
AIMMS Modeling Guide - Sensitivity Analysis

This file contains only one chapter of the book. For a free download of the complete book in pdf format, please visit www.aimms.com or order your hard-copy at www.lulu.com/aimms.

Copyright © 1993–2011 by Paragon Decision Technology B.V. All rights reserved.

Paragon Decision Technology B.V.	Paragon Decision Technology Inc.	Paragon Decision Technology Pte.
Schipholweg 1	500 108th Avenue NE	Ltd.
2034 LS Haarlem	Ste. # 1085	80 Raffles Place
The Netherlands	Bellevue, WA 98004	UOB Plaza 1, Level 36-01
Tel.: +31 23 5511512	USA	Singapore 048624
Fax: +31 23 5511517	Tel.: +1 425 458 4024	Tel.: +65 9640 4182
	Fax: +1 425 458 4025	

Email: info@aimms.com
WWW: www.aimms.com

AIMMS is a registered trademark of Paragon Decision Technology B.V. IBM ILOG CPLEX and sc CPLEX is a registered trademark of IBM Corporation. GUROBI is a registered trademark of Gurobi Optimization, Inc. KNITRO is a registered trademark of Ziena Optimization, Inc. XPRESS-MP is a registered trademark of FICO Fair Isaac Corporation. MOSEK is a registered trademark of Mosek ApS. WINDOWS and EXCEL are registered trademarks of Microsoft Corporation. \TeX , \LaTeX , and $\AMS-\LaTeX$ are trademarks of the American Mathematical Society. LUCIDA is a registered trademark of Bigelow & Holmes Inc. ACROBAT is a registered trademark of Adobe Systems Inc. Other brands and their products are trademarks of their respective holders.

Information in this document is subject to change without notice and does not represent a commitment on the part of Paragon Decision Technology B.V. The software described in this document is furnished under a license agreement and may only be used and copied in accordance with the terms of the agreement. The documentation may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from Paragon Decision Technology B.V.

Paragon Decision Technology B.V. makes no representation or warranty with respect to the adequacy of this documentation or the programs which it describes for any particular purpose or with respect to its adequacy to produce any particular result. In no event shall Paragon Decision Technology B.V., its employees, its contractors or the authors of this documentation be liable for special, direct, indirect or consequential damages, losses, costs, charges, claims, demands, or claims for lost profits, fees or expenses of any nature or kind.

In addition to the foregoing, users should recognize that all complex software systems and their documentation contain errors and omissions. The authors, Paragon Decision Technology B.V. and its employees, and its contractors shall not be responsible under any circumstances for providing information or corrections to errors and omissions discovered at any time in this book or the software it describes, whether or not they are aware of the errors or omissions. The authors, Paragon Decision Technology B.V. and its employees, and its contractors do not recommend the use of the software described in this book for applications in which errors or omissions could threaten life, injury or significant loss.

This documentation was typeset by Paragon Decision Technology B.V. using \LaTeX and the LUCIDA font family.

Chapter 4

Sensitivity Analysis

The subject of this chapter is the introduction of *marginal values* (shadow prices and reduced costs) and *sensitivity ranges* which are tools used when conducting a sensitivity analysis of a linear programming model. A sensitivity analysis investigates the changes to the optimal solution of a model as the result of changes in input data.

This chapter

Sensitivity analysis is discussed in a variety of text books. A basic treatment can be found in, for instance, [Ch83], [Ep87] and [Ko87].

References

4.1 Introduction

In a linear program, slack variables may be introduced to transform an inequality constraint into an equality constraint. When the simplex method is used to solve a linear program, it calculates an *optimal solution* (i.e. optimal values for the decision and/or slack variables), an *optimal objective function value*, and partitions the variables into *basic variables* and *nonbasic variables*. Nonbasic variables are always at one of their bounds (upper or lower), while basic variables are between their bounds. The set of basic variables is usually referred to as the *optimal basis* and the corresponding solution is referred to as the *basic solution*. Whenever one or more of the basic variables (decision and/or slack variables) happen to be at one of their bounds, the corresponding basic solution is said to be *degenerate*.

Terminology

The simplex algorithm gives extra information in addition to the optimal solution. The algorithm provides *marginal values* which give information on the variability of the optimal solution to changes in the data. The marginal values are divided into two groups:

Marginal values

- *shadow prices* which are associated with constraints and their right-hand side, and
- *reduced costs* which are associated with the decision variables and their bounds.

These are discussed in the next two sections.

In addition to marginal values, the simplex algorithm can also provide *sensitivity range* information. These ranges are defined in terms of two of the characteristics of the optimal solution, namely the optimal objective value and the optimal basis. By considering the objective function as fixed at its optimal value, sensitivity ranges can be calculated for both the decision variables and the shadow prices. Similarly, it is possible to fix the optimal basis, and to calculate the sensitivity ranges for both the coefficients in the objective function and the right-hand sides of the constraints. All four types will be illustrated in this chapter.

Sensitivity ranges

Although algorithms for integer programming also provide marginal values, the applicability of these figures is very limited, and therefore they will not be used when examining the solution of an integer program.

Linear programming only

4.2 Shadow prices

In this section all constraints are assumed to be in standard form. This means that all variable terms are on the left-hand side of the (in)equality operator and the right-hand side consists of a single constant. The following definition then applies.

Constant right-hand side

The marginal value of a constraint, referred to as its shadow price, is defined as the rate of change of the objective function from a one unit increase in its right-hand side. Therefore, a positive shadow price indicates that the objective will increase with a unit increase in the right-hand side of the constraint while a negative shadow price indicates that the objective will decrease. For a nonbinding constraint, the shadow price will be zero since its right-hand side is not constraining the optimal solution.

Definition

To improve the objective function (that is, decreases for a minimization problem and increases for a maximization problem), it is necessary to weaken a binding constraint. This is intuitive because relaxing is equivalent to enlarging the feasible region. A “ \leq ” constraint is weakened by *increasing* the right-hand side and a “ \geq ” constraint is weakened by *decreasing* the right-hand side. It therefore follows that the signs of the shadow prices for binding inequality constraints of the form “ \leq ” and “ \geq ” are opposite.

Constraint weakening

When your model includes equality constraints, such a constraint could be incorporated into the LP by converting it into two separate inequality constraints. In this case, at most one of these will have a nonzero price. As discussed above, the nature of the binding constraint can be inferred from the sign of its shadow price. For example, consider a minimization problem with a negative shadow price for an equality constraint. This indicates that the objective will decrease

Equality constraints

(i.e. improve) with an increase in the right-hand side of the equality constraint. Therefore, it is possible to conclude that it is the “ \leq ” constraint (and not the “ \geq ” constraint) that is binding since it is relaxed by increasing the right-hand side.

Table 4.1 presents the shadow prices associated with the constraints in the potato chips example from Chapter 2.

Potato chips model

process	constraint	optimal time [min]	upper bound [min]	shadow price [\$/min]
slicing	$2X_p + 4X_m \leq 345$	315	345	0.00
frying	$4X_p + 5X_m \leq 480$	480	480	0.17
packing	$4X_p + 2X_m \leq 330$	330	330	0.33

Table 4.1: Shadow prices

The objective in the potato chips model is profit maximization with less than or equal process requirement constraints. The above shadow prices can be used to estimate the effects of changing the binding constraints. Specifically, as discussed below, it is possible to deduce from the positive values of the frying and packing constraints that there will be an increase in the overall profit if *these* process times are increased.

It is important to note that the slicing inequality constraint is nonbinding at the optimal solution and hence its associated shadow price is zero. This predicts that there will be no improvement in the objective function value if the constraint is relaxed. This is expected because a sufficiently small change in the right-hand side of such a constraint has no effect on the (optimal) solution. In contrast, a change in the objective function is expected for each sufficiently small change in the right-hand side of a *binding* constraint.

Nonbinding constraint

The benefit of relaxing a binding constraint can be investigated by resolving the LP with the upper bound on the availability of the packer increased to 331 minutes. Solving gives a new profit of \$190.33, which is exactly the amount predicted by the shadow price (\$0.33). Similarly an upper bound of 332 minutes gives rise to a profit of \$190.67. This shows that the shadow price gives the revenue of an extra minute of packing time, which can then be compared with the cost of installing additional packing equipment. The shadow price can therefore be considered to represent the benefit of relaxing a constraint. From comparing the shadow prices of frying and packing, it is fair to conclude that upgrading packing equipment is probably more attractive than frying equipment.

Relaxing binding constraints

If a binding constraint is tightened, the value of the objective will deteriorate by the amount approximated by the shadow price. For the case of lowering the upper bound on the availability of the packer to 329 minutes, profit decreases by \$0.33 to \$189.67. This shows that the shadow price gives the amount of change in both directions.

Tightening constraint




Figure 4.1: Weakening the constraint on packing

In Figure 4.1, there is a graphical illustration of what happens when the packing constraint is weakened. The corresponding line shifts to the right, thus enlarging the feasible region. Consequently, the dashed line segment representing the linear objective function can also shift to the right, yielding a more profitable optimal solution. Notice that if the constraint is weakened much further, it will no longer be binding. The optimal solution will be on the corner where the frying line and the plain chips axis intersect (120,0). This demonstrates that the predictive power of shadow prices in some instances only applies to limited changes in the data.

Picturing the process

In general, the conclusions drawn by analyzing shadow prices are only true for small changes in data. In addition, they are only valid if one change is made at a time. The effects of changing more data at once cannot be predicted.

Shadow price limitations

In the two decision variable example to date, there have only been two binding constraints. However, if there were three or more constraints binding at the optimal solution, weakening one requirement may not have the effect suggested by the shadow prices. Figure 4.2 depicts this situation, where the bold lines can be interpreted as three constraints that intersect at the optimal solution. This condition is referred to as *degeneracy* and can be detected when one or more of the variables (decision and/or slack variables) in the basis are at one of their bounds. In this example, one of the three slack variables will be in the basis at their bound of zero. In the presence of degeneracy, shadow

Degeneracy

prices are no longer unique, and their interpretation is therefore ambiguous. Under these conditions, it is useful to analyse sensitivity ranges as discussed in Section 4.4.




Figure 4.2: Non-unique solutions illustrated

In general, the information provided by shadow prices should only be used as an indication of the potential for improvement to the optimal objective function value. However, there are some circumstances where a slightly stronger conclusion can be drawn. Specifically, if there are shadow prices that are large relative to others, then it is possible that the optimal solution is overly sensitive to changes in the corresponding data. Particular care should be taken when this data is not known exactly. Under these conditions, it might be wise to use methods specifically designed for handling uncertainty in data, or to run a set of experiments investigating the exact effect on the (optimal) solution with particular data modifications.

Conclusion

4.3 Reduced costs

A decision variable has a marginal value, referred to as its reduced cost, which is defined as the rate of change of the objective function for a one unit increase in the bound of this variable. If a nonbasic variable has a positive reduced cost, the objective function will increase with a one unit increase in the binding bound. The objective function will decrease if a nonbasic variable has a negative reduced cost. The reduced cost of a basic variable is zero since its bounds are nonbinding and therefore do not constrain the optimal solution.

Definition

By definition, a nonbasic variable is at one of its bounds. Moving it off the bound when the solution is optimal, is detrimental to the objective function value. A nonbasic variable will improve the objective function value when its binding bound is relaxed. Alternatively, the incentive to include it in the basis can be increased by adjusting its cost coefficient. The next two paragraphs explain how reduced cost information can be used to modify the problem to change the basis.

Improving the objective

A nonbasic variable is at either its upper or lower bound. The reduced cost gives the possible improvement in the objective function if its bound is relaxed. Relax means decreasing the bound of a variable at its lower bound or increasing the bound of a variable at its upper bound. In both cases the size of the feasible region is increased.

Bound relaxation

The objective function value is the summation of the product of each variable by its objective cost coefficient. Therefore, by adjusting the objective cost coefficient of a nonbasic variable it is possible to make it active in the optimal solution. The reduced cost represents the amount by which the cost coefficient of the variable must be *lowered*. A variable with a positive reduced cost will become active if its cost coefficient is lowered, while the cost coefficient of a variable with a negative reduced cost must be increased.

Objective coefficient

Table 4.2 gives the reduced costs associated with the optimal solution of the potato chips model. In this problem the decision variables are the quantity of both types of chips to be included in the optimal production plan. The reduced costs of both chip types are zero. This is expected since neither chip type is at a bound (upper or lower) in the optimal solution. It is possible to make one variable nonbasic (at a bound) by modifying the data.

Potato chips model

chip type	optimal value	reduced costs
	[kg]	[\$/kg]
plain	57.5	0.0
Mexican	50.0	0.0

Table 4.2: Reduced costs in the potato chips model

A modification to the potato chips model which results in a nonzero reduced cost, is to lower the net profit contribution of Mexican chips from 1.50 to 0.50 \$/kg. Solving the model gives the optimal production plan in Table 4.3, where Mexican chips production is now nonbasic and at its lower bound of zero. As a result, there is a reduced cost associated with the production of Mexican chips. The profit has dropped from \$190 to \$165, which is the best achievable with these profit contributions.

The modified potato chips model

chip type	optimal value	reduced costs
	[kg]	[\$/kg]
plain	82.5	0.0
Mexican	0.0	-0.5

Table 4.3: Reduced costs in the modified potato chips model

In Table 4.3, the zero reduced cost of plain chips production reflects that it is basic. The nonzero reduced cost for Mexican chips indicates that it is at a bound (lower). Its negative reduced cost indicates that the objective function value will decrease should the quantity of Mexican chips be increased by one unit. Given that it is a maximization problem, this is unattractive and hence is consistent with the decision to place the variable at its lower bound.

Interpretation

The optimal Mexican chips production is at its lower bound because it is more attractive to produce plain chips. However, by adjusting its objective cost coefficient it is possible for Mexican chips to become active in the optimal solution. From Table 4.3, and using the fact that the potato chips model is a maximization model, it can be concluded that if the profit contribution from Mexican chips is increased by at least 0.5 \$/kg, then Mexican chips will become basic.

Adjusting objective coefficient

Changing coefficients in the objective can be regarded as changing the slope of the objective function. In Figure 4.3, profit lines corresponding to different profit contributions from Mexican chips are given. It can easily be seen that the slope of the objective determines which corner solution is optimal. Reduced costs give the minimal change in a coefficient of the objective such that the optimal solution shifts from a corner on one of the axes to another corner of the feasible region (possibly on one or more of the other axes). Note that the slope of line (2) is parallel to the slope of the constraint on packing, thus yielding multiple optimal solutions.

Picturing the process

One might conclude that a model never needs to be solved on the computer more than once since all variations can be derived from the reduced costs and shadow prices. However, often it is useful to conduct some further sensitivity analysis. In general, shadow prices and reduced costs are only valid in a limited sense, but that they are useful when their values are large relative to others. Their exact range of validity is not known a priori and their values need not be unique. It is a result of this limitation that the study of sensitivity ranges becomes useful.

Conclusion




Figure 4.3: Varying the slope of the objective function

4.4 Sensitivity ranges with constant objective function value

Optimal decision variables and shadow prices are not always unique. In this section the range of values of optimal decision variables and optimal shadow prices for the potato chips model is examined. In ATMMS there are in-built facilities to request such range information.

This section




Figure 4.4: Decision variable ranges illustrated

Figure 4.4 illustrates the sensitivity ranges for decision variables if the three bold lines are interpreted as different objective contours. It is clear that for contour (1) there is a *range* of values for the amount of plain chips ([32.5, 57.5] kg) and, a corresponding range for the amount of Mexican chips ([50.0, 70.0] kg) that can yield the same objective value. Contour (3) also exhibits this behavior but the ranges are different. For objective contour (2), there is a unique optimal decision.

*Decision
variable ranges*

The bold lines in Figure 4.4 were initially interpreted as constraints that intersect at the optimal solution. In this case, the shadow prices are not unique and the situation is referred to as a case of degeneracy. The potato chip problem to date does not have a constraint corresponding to line (2) but a new constraint can easily be added for illustrative purposes only. This constraint limits the objective value to be less than its optimal value. Thus, the contours in Figure 4.4 can also be interpreted as follows:

*Shadow price
ranges*

1. frying constraint,
2. new constraint limiting the optimal value, and
3. packing constraint.

Examine the shadow prices for values of the bounds in a very small neighborhood about their nominal values. This helps to see that there are multiple solutions for the shadow prices. If constraint (2) in Figure 4.4 is binding with shadow price equal to 1.0 \$/min, then the shadow prices on constraints (1) and (3) will necessarily be zero. By relaxing constraint (2) a very small amount, it becomes non-binding. Its shadow price will go to zero, and as this happens, constraints (1) and (3) become binding with positive prices equal to the optimal values from Table 4.1. This means that in this special case there is a range of shadow prices for all three constraints where the optimal objective value remains constant.

*Examining their
values*

1. frying constraint has shadow price range [0.0, 0.17]
2. new constraint has shadow price range [0.0, 1.0], and
3. packing constraint has shadow price range [0.0, 0.33].

4.5 Sensitivity ranges with constant basis

The optimal basis does not always remain constant with changes in input data. In this section the ranges of values of objective function coefficients and right-hand sides of the original potato chips model are examined with the requirement that the optimal basis does not change. In AIMMS there are in-built facilities to request such range information.

This section

Changing one or more coefficients in the objective has the effect of changing the slope of the objective contours. This can be illustrated by interpreting the bold lines in Figure 4.4 as the result of

Ranges of objective coefficients

1. decreased plain chip profits (1.2 \$/kg)
2. nominal plain chip profits (2.0 \$/kg), and
3. increased plain chip profits (3.0 \$/kg),

Note that the optimal basis for the nominal profits is still optimal for the other two objectives. Therefore, the range of objective coefficient values defined by contours (1) and (3) represent the amount of plain chips for which the optimal basis remains constant. Outside this range, there would be a change in the optimal basis (movement to a different extreme point).

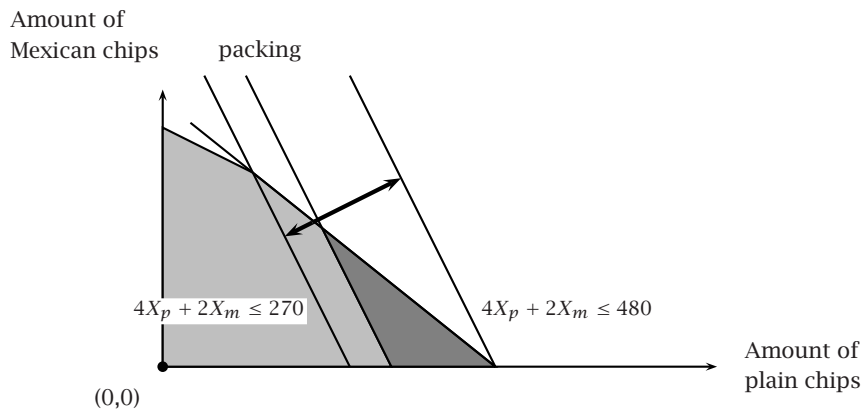


Figure 4.5: Right-hand side ranges illustrated

The potato chip model uses less than or equal constraints, but the following analysis also holds for greater than or equal constraints. The nominal solution of the potato chip problem has the packing and frying constraints binding. These binding constraints represent a basis for the shadow prices. By changing the right-hand side on the packing constraint, it will shift as can be seen in Figure 4.5.

Ranges of right-hand sides

The right-hand side can shift *up* to 480.0 minutes, where it would become redundant with the lower bound of zero on the amount of Mexican chips. The solution is then degenerate, and there are multiple shadow price solutions. This can also be interpreted as a change in the basis for the shadow prices. The right-hand side can shift *down* to 270.0 minutes, where it becomes redundant with the slicing constraint, and another change in the shadow price basis can occur. Through this exercise, it has been shown that the right-hand side on

Examining their values

the packing constraint has a range of [270.0, 480.0] minutes over which the shadow price basis does not change. Any extension of this range will force a change in the binding constraints at the optimal solution. Changing the right-hand side of non-binding constraints can make them become binding. The non-binding constraints in the potato chip problem are the slicing constraint and the two non-negativity constraints on the decision variables.

4.6 Summary

In this chapter, the concepts of marginal values and ranges have been explained using the optimal solution of the potato chips model. The use of both shadow prices and reduced costs in sensitivity analysis has been demonstrated. Sensitivity ranges have been introduced to provide validity ranges for the optimal objective function value and optimal basis. Although there is some benefit in predicting the effect of changes in data, it has been shown that these indicators do have their limits. Repeated solving of the model provides the best method of sensitivity analysis, and the AIMMS modeling system has some powerful facilities to support this type of sensitivity analysis.

Bibliography

- [Ch83] V. Chvátal, *Linear programming*, W.H. Freeman and Company, New York, 1983.
- [Ep87] D.D. Eppen, F.J. Gould, and C.P. Schmidt, *Introductory management science*, 2nd ed., Prentice-Hall, Englewood Cliffs, N.J., 1987.
- [Ko87] J.S.H. Kornbluth and G.R. Salkin, *The management of corporate financial assets: Applications of mathematical programming models*, Academic Press Ltd., London, 1987.