB
spawning_part_mort	Proportion of mortality in the time step before recording SSB
spawning_areas	Area for recording SSB, for each stock
spawning_all_areas	Is SSB recorded for all areas combined?
spawning_ps	Spawning proportions by age/size class
spawning_p	Spawning proportion
spawning_use_total_B	Should SSB be defined as total biomass rather than mature biomass?
n_growths	Number of growth episodes per year
growth_times	Time step in which each growth episode occurs
aging_time	Time step when age is incremented
growth_props	Proportion of growth that has occurred by each time step
M_props	Proportion of natural mortality that occurs in each time step
baranov	Should fishing mortality be applied simultaneously with natural mortality using the Baranov equation, rather than instantaneously?
midmortality_partition	Method to calculate mortality within the time step
fishery_names	Names of the fishery
fishery_times	Time step when each fishery occurs
fishery_areas	Area when each fishery occurs
n_migrations	Number of migrations in each year
migration_names	Names of the fishery
migration_times	Time step of each migration
migrate_from	Area from which each migration departs
migrate_to	Area where each migration arrives
n_maturations	Number of maturation episodes in each year
maturation_times	Time step of each maturation episode
disease_mortality_time	Time step to apply disease mortality

Defining recruitment

@y_enter	Number of years after which a year class enters the partition
@standardise_YCS	Use the Haist parameterisation for YCS?
@use_mean_YCS	Use the Francis parameterisation for YCS?
@recruitment	Recruitment block command
YCS	Year class strengths for the stock
YCS_years	Years for which YCS are provided
n_initial	Number of years for which Rinitial is to be used as the YCS
SR	Stock-recruitment relationship
steepness	Steepness parameter of the stock-recruitment relationship
CR	Climate-recruitment relationship
Ts	Climate variable T
Ts_years	Years for which the climate variable T is provided
CR_alpha, CR_beta	Climate-recruitment parameters alpha and beta
CR_beta2	Climate-recruitment parameter beta2
CR_p	Climate-recruitment parameter p
initial_size_mean	Mean size at recruitment (both sexes)
initial_size_cv	C.v. of size at recruitment (both sexes)
initial_size_mean_male	Mean size at recruitment (male)
initial_size_mean_female	Mean size at recruitment (female)
initial_size_cv_male	C.v. of size at recruitment (male)
initial_size_cv_female	C.v. of size at recruitment (female)
p_male	Proportion of recruits that are male
growthpaths	Proportion of recruits on each growth path
first_free, last_free	Range of YCS defining R_0 , with the Haist or Francis YCS parameterisation

Defining recruitment variability

@randomisation_method	Randomisation method for recruitment variability in stochastic simulations and projections
@first_random_year	For projections, the first year for which YCS are randomised
<i>@recruitment sub-commands that define recruitment variability</i>	
sigma_r	Standard deviation on the log scale of randomised YCS
T_sigma_r	Standard deviation on the log scale of randomised climate data <i>T</i>
rho	Lag-1 log-scale autocorrelation of randomised YCS
T_rho	Lag-1 log-scale autocorrelation of randomised climate data <i>T</i>
year_range	Year range from which randomised YCS are resampled
T_year_range	Year range from which randomised climate data <i>T</i> are resampled
T_mean	Mean on the linear scale of randomised climate data <i>T</i>

Defining recruitment in yield simulations

simulation_SR	The stock-recruitment relationship to be used in yield simulations
simulation_steepness	Steepness parameter of the stock-recruitment relationship to be used in yield simulations

Defining growth (in a size based model)

@growth	Growth block command
stock	Stock that the growth episode applies to
type	Growth model used by the growth episode
g	Reference growths for the basic growth model
l	Reference sizes for the basic growth model
cv	C.v. for the basic growth model
minsigma	Lower bound on sigma for the basic growth model
g_male, g_female	Reference growths for the basic growth model
l_male, l_female	Reference sizes for the basic growth model
cv_male, cv_female	C.v. for the basic growth model
minsigma_male, minsigma_female	Lower bound on sigma for the basic growth model
g_mature, g_immature	Reference growths for the basic growth model
l_mature, l_immature	Reference sizes for the basic growth model
cv_mature, cv_immature	C.v. for the basic growth model
minsigma_mature	
minsigma_immature	Lower bound on sigma for mature/immature
g_male_mature, etc.	Reference growths for the basic growth model
l_male_mature, etc.	Reference sizes for the basic growth model
cv_male_mature, etc.	C.v. for the basic growth model
minsigma_male_mature, etc.	Lower bound on sigma

Defining maturation (when maturity is in the partition)

@maturation	Maturation block command
stock	Stock that the maturation episode applies to
area	Area that the maturation episode applies to
rates_all	Rates of maturation by age or size class
rates_male, rates_female	Rates of maturation by sex and age or size class

Defining maturity (when maturity is not in the partition)

@maturity_props	Maturity proportion block command
all	Maturity proportions by age/size class
male, female	Maturity proportions by sex and age/size class

Defining migrations

@migration	Migration block command
stock	Stock that migrates
migrators	Whether mature, immature, or both kinds of fish migrate
prop	Proportion of applicable fish that migrate
rates_all	Proportion of applicable fish that migrate, by age/size class

rates_male, rates_female	Proportion of applicable fish that migrate, by sex and age/size class
density_dependent	Are the migration rates density-dependent?
S, D	Source and destination density-dependence parameters
wave	Is this one wave of a 2-wave migration. If so, is it the first or second?
pwave	The proportion of fish in the first wave of a 2-wave migration
Defining natural mortality	
@natural_mortality	Natural mortality block command
all	The overall natural mortality rate
male, female	The male and female natural mortality rates
avg, diff	The male/female average and male-female difference in natural mortality rates
mature, immature	The mature and immature natural mortality rates
male_mature, etc.	Natural mortality rates by sex and maturity
ogive_all	The overall natural mortality rate as an ogive
ogive_male, ogive_female	The male and female natural mortality rates as ogives
ogive_avg, ogive_diff	The male/female average and male-female difference in natural mortality rates, as ogives
ogive_mature, ogive_immature	The mature and immature natural mortality rates, as ogives
ogive_male_mature, etc.	Natural mortality rates by sex and maturity, as ogives
Defining fishing mortality	
@fishery	Fishery block command
catches	Catches by year
years	Years for which catches are provided
selectivity	Name of the selectivity to use
F_max	Maximum fishing pressure (Baranov mortality)
U_max	Maximum fishing pressure (instantaneous mortality)
Fs	Instantaneous mortality F by year
Fs_years	Years for which F 's are provided
future_catches	Catches by year in the projection period
future_years	Years for which catches are provided
Defining disease mortality	
@disease_mortality	Disease mortality block command
DM	Disease mortality rate
selectivity	The selectivity ogive
years	Years to apply the disease mortality
index	Relative value of the disease mortality by year
Defining selectivities	
@selectivity_names	List of selectivity names
@selectivity	Selectivity block command
all	The selectivity ogive
male, female	The selectivity ogives by sex
mature, immature	The selectivity ogives by maturity
male_mature, etc.	The selectivity ogives by sex and maturity
shift_E	Exogenous selectivity shift variable E
shift_years	Years for which exogenous selectivity shift variable E is provided
shift_a	Exogenous selectivity shift parameter a
Setting the initial state	
@n_equilibrium	Number of years of running the equilibrium model
@Rinitial_is_deviat	Is $R_{initial}$ supplied relative to R_0 ?
@initialization	Initialization block command
B0	Equilibrium abundance B_0
R0	Equilibrium recruitment R_0

Bmean	Equilibrium abundance B_{mean} corresponding to R_{mean}
Rmean	Expected recruitment in any year R_{mean}
Binitial	Initial abundance $B_{initial}$
Rinitial	Initial recruitment $R_{initial}$
Cinitial	Initial number in each age/size class $C_{initial}$
Cinitial_male, Cinitial_female	Initial number in each age/size class $C_{initial}$ for each sex
@B0_total, @log_B0_total, @R0_total, @log_R0_total, @B0_prop_stock1, @R0_prop_stock1	Two stock model options for specifying initial state

Defining ogive preferences

@n_quant	Number of points at which to evaluate size-based ogives in an age-based model
----------	---

Defining size-at-age

@size_at_age_type	Size-at-age model type
@size_at_age_years	Years for which mean-size-at-age data are provided
@size_at_age_dist	Distribution of sizes-at-age around the mean
@size_at_age	Size-at-age block command
k, t0, Linf	von Bertalanffy parameters
k_male, t0_male, Linf_male, k_female, t0_female, Linf_female	von Bertalanffy parameters by sex
y1, y2, tau1, tau2, a, b	Schnute parameters
y1_male, y2_male, tau1_male, tau2_male, a_male, b_male, y1_female, y2_female, tau1_female, tau2_female, a_female, b_female	Schnute parameters by sex
male_[year], female_[year]	Mean-size-at-age of male and female fish in [year]
all_[year]	Mean-size-at-age of fish of both sexes in [year]
cv	c.v. of sizes-at-age around the mean
cv_male, cv_female	c.v. of sizes-at-age around the mean, by sex
@annual_growth	Use annual growth variation: amount of an average year's growth that occurs in each year
@annual_growth_years	Years for which annual growths are provided

Defining the size-weight relationship

@size_weight	Size-weight block command
type	The size-weight relationship function
a, b	The size-weight parameters a and b
a_male, b_male, a_female, b_female	The size-weight parameters a and b , by sex
verify_size_weight	Verify the supplied size-weight relationship and units
@weightless_model	Is this a model which does not involve fish weight?

List of commands and sub-commands in the `estimation.cs1` data input file

Defining the estimation method

@estimator	Choice of estimation method
@weighting	Choice of least-squares weighting
@k	Robustifying constant for least-squares weighting
@ko, @kp	Robustifying constants for least-squares weighting

Defining point estimation

@max_iters	Maximum number of iterations in the minimiser
@max_evals	Maximum number of evaluations in the minimiser
@max_iters_intermediate	Maximum number of iterations in early phases
@max_evals_intermediate	Maximum number of evaluations in early phases
@grad_tol	Minimiser convergence threshold

Defining likelihood or posterior profiling

@profile	Profile block command
parameter	Name of the parameter to be profiled
n	Number of values at which to profile the parameter
l, u	Range of values at which to profile the parameter

Defining MCMC

@MCMC	MCMC block command
start	Covariance multiplier for the starting point of the Markov chain
length	Length of the Markov chain
keep	Spacing between recorded values in the chain
max_cor	Maximum absolute correlation in the covariance matrix of the proposal distribution
stepsize	Initial stepsize (as a multiplier of the approximate covariance matrix)
adaptive_stepsize	Should the MCMC stepsize be altered during the chain?
adapt_at	At which iteration numbers can the MCMC stepsize be altered?
proposal_t	Should the proposal distribution be multivariate t ?
df	Degrees of freedom of the multivariate t proposal distribution.
burn_in	Number of samples to be discarded for the burn-in period
subsample_size	Size of the sub-sample to be generated
systematic	Should sub-sampling be systematic?
prior_reweighting	Should the sub-sample be generated using prior reweighting?
adaptive_covariance	Should the MCMC covariance matrix be altered during the chain?
adapt_covariance_at	At which iteration numbers can the MCMC covariance matrix be altered?
adaptive_covariance_discard	If the MCMC covariance matrix is altered during the chain, how many observations should be discarded from the start of the chain when taking a subsample for estimating the new covariance matrix?
adaptive_covariance_transitions	If the MCMC covariance matrix is altered during the chain, how many transitions must occur in the part of the chain used to estimate the new covariance matrix?
adaptive_covariance_stepsize	If the MCMC covariance matrix is altered during the chain, what is the new stepsize value?
@trivariate_normal_test	Test MCMC algorithm with a simple trivariate normal example

Defining the free parameters and priors

@estimate	Free parameter block command
parameter	Name of the parameter to be estimated
same	Names of the other parameters which are constrained to have the same value
phase	Phase at which this parameter should be estimated, in point estimation
lower_bound, upper_bound	Bounds on this scalar parameter
lower_bound, upper_bound	Bounds on this vector parameter
MCMC_fixed	Should this parameter be fixed during MCMC?
prior	What type of prior does this parameter have?
mu, cv	What are the mean and c.v. of this normal or lognormal prior on a scalar parameter?
mu, stdev	What are the mean and standard deviation of this normal-by-standard deviation or beta prior on a scalar parameter?
A, B	What are the lower and upper values for the range parameters of the beta prior?
m, s	What are the mean and standard deviation of the log of this scalar parameter, under the normal-log prior?
mu, cv	What are the mean and c.v. of each element of this normal or lognormal prior on a vector parameter?

mu, stdev	What are the mean and standard deviation of each element of this normal-by-standard deviation or beta prior on a vector parameter?
A, B	What are the lower and upper values for each element of the range parameters of the beta prior?
m, s	What are the mean and standard deviation of each element of the log of this vector parameter, under the normal-log prior?
mu, cv, rho	What are the mean, c.v., and ρ of this normal-AR prior on a vector parameter?
m, s, r	What are the log-scale mean, standard deviation, and ρ of this normal-log-AR prior on a vector parameter?
s, r	What are the log-scale standard deviation and ρ of this normal-log-mean1-AR prior on a vector parameter?

Defining the relativity constants q

@q_method	Method used for relativity constants q
@q	Relativity constant q block command
q	Value of the q parameter
b	Curvature parameter b associated with the q

Defining the observations

@abundance	Absolute abundance block command
@relative_abundance	Relative abundance block command
years	Years of the time series
step	Time step in which the observations occur
proportion_mortality	Proportion of the step's mortality, after which the observations occur
area	Area in which the observations occur
q	Relativity constant q to use
curvature	Should a curvature parameter be used?
biomass	Are the observations biomass rather than numbers of fish?
ogive	Which selectivity ogive should be applied?
[year]	Abundance for [year]
mature_only	Do these observations include mature fish only?
stock	Which stock do these observations relate to?
all_areas	Do these observations cover all areas in the model?
@numbers_at	Numbers_at block command
@relative_numbers_at	Relative numbers-at block command
@proportions_at	Proportions-at block command
years	Years of the time series
step	Time step in which the observations occur
proportion_mortality	Proportion of the step's mortality after which the observations occur
at_size	Are the observations by size?
sexed	Are the observations sexed?
area	Area in which the observations occur
q	Relativity constant q to use
ogive	Which selectivity ogive should be applied?
class_mins	What are the size bins of the observations (in an age-based model)?
class_nums	What are the class numbers for the observations (in a size-based model)?
class_nums_male, class_nums_female	What are the class numbers for the sexed observations (in a size-based model)?
plus_group	Is the last age or size class a plus group?
min_class, max_class	Which age/size classes are covered by the observations?
sum_to_one	Should the proportions sum to 1?
ageing_error	Should ageing error be applied to these observations?
[year]	Numbers or proportions for [year]

@catch_at	Catch_at block command
years	Years of the time series
fishery	Fishery or fisheries covered by the observations
at_size	Are the observations by size?
sexed	Are the observations sexed?
class_mins	What are the size bins of the observations (in an age-based model)?
class_nums	What are the class numbers of the observations (in a size-based model)?
class_nums_male, class_nums_female	What are the class numbers of the sexed observations (in a size-based model)?
plus_group	Is the last age or size class a plus group?
min_class, max_class	Which age/size classes are covered by the observations?
sum_to_one	Should the proportions sum to 1?
ageing_error	Should ageing error be applied to these observations?
[year]	Numbers or proportions for [year]
@proportions_mature	Proportions_mature block command
years	Years of the time series
step	Time step in which the observations occur
proportion_mortality	Proportion of the step's mortality, prior to when the observations occur
sexed	Are these observations sexed?
females_only	Are these observations for females only?
at_size	Are the observations by size?
area	Area in which the observations occur
ogive	Which selectivity ogive should be applied?
class_mins	What are the size bins of the observations (in an age-based model)?
plus_group	Is the last age or size class a plus group?
min_class, max_class	Which age/size classes are covered by the observations?
ageing_error	Should ageing error be applied to these observations?
[year]	Proportions mature for [year]
@proportions_migrating	Proportions_migrating block command
years	Years of the time series
migration	Migration to which the observations apply
sex	Which sex do the observations apply to?
at_size	Are the observations by size?
area	Area in which the observations occur
ogive	Which selectivity ogive should be applied?
class_mins	What are the size bins of the observations (in an age-based model)?
plus_group	Is the last age or size class a plus group?
min_class, max_class	Which age/size classes are covered by the observations?
ageing_error	Should ageing error be applied to these observations?
[year]	Proportions migrating for [year]
@age_size	Age_size block command
year	Year in which the data were collected
step	Time step in which the data were collected
proportion_mortality	Proportion of the step's mortality, prior to when the observations occur
area	Area in which the observations occur
stock	Stock for which the data were collected
sample	Sampling method under which the observations were generated
ogive	Which selectivity ogive should be applied?
ageing_error	Should ageing error be applied to these observations?
ages	Age data
sizes	Size data
sexes	Sex data
@age_at_maturation	Age_at_maturation block command

sexed	Are these observations sexed?
sampled_ages	What were the estimated ages of these fish at sampling?
maturation_ages	What were the estimated ages of these fish at maturation?
sexes	What were the sexes of these fish?
stock	Which stock do these observations relate to?
ageing_error	Should ageing error be applied to these observations?
k	After how many years can maturation be detected?

Defining the objective function associated with the observations

weight	Weight of this time series, in the Cordue weighted least-squares scheme
cv_[year]	C.v.s by year for this time series, in the Cordue weighted least-squares scheme
dist	Likelihood of the observations
r	Robustifying constant
cv	C.v. for all observations in this time series, used with likelihoods
cv_[year]	C.v.s by year for this time series, used with likelihoods
cvs_[year]	C.v.s by observation and year for this time series, used with likelihoods
cv_process_error	Process error c.v. for this time series, used with likelihoods parameterised by the c.v.
stdev	Standard deviation for all observations in this time series, used with normal-by-standard deviation likelihoods
stdev_[year]	Standard deviation by year for this time series, used with normal-by-standard deviation likelihoods
stdevs_[year]	Standard deviation by observation and year for this time series, used with normal-by-standard deviation likelihoods
stdev_process_error	Process error standard deviation for this time series, used with likelihoods parameterised by the standard deviation
N	N for all years in this time series, used with proportions likelihoods
N_[year]	N 's by year for this time series, used with proportions likelihoods
Ns_[year]	N 's by observation and year for this time series, used with proportions likelihoods
N_process_error	Process error N for this time series, used with likelihoods parameterised by the effective sample size.

Defining the penalties

@ogive_smoothing_penalty	Ogive smoothing penalty block command
label	The name of the penalty
ogive	The name of the ogive parameter to which the penalty is applied
r	Penalty is applied to r th differences
lower_bound, upper_bound	Penalty is applied for age or size classes lower_bound to upper_bound
multiplier	Multiply the penalty by this factor
@catch_limit_penalty	Catch limit penalty block command
label	The name of the penalty
fishery	The label of the fishery to which the penalty is applied
log_scale	Should sums of squares be calculated on the log scale?
multiplier	Multiply the penalty by this factor
@vector_average_penalty	Vector average penalty block command
label	The name of the penalty
vector	The name of the vector parameter to which the penalty is applied
k, l, m	Vector should average arithmetically to k or m , or geometrically to l .
multiplier	Multiply the penalty by this factor

@vector_smoothing_penalty	Vector smoothing penalty block command
label	The name of the penalty
ogive	The name of the vector parameter to which the penalty is applied
r	Penalty is applied to rth differences
lower_bound, upper_bound	Penalty is applied for index lower_bound to upper_bound
multiplier	Multiply the penalty by this factor
@element_difference_penalty	Element difference penalty block command
label	The name of the penalty
vector1, vector2	The name of the vector parameters to which the penalty is applied
i	Penalise differences in the ith elements of the two vectors
multiplier	Multiply the penalty by this factor
@YCS_difference_penalty	YCS difference penalty block command
label	The name of the penalty
stock1, stock2	The names of the stocks to which the penalty is applied
year	Year for which the penalty is to be applied
multiplier	Multiply the penalty by this factor
@similar_qs_penalty	Similar q 's penalty block command
label	The name of the penalty
q1, q2	The names of the two q 's to which the penalty is applied
multiplier	Multiply the penalty by this factor
@ogive_comparison_penalty	Ogive comparison penalty block command
label	The name of the penalty
ogive1, ogive2	The name of the ogive parameters to which the penalty is applied
lower_bound, upper_bound	Penalty is applied for age or size classes lower_bound to upper_bound
multiplier	Multiply the penalty by this factor
@ogive_difference_penalty	Ogive difference penalty block command
label	The name of the penalty
ogive1, ogive2	The name of the ogive parameters to which the penalty is applied
class	The age, or size class number, to which the penalty is applied
multiplier	Multiply the penalty by this factor

Defining the ageing error

@ageing_error	Ageing error block command
type	Type of ageing error model
p1, p2, k	Parameters of the "off_by_one" ageing error model
c	Parameter of the "normal" ageing error model
[age]	Row of the misclassification matrix for the "misclassification_matrix" ageing error model.

CASAL extensions

@user_parameterisation	Is there a reparameterisation of the population section using user.parameterisation.C?
@user_components	Lists the names of the user-defined priors or penalties calculated in user.prior_penalty.C

List of commands and sub-commands in the output .cs1 data input file

Defining the printouts

@print	Printouts block command
parameters	Print the population, estimation, and output parameters?
unused_parameters	Print a list of the parameters that were never used?
population_section	Print a description of the population section?
requests, results	Print a description of the requests sent to the population section and the corresponding results?

initial_state, state_annually, state_every_step, final_state	Print the state of the population?
estimation_section	Print a description of the estimation section?
fits, resids, pearson_resids, normalised_resids	Print the fits, residuals, and standardised residuals?
covariance	Print the approximate covariance matrix of the free parameters
yields	Print a description of the yield calculations?
fits_every_eval, objective_every_eval, parameters_every_eval, parameter_vector_every_eval	Print the fits, objective function, or parameters at every function evaluation?
@print_sizebased_ogives_at	Sizes for which size-based ogives should be printed in an age-based model

Defining the output quantities

@quantities	Output quantities block command
all_free_parameters	Output quantities include all free parameters?
scalar_parameters, vector_parameters, ogive_parameters	Output quantities include these parameters
ogive_arguments	Output quantities include these ogive arguments
nuisance_qs	Output quantities include nuisance q 's?
B0, R0, Bmean, Rmean Binitial, Rinitial	Output quantities include B0, R0, Bmean, Rmean, Binitial, Rinitial?
SSBs	Output quantities include SSBs?
actual_catches, actual_catches_by_stock	Output quantities include actual catches (by stock)?
recruitments	Output quantities include recruitments?
YCS	Output quantities include YCS?
true_YCS	Output quantities include 'true YCS' = $YCS \times SR \times CR$?
Ts	Output quantities include climate variable T ?
fishing_pressures	Output quantities include fishing pressures?
stock_crash	Output quantities include 'stock crash'?
fits	Output quantities include the 'fits' or the expected value for each observation?
resids	Output quantities include the 'residuals' or the observed less expected value for each observation?
pearson_resids	Output quantities include the 'Pearson residuals' for each observation?
normalised_resids	Output quantities include the 'normalised residuals' for each observation?
@selectivity_at	Selectivity-at block command
ogive	Which selectivity should be applied?
years	Years of the time series
step	Time step in which the observations occur
proportion_mortality	Proportion of the step's mortality after which the observations occur
sexed	Are the observations sexed?
area	Area in which the observations occur
mature_only	Do these observations include mature fish only?
stock	Which stock do these observations relate to?
[year]	Numbers or proportions for [year]

Defining projections

@n_projections	Number of projections to be done (from a point estimate)
----------------	--

Defining yield calculations

@catch_split	Catch split used in yield simulations
@deterministic_yields_mortality_rate, @MCY_CAY_mortality_rate	Definition of the mortality rate F used in deterministic yield calculations and in CAY calculations
@B_pre	Pre-fishery biomass B_{pre} block command

mature_only	Pre-fishery biomass B_{pre} is mature fish only?
area	Area for which pre-fishery biomass B_{pre} is calculated
step	Time step in which pre-fishery biomass B_{pre} is calculated
proportion_mortality	Proportion of the time step's mortality after which pre-fishery biomass B_{pre} is calculated
selectivity	Selectivity ogive with which pre-fishery biomass B_{pre} is calculated
Defining deterministic yields	
@per_recruit	Per-recruit block command
do_YPR_SPR	Supply data to plot a YPR or SPR curve?
F	F's at which to calculate YPR & SPR
do_Fmax	Calculate F_{max} ?
do_F0_1	Calculate $F_{0.1}$?
do_Fx	Calculate $F_{x\%}$?
x	x at which to calculate $F_{x\%}$
guess	First guess at F_{max} , $F_{0.1}$, $F_{x\%}$
@deterministic_MSX	Deterministic MSY block command
do_MSX	Calculate deterministic MSY?
do_yield_vs_SSB	Supply data to plot a yield versus SSB curve?
F	F's at which to calculate yield and SSB
guess	First guess at F_{MSYdet}
Defining stochastic yields	
@MCY_CAY	MCY/CAY block command
do_MCY	Calculate MCY?
do_CAY	Calculate CAY?
interactive	Should MCY/CAY be calculated interactively?
p, q	Risk constraints in MCY/CAY analysis
n_simulations	Number of simulations for each harvest rate H in MCY/CAY analyses (point-based only)
n_discard, n_keep	Number of years to discard and to keep in each MCY/CAY simulation.
max_upper_iter	Maximum number of times upper bound on H can be increased when searching for an optimal yield.
MCY_uncertainty_dist	Distribution of uncertainty in virgin biomass (point-based only)
MCY_uncertainty_cv	C.v. of uncertainty in virgin biomass (point-based only)
MCY_guess	First guess at MCY
F_CAY_guess	First guess at F_{CAY}
@CSP	CSP block command
do_CSP	Calculate CSP?
individual_stocks	Calculate CSP by individual stock?
CSP_guess	First guess at CSP
@B_post	Post-fishery biomass B_{post} block command
mature_only	Post-fishery biomass B_{post} is mature fish only?
area	Area for which post-fishery biomass B_{post} is calculated
step	Time step in which post-fishery biomass B_{post} is calculated
proportion_mortality	Proportion of the time step's mortality after which post-fishery biomass B_{post} is calculated
selectivity	Selectivity ogive with which post-fishery biomass B_{post} is calculated