

MARS Leeds Flight Simulator –How to use the optimisation facilities in VENSIM.

Author : Simon Shepherd, Institute for Transport Studies, University of Leeds, UK.

Contact : S.P.Shepherd@its.leeds.ac.uk

March 2008

The purpose of this note is to demonstrate how the optimisation facilities within the VENSIM tool can be used with the MARS model to look at integrated transport packages. It describes two basic approaches. The first is using the “policy” optimisation facility where an objective function is given to the model (in this case a simple MCA score) and the policy combination selected is that which maximises the given objective. The second approach uses the “calibration” mode of the optimisation tool to generate a policy which gives the best fit to a path of CO₂ emission targets over the planning period. It should be noted that while these optimisation features would only be available to those with a full VENSIM license, the LeedsFS can be used to simulate policy combinations which can then be compared in terms of performance against the CO₂ trajectories or the “Test MCA” described below as these indicators are included in the standard simulator outputs.

Policy Optimisation with an example MCA

In MARS a policy profile is used to implement a change in policy over the planning period (in this case 30 years). A policy “level” can be changed in the short-run year or in the long run year with the intermediate values determined by interpolation between these levels. In our example the short-run year is year 5 and the long run year is year 30 or the end of the simulation period.

Table 1 shows the policy instruments and lower and upper bounds used in the example.

Instrument	Lower bound	Upper bound
Cordon charge peak	0 euros	10 euros
Cordon charge off-peak	0 euros	5 euros

Bus fares peak	-50%	+100%
Bus fares off-peak	-50%	+100%
Bus frequency peak	0%	+100%
Bus frequency off peak	0%	+100%
Bus Quality	-20 cents	+20 cents
Road capacity	-5%	+5%

Table 1 : Instruments and bounds

As there are eight instruments and each has a short run and long run value then the optimisation process is dealing with sixteen variables.

The MCA should in principle be developed with stakeholders using appropriate indicators, targets and weight setting procedures. As this is only an example of how to apply the optimisation process within VENSIM we have taken three indicators from DISTILLATE Project C to represent economic efficiency, safety and the environment. These are total time spent travelling in the peak, total accidents and total CO₂ emitted from well to wheel for all modes respectively.

Rather than go to stakeholders to determine weights for this simple example we simply use the following sum of relative changes in the indicators over time to score the strategies :-

$$MCA = \sum_i \sum_t w_i \left(\frac{y_{it} - y_{i0}}{y_{i0}} \right)$$

Where y_{it} is the value of indicator t at time t and w_i is the weight given to indicator i .

To begin with we use equal weights between indicators. Note that a higher value at time t is assumed to be worse so we wish to minimise the MCA value.

Figure 1 shows the simple view in VENSIM where this MCA value can be calculated via existing indicators. It also highlights the optimisation tool which initiates the set up for a calibration or optimisation.

Figure 2 shows how simple it is to change an equation within VENSIM – it shows the equation editor view when the variable “TestMCA” from figure 1 is opened with the equation tool. As can be seen here this screen shot was taken with a weight of 1.9 applied to relative change in total time in the peak from year 0 (weights are varied for total travel time in what follows for illustrative purposes).

Figures 3 and 4 show the optimisation set up views which appear when the optimisation tool button is selected (see figure 1). The first view is used to select either a policy or calibration type of optimisation. In this case we are using a policy optimisation and the user can select a number of variables and set different weights for each variable using the select and edit buttons. Here we have selected “TestMCA” and assigned a negative weight – as the optimisation actually maximises the objective function given so we wish to maximise the negative of the MCA value which is equivalent to minimising the MCA value.

Figure 4 shows the second set up screen where the variables to be optimised and their lower and upper bounds are entered as in table 1 above. Again this process is done via the edit and select buttons but once completed the parameters can be saved in a control file (see top left filename). This avoids having to repeat the task for similar optimisation runs.

Figure 5-8 show the indicators for accidents, total time in peak, CO2 well to wheel and MCA values over the 30 year period for the Do-nothing and the optimal solution. The optimal solution included full 50% fare reductions in both peak and off peak, no changes in bus frequency, the maximum allowed increase for cordon charges (peak and off-peak) and bus quality and a reduction of 5% in road capacity in both short and long run years. All values are on a lower or upper bound in this case. It is worth pointing to the results for changes in bus frequency which are left unchanged – this is due to increases in CO2 emitted when frequencies are increased which are not out-weighted by reductions in accidents or time spent in the peak.

From figures 5-7 we can see that with equal weights given to the three indicators we have a reduction in accidents and CO2 whilst the time spent in the peak is increased. There are a number of reasons for this result – firstly if average speed increases then accidents increase in this model, secondly if the car share were to increase then CO2 emitted would also increase so there are trade-offs being made during the optimisation. Increasing the weight

on total time spent in the peak results in the same policy combination until a weight greater than 1.75 is used. With a weight=2 for total time in the peak the policy combination is the same as above but with increased fares in the peak and slight increases in road capacity. This results in a decrease in total time in the peak at the expense of accidents which now increase (see figures 9 and 10).

Calibration option using CO2 target paths

A similar process to the optimisation above can be applied using the calibration option. This time the calibration radio button is used in the set up view (see figure 11) and an alternate payoff function is included – the total CO2 emitted well to wheel.

The other set up views (not shown) basically define the same variables and bounds to be optimised and in the final view the target trajectory for CO2 emissions is included in a data file.

Two target trajectories were used in this case study based on Option 2 of Tight et al (2005) who set target reductions for the transport sector of 33% and 67% by 2050. These targets allow for a growth in the transport share of emissions of CO2 for the UK over the period. Table 2 shows the target trajectories assumed between 2010 and 2050 noting that these were not considered in the paper by Tight et al so they are used as examples only here.

Percentage reduction relative to 2000 levels	2010	2020	2030	2040	2050
Low target	10	15	20	28	33
High target	20	25	30	40	67

Table 2 : Target trajectories assumed for total CO2 well to wheel from transport

Note that the MARS model only models the period from 2000 to 2030 and that policies are assumed to start in 2005. Hence only the values from 2010-2030 are relevant here and it should also be noted that these values were actually implemented as changes relative to the base year 2000 rather than 1997 as in Tight et al.

Reading in the target trajectories can be done by first of all using the “Model-import dataset command” which converts the data from an excel workbook to a VENSIM compatible vdf file. This can then be viewed for the relevant variable and is also used in the final set up screen of the optimisation process to define the comparator for the “calibration”.

The calibration option effectively minimises the difference between the output CO2 path and the target path at the given points by varying the policy combinations. Figure 12 shows the Do-nothing, the high target path and the optimal solution. From this we can see that for the given policy instruments (which are the same as for the MCA example) the optimal policy cannot meet the high targets for CO2. In fact the optimal policy is the same as for the initial MCA optimisation described above. It may be that other policy instruments such as fuel tax may be able to meet these targets.

Figure 13 shows the results of applying the same optimisation/calibration but using the lower CO2 targets. This time the targets in 2010 and 2020 are almost met. In fact the 2010 target is almost achieved in the Do-nothing as a result of improved vehicle fleet and fuel efficiency. This means that the optimal policy differs initially from earlier runs. For example the initial values for fares in the peak remain near zero but eventually fares are reduced by 50%. In addition the peak cordon charge begins in year 5 with a value of 1 euro rising to 10 euros by year 30. Bus quality also rises from around 0 to 20 cents by year 30. By the end of the period it is still the case that the target cannot be met and the final policy values are the same as in the previous runs.

This analysis shows that either greater improvements in fleet technology or other policies must be implemented beyond 2020 to meet even the lowest of CO2 emission targets.

Summary

This note has shown how it is possible to use the VENSIM optimisation facilities with simple examples whereby combinations of instruments can be found which either maximise some objective function (in this case an MCA) or minimise the difference in some output indicator trajectory compared to a required or target path (CO2) using the calibration mode. It is not intended as a policy paper but shows the potential of the tool. It should be noted that a full VENSIM license is required to conduct full optimisations and that each of the above optimisations required around 200+ MARS runs taking around 2-3 hours on a standard PC.

Even without a full licence the MARS FS can be used to examine user defined combinations of policies and standard outputs include the indicators used in the above examples.

References

Tight, M.R, Bristow, A.L., Pridmore, A. and May, A.D. (2005) What is a sustainable level of CO2 emissions from transport activity in the UK in 2050? *Transport Policy* 12 (2005) 235–244.

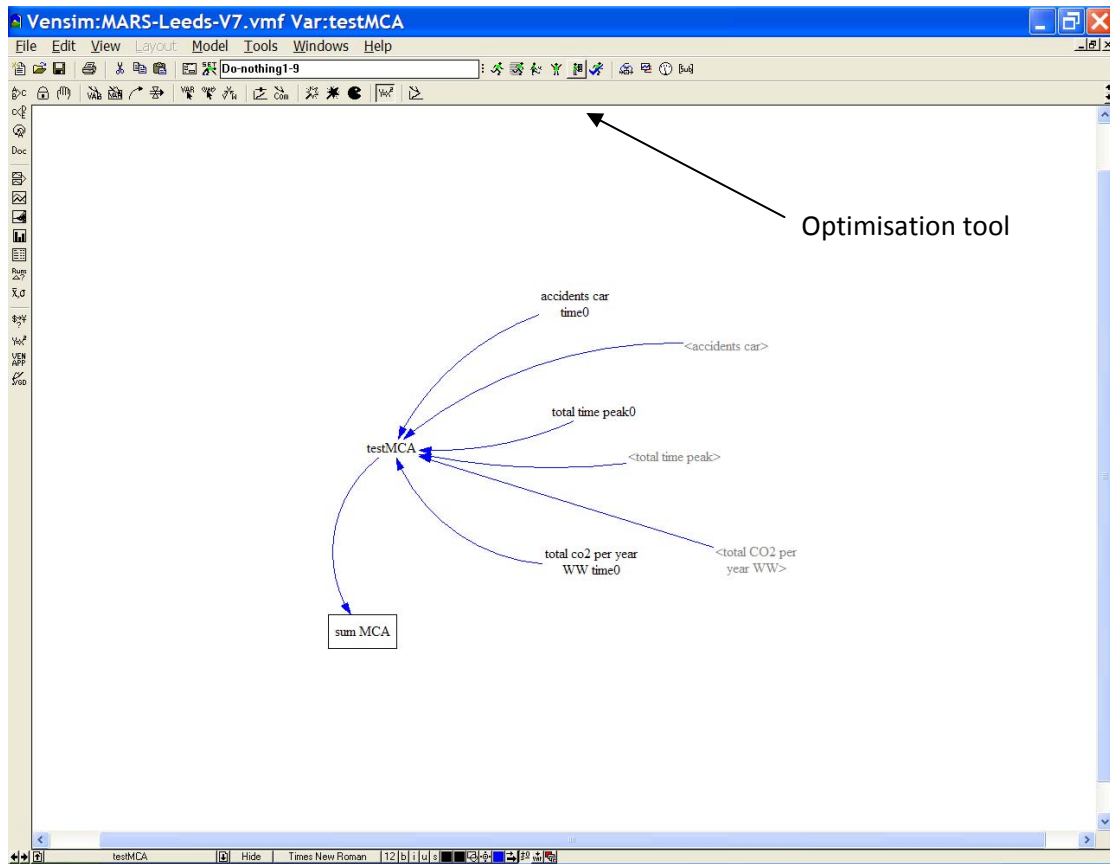


Figure 1 : VENSIM view of simple MCA and optimisation tool button

The 'Editing equation for - testMCA' dialog box shows the following equation:

$$= \frac{\text{accidents car} - \text{accidents car time0}}{\text{accidents car time0}} + \frac{\text{total CO2 per year WW} - \text{total co2 per year WW time0}}{\text{total co2 per year WW time0}} + 1.9 \cdot \frac{\text{total time peak} - \text{total time peak0}}{\text{total time peak0}}$$

The dialog includes a variable list with the following items: accidents car, accidents car time0, total CO2 per year WW, total co2 per year WW time0, total time peak, and total time peak0. The 'Errors' field shows 'Equation OK'.

Figure 2 : Equation editor for inputting MCA calculation

Optimization Setup

Payoff Definition. Edit the filename to save changes to a different control file

Filename:

Type Calibration Policy

Payoff Elements

testMCA/-1

Variable Compare to is used only for calibration payoffs

Compare to

Weight

The weight should be positive for calibration. For policy optimizations use a positive number when more is better and a negative number when less is better.

Figure 3 : Optimisation set up view 1

Optimization Setup

Optimization Control. Edit the filename to save changes to a different control file

Filename:

Output Level Trace Sensitivity =

Multiple Start Random type Seed

#Restart Optimizer Max Iterations Max Sims

Pass Limit Fractional Tolerance Tolerance Multiplier

Absolute Tolerance Scale Absolute Vector Points

Currently active parameters (drag to reorder)

- 50<=start value[bus fares peak]<=100
- 50<=end value[bus fares peak]<=100
- 0<=end value[road charge peak]<=10
- 50<=start value[bus fares opeak]<=100
- 50<=end value[bus fares opeak]<=100
- 0<=start value[road charge opeak]<=5
- 0<=end value[road charge opeak]<=5
- 0<=start value[bus freq peak]<=100

<= = <=

Model value of constant -- =

Figure 4 : Optimisation set up view 2

Car Accidents




Figure 5 : Accidents by car for the do-nothing and optimal solution

Total CO2 emissions WW

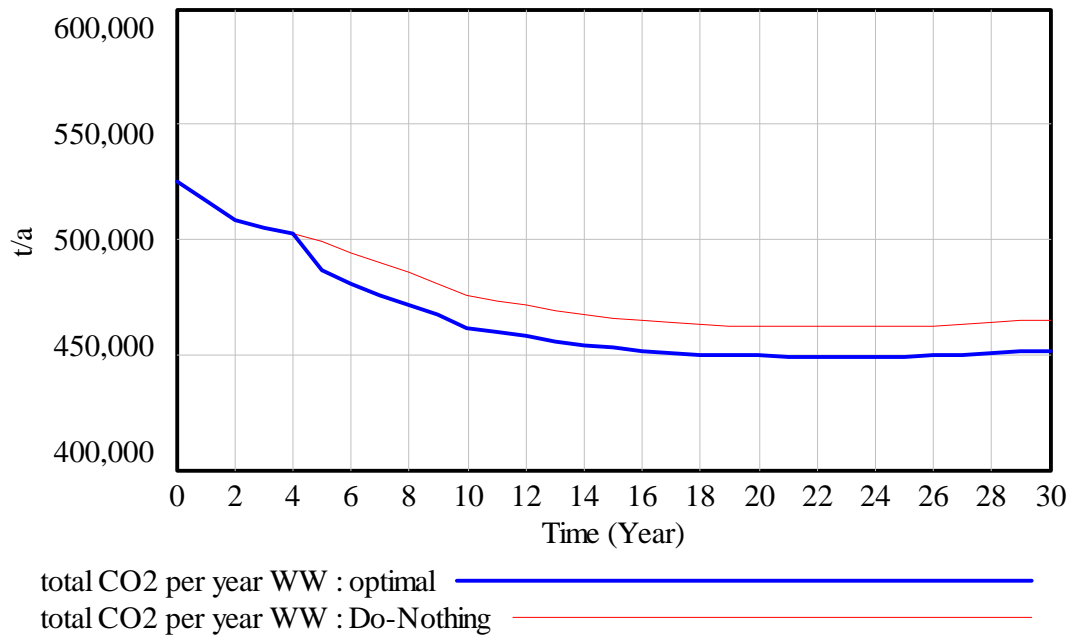


Figure 6 : CO2 emissions well to wheel

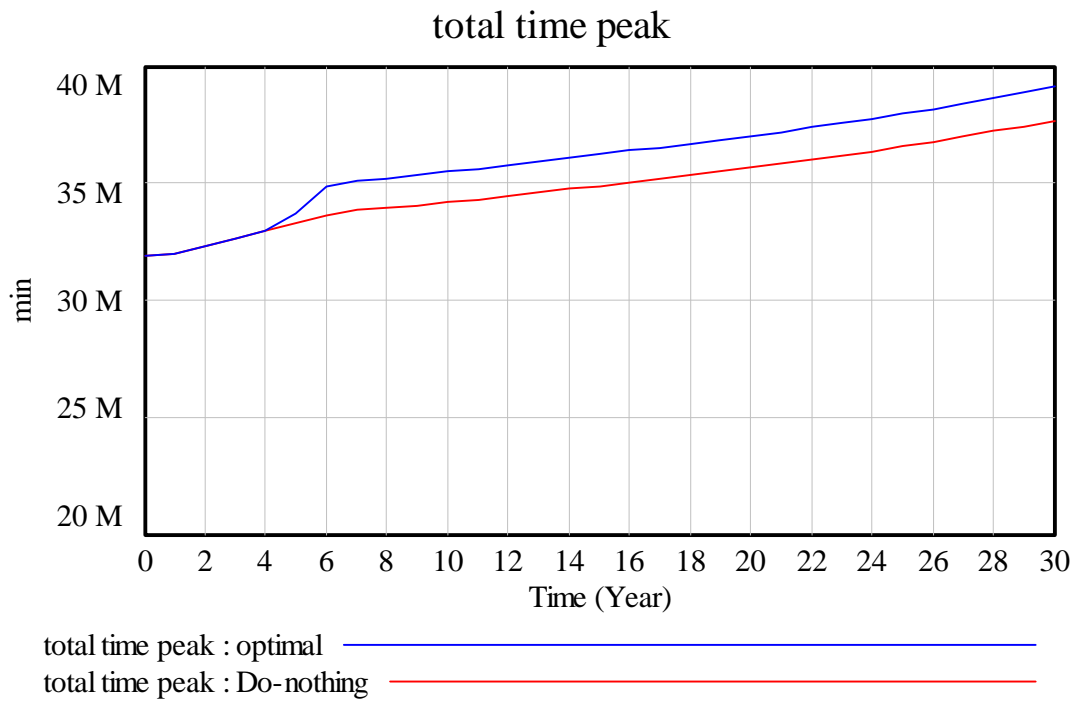


Figure 7 : Total time in peak for all modes

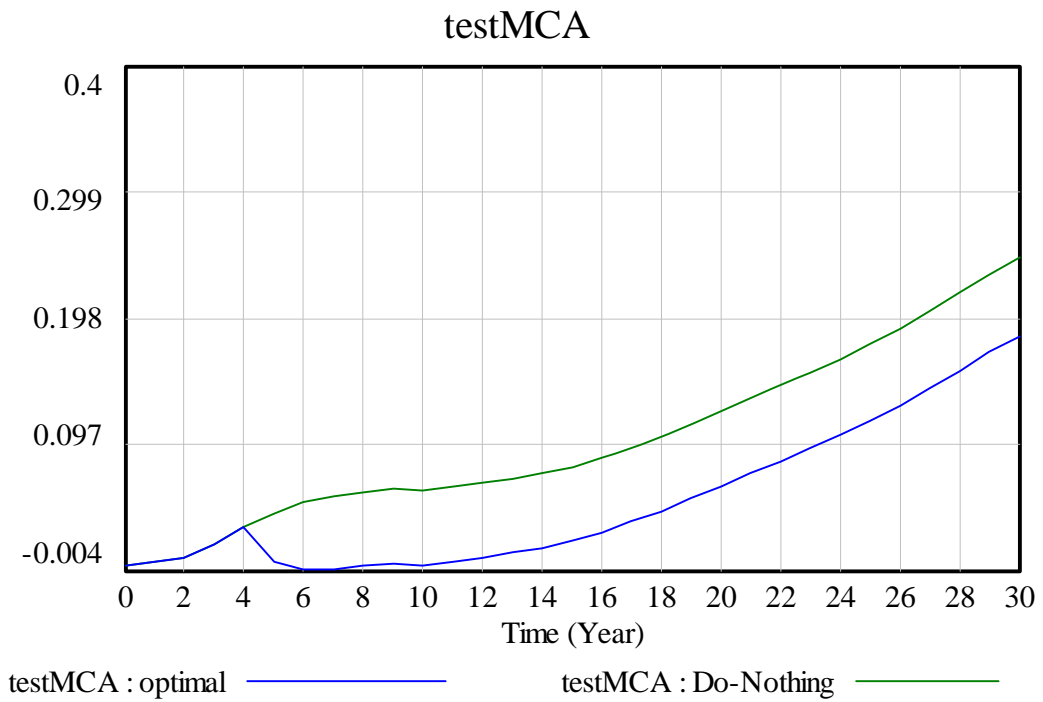


Figure 8 : Comparison of MCA values over time for Optimal policy and do-nothing.

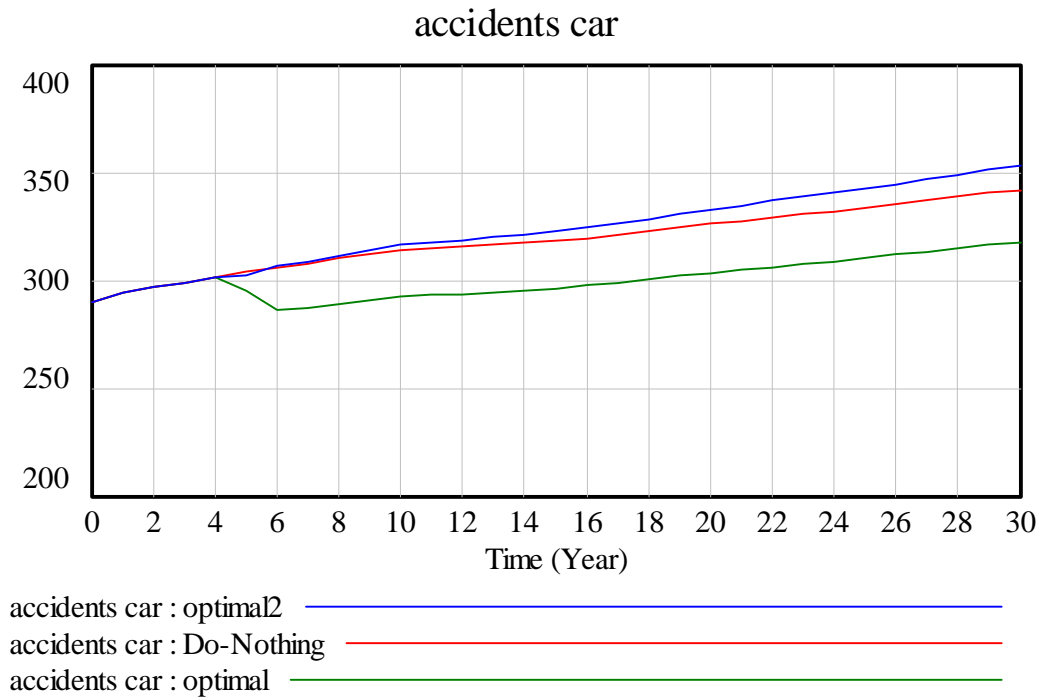


Figure 9 : Accidents with optimal2 weight=2 on total time peak, optimal weights=1

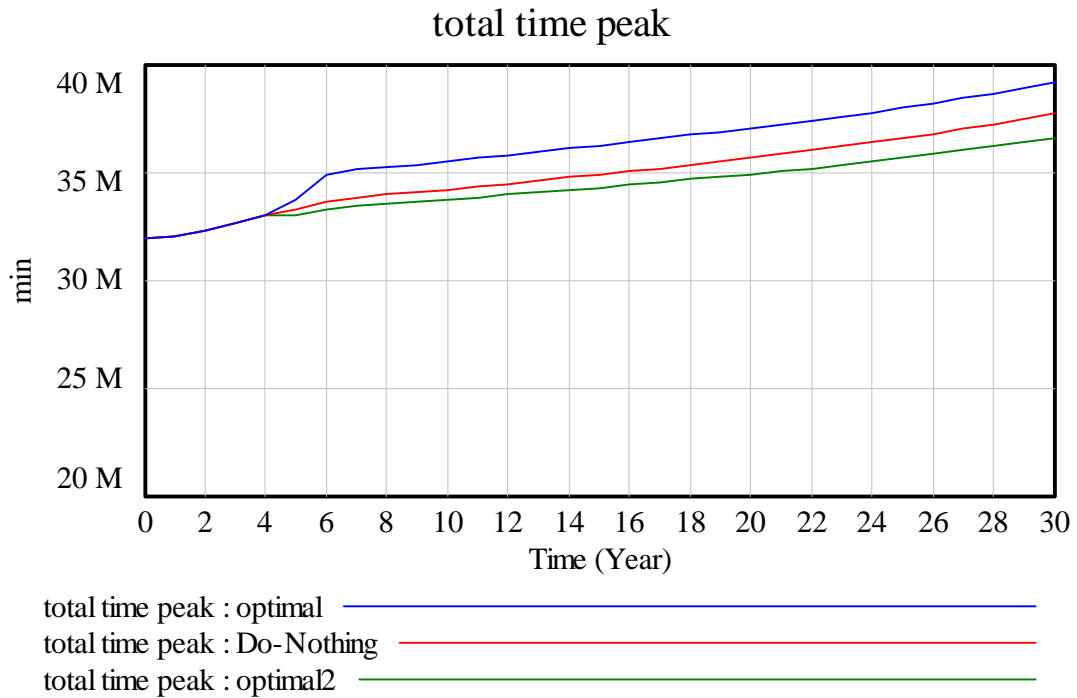


Figure 10 : Total time peak with optimal2 weight=2 on total time peak, optimal weights=1

Optimization Setup

Payoff Definition. Edit the filename to save changes to a different control file

Filename:

Type Calibration Policy

Payoff Elements

total CO2 per year WW/1

Variable Compare to is used only for calibration payoffs

Compare to

Weight

The weight should be positive for calibration. For policy optimizations use a positive number when more is better and a negative number when less is better.

Figure 11 : Calibration set up view




Figure 12 : High target scenario with target trajectory (test-calib), do-nothing, optimal solution (which cannot meet the target path).

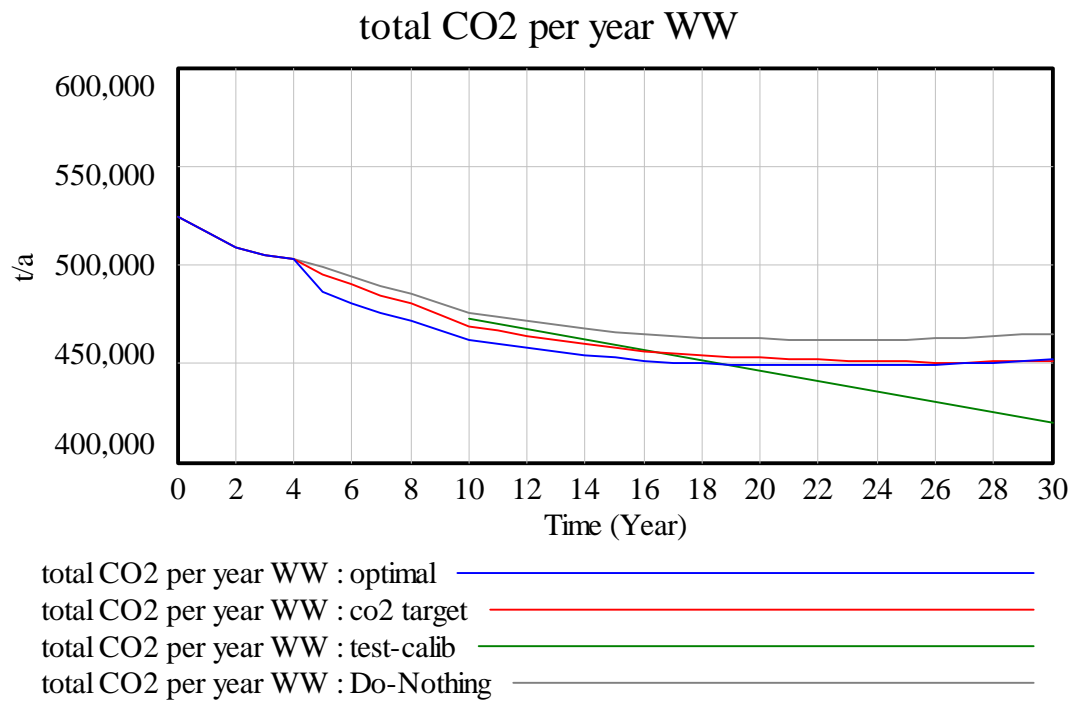


Figure 13 : Low target scenario with target trajectory (test-calib), do-nothing, optimal solution from MCA and optimal solution for CO2 target trajectory.