Interfaces
Vol. 41, No. 3, May–June 2011, pp. 278–288
ISSN 0092-2102 | EISSN 1526-551X | 11 | 4103 | 0278
DOI 10.1287/inte.1100.0542

A Mathematical Programming Approach to Applicant Selection for a Degree Program Based on Affirmative Action

Guillermo Durán
Departmento de Ingeniería Industrial, Universidad de Chile, 8370439 Santiago, Chile;
Departmento de Matemática, FCEN, Universidad de Buenos Aires, Argentina;
CONICET, Argentina, gduran@dii.uchile.cl

Rodrigo Wolf-Yadlin
Departmento de Ingeniería Industrial, Universidad de Chile, 8370439 Santiago, Chile, rwolf@dii.uchile.cl

In 2007, the Department of Industrial Engineering at the University of Chile allied with a major Chilean mining company to inaugurate a master's degree program in globalization management. The program's objective is to address the challenges Chile faces in its development of human and social capital by training young professionals. This paper describes the use of mathematical programming models in the program's applicant selection procedure for the first three entering classes, subject to equity criteria on gender, regional origin, and socioeconomic background. The models generated robust solutions in minutes, an achievement practically impossible using manual methods. The application's success demonstrates how mathematical programming and operations research can contribute to social policy.

Key words: affirmative action; equity; integer programming; operations research; robust selection.
History: This paper was refereed. Published online in *Articles in Advance* June 1, 2011.

In 2007, the University of Chile inaugurated a master's degree program in globalization management. The program is run by an alliance of the university's Department of Industrial Engineering and one of Chile's largest mining companies. Its goal is to address the challenges Chile faces in its development of human and social capital by providing young professionals with an excellent education in business administration. These students come from a wide spectrum of socioeconomic backgrounds and must have the potential to perform effectively in globalized businesses. All students who are admitted are eligible for a grant allowing them to study full time. The 18-month program includes courses given in Chile and internships at universities abroad. Applicants must meet a set of requirements regarding age, educational background, and work experience.

For the first entering class (2007), the program directors set the number of admitted students at 53; they reduced this number to 51 for the 2008 class and to 47 for the 2009 class. They also decided to apply equity (i.e., affirmative action) selection criteria based on gender, region of origin, and socioeconomic background. This policy reflects another central objective of the program, which is to ensure genuine equality of opportunity. The directors were seeking to reverse Chile's traditional concentration of resources on men from the Santiago (capital) region and in the top-income quintile. Thus, the program directors decided that at least 30 percent of total admissions would be women, 60 percent would come from non-Santiago regions, and 80 percent would belong to the lower-four income quintiles. In 2008, program organizers slightly altered these criteria by lowering these percentages to 30, 55, and 70 percent, respectively. In 2009, they reinstated the percentages to 30, 60, and 80 percent, respectively; they also added a 30 percent minimum nonengineer requirement to increase diversity, because data for the first two classes indicated that successful applicants were almost exclusively engineers.

In each year, more than 600 applicants met the minimum requirements and entered the first stage of the selection process, in which they were each assigned a number of points based on their academic and work backgrounds. Approximately 500 applicants (of the 600) advanced to the second stage, which required that they take a series of aptitude tests in various fields of knowledge and a psychometric evaluation. These results were combined with the first-stage point total to arrive at a new score. Based on that score, 160 applicants progressed to the third stage, in which they were given a psychological evaluation; in 2008 and 2009, an English test was added to the third stage. Students who passed the third stage (87 in 2007, 83 in 2008, and 86 in 2009) formed the short list from which the final group of admitted applicants would be selected; the final list also included a 20-candidate waiting list to be used if any admitted student declined the admission offer. In 2009, the psychological evaluation was used to reduce the number of applicants and to define a series of scenarios for the final-group selection instead of only the one scenario used in 2007 and 2008. Each scenario considered different ways to calculate the candidate's score and to take into account the results of the psychological evaluation.

The program organizers defined the method of evaluating the scores and minimum conditions for determining the short list, which we will not discuss in this paper. Our focus is on the use of a novel selection process using mathematical methods to identify the final group of candidates; this reflects the program organizers' desire to choose applicants with the best qualifications, while ensuring that the advantages of gender, regional origin, and socioeconomic background would not be decisive. Program officials also precisely identified the lower quintiles and non-Santiago region status (based on private school attendance or place of secondary school completion). During the 2007 selection process, two quintile definitions were used. The first defined the top quintile as applicants who had attended private secondary schools; the second added the condition that at least one of the applicant's parents was in a traditionally high-income profession (e.g., medicine, engineering, law, economics). The latter definition, which was more restrictive and therefore broadened the base of the lower quintiles, was finally selected and was maintained in 2008 and 2009.

Our objective in this paper is to show how integer linear programming (ILP) models were used to select, from the short list, the applicants who best fit the qualifications profile of the master's program while satisfying the equity constraint minima. Our goal was to obtain a definitive solution that was robust in that it would not vary greatly because of small variations in the admission criteria. Achieving this with a manual procedure in a reasonable period would have been practically impossible; this originally prompted our choice of using mathematical modeling for applicant selection. The ILP models we used also brought greater transparency to the selection process. Our study demonstrates the potential of operations research (OR) to contribute to affirmative action social policies and, more particularly, to strengthen equality of opportunity in graduate-level education. The programming tool we describe was used in each of the 2007–2009 selection processes.

Affirmative action is a controversial topic in public discussions worldwide. In particular, the OR community has addressed it in a special issue of *OR/MS Today* (Barnett 1996, Caulkins 1996, Horner 1996, Pollack 1996), which provides differing perspectives about this topic.

In the literature, the use of operations management and OR techniques in selection processes is mainly associated with applications of the analytic hierarchy process (AHP) (Saaty 1980). Examples include articles in the fields of health (Ross and Nydick 1992), education (Grandzol 2005), and business management (Chan 2003). Ross and Nydick (1992) discuss a study by a pharmaceutical company to determine which projects to implement in its search for new cancer drugs. The findings demonstrate the advantages of using AHP in the decision-making process, particularly when using various criteria, and even more so when the criteria are subjective. In Grandzol (2005), AHP is used to choose new instructors at a higher-education institution. The selection process generated wastes no resources (e.g., selection committee time), considers all criteria, and applies a procedure that is fair to all participants. Chan (2003) presents another AHP application that uses both quantitative and qualitative factors to determine the best way to select a manufacturing firm's suppliers.

The application of decision models for project selection also relates to selection processes. Greiner et al. (2003) present a hybrid approach that combines AHP and integer programming (IP) to screen weapon-systems development projects. The paper describes the use of AHP to allow a decision maker to incorporate qualitative and intangible criteria into the decision-making process. It then solves a knapsack problem using the priorities defined by AHP as the coefficients of the objective function in the optimization problem. This paper also provides interesting extensions of the application.

Gottlieb (2001) discusses an application of an IP model for college admissions. In this paper, the author proposes to divide the candidates into homogeneous groups and then select the number of candidates in each group by considering, for example, constraints related to a student's academic level, the school budget (schools must accept some students who can pay more), and the relationship between enrolled and admitted applicants (i.e., the number of applicants who enroll divided by the number of applicants who are admitted). The model uses the probability that an applicant will enroll if admitted—a probability that can be calculated using historical data.

To the best of our knowledge, the use of ILP techniques in a selection process that incorporates affirmative action policies is a novel feature of our work. We adapted the concept of combining the results of multiple scenarios to arrive at a final robust decision from Epstein et al. (2002), in which the authors apply a mathematical model to create a set of winning offers in a combinatorial auction for supplying school meals in Chile.

In the remaining sections of this paper, *Mathematical Models* describes the three mathematical models used in the selection process, *Selection Algorithm* discusses the development of a selection algorithm for combining the models to obtain a more robust solution. *Results* and *Discussion and Conclusions* follow. Two appendices (*Appendix A*: *Final Results of Selection Processes*: 2007, 2008, *and* 2009 and *Appendix B*: *Application of Selection Algorithm to* 2008 *Process*) are included in an electronic companion to this paper. The electronic companion is available as part of the online version that can be found at http://interfaces.pubs.informs.org/ecompanion.html.

The tables in Appendix A provide the final results of the three selection processes (2007, 2008, and 2009). Appendix B illustrates a specific example of the selection algorithm for the 2008 process.

Mathematical Models

Each of the three mathematical models that we developed for the selection process incorporates a different selection criterion. The first model maximizes the sum of the scores assigned to the selected applicants, the second minimizes the sum of their rankings, and the third minimizes the ranking of the last candidate selected. In each model, applicants must satisfy the gender, lower-income-quintile, and non-Santiago-region criteria and, for 2009, the professional-background criterion. Below, we list the notation, decision variables, and constraints common to all the models and then describe the specifics of each model.

Notation

Let N be the number of persons to be admitted, K the set of short-listed applicants, M the set of all female applicants, R the set of all non-Santiago-region applicants, Q the set of lower-income-quintile applicants, P the set of nonengineers, and p_i the score of applicant i. (Without loss of generality, we may assume that the scores are ordered from high to low.)

Decision Variables

$$x_i = \begin{cases} 1 & \text{if applicant } i \text{ is selected;} \\ 0 & \text{otherwise.} \end{cases}$$

Constraints

1. The total number of applicants to be selected is predetermined by the program organizers (the value of N was 53 in 2007, 51 in 2008, and 47 in 2009).

$$\sum_{i \in K} x_i = N.$$

2. At least m percent of the selected applicants must be women. The value of m used in each year was 30.

$$\sum_{i \in M} x_i \geq \frac{m}{100} \cdot N.$$

3. At least r percent of the selected applicants must be from non-Santiago regions. The values of r used were 60 in 2007 and 2009, and 55 in 2008.

$$\sum_{i\in R} x_i \geq \frac{r}{100} \cdot N.$$

4. At least q percent of the selected applicants are from the lower-four income quintiles. The values of q used were 80 in 2007 and 2009, and 70 in 2008.

$$\sum_{i\in Q} x_i \geq \frac{q}{100} \cdot N.$$

5. At least p percent of the selected applicants must not be engineers. The value of p used was 30 (this constraint was added for 2009).

$$\sum_{i\in P} x_i \geq p/100 \cdot N.$$

We now describe the objective functions of each model; for the third model, we describe an additional decision variable and constraint.

Model 1

The objective is to maximize the sum of the selected applicants' scores to find a global optimum score.

Objective function

$$\max \sum_{i\in K} x_i \cdot p_i.$$

Model 2

The concept behind model 2 is similar to that for model 1; however, in model 2 we consider the candidates' ranking order, not their scores. The objective is to minimize the sum of the selected applicants' rankings.

Objective function

$$\min \sum_{i\in K} i \cdot x_i.$$

If two applicants have the same score, we attribute a better ranking to the applicant who satisfies a greater number of the equity and professional-background characteristics that the program seeks to favor (i.e., women, lower-income quintiles, non-Santiago regions of origin, nonengineers). If the tie persists, we define the ranking randomly but record the details; if one of the tied applicants is among those admitted in the final-selection stage or placed on the waiting list and the other one is not selected, the organizers make the final decision based on qualitative criteria that they consider appropriate. These conditions to break a tie are also valid for the other two models.

Model 3

Model 3 aims to impose the condition that the last applicant selected has the best possible ranking. The objective is to minimize the ranking of the last-selected applicant. The model contains an additional decision variable y (positive real), which does not appear in the other models, whose value is greater than or equal to the ranking of all the selected applicants; once minimized, this will be the ranking of the last-chosen candidate. The model also incorporates an extra constraint that requires the new variable to be greater than or equal to the position on the (ordered) list of all selected applicants. The objective function value minimizes the sum of this variable's value and that of the objective function value of model 2, the latter multiplied by a very small number. We do this so that given two sets of candidates who are tied in the ranking of the last-chosen candidate, the set of best-ranked applicants is selected. Clearly, the second term of the sum will not affect the result if the rankings of the two sets of candidates differ from that of the last-chosen applicant.

Decision variable

y: the relative position greater than or equal to selected applicants.

Constraint

$$i \cdot x_i \leq y \quad \forall\, i.$$

Objective function

$$\min\left\{ y + \left(0.0002 \cdot \sum_{i\in K} i \cdot x_i\right)\right\}.$$

From these models, we can easily construct examples in which each identifies a different group of admitted applicants. The computational results in the *Results* section, which show that the admissions-list cutoff varies slightly between the models, confirm this. To achieve a more robust solution, we therefore use all three models instead of only one in a procedure that combines their results and uses the selection algorithm described in the next section.

Selection Algorithm

The procedure defined by the selection algorithm runs the models a set number of times before combining the best solutions generated to produce a single final solution. The number of runs is a parameter chosen by the user. We used the three best solutions of each model; run 1 yielded the best solution, run 2 the second best, and run 3 the third best. We obtained the second-best solution by adding a constraint to the models to render the best solution infeasible; we derived the third-best solution by similarly also eliminating the second-best one. If unique-best, second-best, and third-best solutions exist, we assign the applicants in each one the coefficients 1, 0.6, and 0.3, respectively. We then sum these values across all three solutions (i.e., runs) and models for each applicant. If, for example, an applicant is selected in run 1 of models 1 and 2, run 2 of models 1 and 2, and run 3 of model 1, he or she is assigned a general weighting coefficient of 3.5. Finally, we multiply this value by the applicant's point total to determine a new score.

More specifically, the algorithm's five steps are as follows.

1. *First selection*: The applicants appearing in the optimal solution (i.e., run 1) of all three models are identified. These candidates are immediately admitted to the program. If the three models return the same optimal solution, the admissions list is complete, we calculate the new scores for each applicant not selected, and the algorithm jumps to Step 5 to identify the applicants on the waiting list. Otherwise, it goes to Step 2.

2. *New score*: The coefficients described above are now calculated for each applicant not selected in Step 1 and then multiplied by their respective point totals to generate new scores.

3. *Second selection*: The composition of the admissions decided in Step 1 in terms of candidates from the three equity categories (i.e., women, non-Santiago regions, and lower-income quintiles) is evaluated to determine how many more of each category are needed to meet the required percentage minima. Model 2 is then run using the scores obtained in Step 2 with constraints that ensure, first, that it selects at least the number of candidates required to meet these minima, and second, that those selected equal the number lacking in the first selection to satisfy N, the program total. The algorithm then checks whether the optimal solution is unique. If not, the algorithm proceeds to Step 4; if it is, the applicants in the solution are selected, thus completing the admissions list, and the algorithm jumps to Step 5.

4. *Third selection:* The sum of the scores of each solution found in Step 3 is calculated (i.e., model 1 is applied). The group with the highest point total completes the list of admitted applicants. If two or more solutions provide the same score, all alternatives are presented to the program organizers for a final decision.

5. *Waiting list*: If the number of applicants who appear in any of the nine runs but are not admitted is greater than 20, the top scorers (after weighting) of this group are placed on the waiting list. If the number is less than 20, all these applicants are placed on the waiting list, and the additional applicants needed to complete it are chosen from the best scorers (before the weighting) among those who were not selected in the best solutions of any model.

The identification of the waiting-list candidates does not consider the equity criteria. However, if any admitted applicant later declines to enter the program, that applicant is replaced with the highest scorer from the waiting-list applicants in such a way that the equity-category minima are met.

The selection algorithm guarantees the robustness of the final solution in the sense that the applicants admitted to the program will have all been present in various best solutions of each model. This clearly shows the value of using mathematical programming models. Furthermore, our programming tool gives the process a high transparency level. In Appendix B, we illustrate how the algorithm actually functions by applying it to the 2008 selection process.

Results

In the selection processes for the three years, we used the models in the first two stages only to ensure enough applicants from the three equity categories. The models acted simply as a support tool for determining which candidates would advance to the next stage; we did not use the selection algorithm. The results we present in this section relate to the final stage of the three processes, the stage in which the

Model	Best solution	O.F. value (first-quintile definition)	O.F. value (second-quintile definition)	Bound for the O.F.
1	1	3,334.0798	3,392.6797	3,399.414
1	2	3,334.0717	3,392.4	
1	3	3,333.6946	3,391.9229	
2	1	1,792	1,470	1,431
2	2	1,795	1,473	
2	3	1,795	1,476	
3	1	71.3702	64.294	53.2862
3	2	71.3714	64.2946	
3	3	71.3716	64.2952	

Table 1: The table shows the objective function (O.F.) results for the three best solutions of each model using the two income-quintile definitions the program organizers used in 2007.

models are used with the selection algorithm. Note that in each of the three years, the program organizers adopted the admissions list thus determined as the definitive list.

Selection Process for 2007

We describe the selection process for 2007 below.

In Table 1, we can observe that the second-quintile definition leads to superior objective function values using all three models, because the set of applicants in quintiles other than the first is larger under this definition. Note also that an upper bound for the objective function in model 1, which would be obtained if the only constraint were the selection of the top 53 applicants to fill the program without any equity criteria restrictions, is 3,399.414.

The corresponding theoretical lower bounds for models 2 and 3 are 1,431 and 53.2862, respectively (the decimals in the latter figure are used to break a tie if two solutions select the same candidate as the last-admitted applicant). The function would have these values if the best 53 scorers were selected (i.e., no equity constraints existed). Again, the constraints have a major impact on the objective function values when we apply the first-quintile concept.

Under the second-quintile definition, the selection algorithm jumps directly from Steps 1 to 5—the admitted applicants are the same using all three models. However, if we use the first definition, the algorithm must execute Steps 2, 3, and 4 before going to Step 5, because the models' best solutions coincide on only 48 (of a possible 53) selected candidates. These cases illustrate the importance of designing a transparent process that combines the results of the different models and completes the admissions list by applying the equity constraints. This process also ensures the robustness of the final solution. In this case, for example, the five candidates selected to complete the list of 53 must appear in several of each model's best solutions.

Under the first-quintile definition, nine applicants appeared in the run solutions but were not admitted. Therefore, to complete the waiting list, 11 more candidates had to be chosen from among those who did not appear in any solution. Under the second definition, only three were present in the solutions but were not admitted, leaving 17 additional waiting-list applicants to be selected. These data indicate that under the first definition, 62 applicants appear in the nine runs; the number falls to 56 under the second definition (i.e., using the second definition, the models coincide to a high degree in their best, second-best, and third-best solutions). The second-best solutions coincide perfectly, as do the third best for models 1 and 3; the latter solutions differ from model 2 on only one candidate. The program organizers ultimately opted for the second-quintile definition (using it again in 2008 and 2009) to improve the academic quality of the set of chosen candidates. In the rest of this paper, we will therefore exclusively apply the second-quintile definition.

Selection Process for 2008

Table 2 shows the objective function results for the three best solutions of each model in the 2008 selection process.

Table 2 shows that the best possible value for model 1 is 3,325.4, the sum of the scores for the top 51 candidates (the number admitted in 2008) with no other constraints applied. For model 2, with the reduction of admissions from 53 to 51, the objective function value, assuming no other constraints, is 1,326; for model 3, it is 51.2662. Thus, in 2008, the optima for the three models are closer to their ideal values than in 2007. We can attribute this to the less restrictive socioeconomic-level and regional-origin constraints in 2008 than in 2007.

When the selection algorithm was run, the best solutions of the three models coincided on 49 (of a

Model	Best solution	O.F. value	Bound for the O.F.
1	1	3,322.65	3,325.4
1	2	3,322.6	
1	3	3,322.5	
2	1	1,351	1,326
2	2	1,353	
2	3	1,353	
3	1	55.2726	51.2662
3	2	55.273	
3	3	55.2734	

Table 2: The table shows the objective function values for the three best solutions of each model in the 2008 selection process.

possible 51) applicants in Step 1. Therefore, the algorithm had to execute Step 3 before jumping to Step 5 (see Appendix B).

Note that because four of the waiting list applicants were present in a run solution, the total number present in any of the nine runs was 55. Note also that model 2 generates two second-best solutions, and both models 1 and 2 yield the same optimal solution, which differs from model 3 in two applicants.

Selection Process for 2009

For the 2009 process, the program organizers implemented two changes to broaden the criteria used for deriving candidates' scores. The first was to define another pair of weighting factors that placed slightly more importance on work history and slightly less on academic background. Therefore, in 2009 each applicant had two initial point totals, one based on the newly defined weights and the other on the original weights used in 2007 and 2008.

The second change was that applicants who passed the psychological test, the last stage in defining the short list, were grouped into three aptitude categories: more qualified (I), qualified (II), and less qualified (III). Their point totals were then multiplied by a coefficient corresponding to the category to which they were assigned. Three sets of such coefficients were defined: 1, 0.95, 0.9; 1, 0.97, 0.92; and 1, 0.92, 0.85.

The combined effect of the changes resulted in the creation of six final-score totals (i.e., scenarios). Scenario 1 combined the 2007 and 2008 point-total weights with the first set of coefficients, scenario 2 combined them with the second set, and scenario 3 with the third set. Scenarios 4, 5, and 6 were formed analogously by combining the 2009 point-total weights with the first, second, and third coefficient sets, respectively.

For each of the six scenarios, the same selection procedure used in 2007 and 2008 (with the application of the selection algorithm described in the *Selection Algorithm* section) was applied to determine which of the short-listed applicants would be admitted. The applicants appearing in the final solution of all six scenarios were identified and immediately admitted to the program. The admissions list was completed by maximizing the number of scenarios in which an applicant was selected subject to the equity and nonengineer constraints. Candidates chosen in any scenario who were not in the final group were placed on the waiting list, which was filled with the 20 top scorers among those who had not been selected in any scenario. For this purpose, the point totals used were the averages over the six scenarios.

Table 3 shows the objective function results for the best solutions of each model and each scenario in the 2009 selection process. Note that the results of the second-best and third-best solutions of each model and each scenario are omitted in this table.

Note that because the applicants' scores vary in the scenarios, the upper bound for model 1 also varies by scenario. However, in models 2 and 3 the scores do

Scenario	Model	O.F. value	Bound for the O.F.
S1	1	2,846.306	2,930.316
S1	2	1,556	1,128
S1	3	75.3276	47.2256
S2	1	2,880.602	2,964.523
S2	2	1,559	1,128
S2	3	74.3256	47.2256
S3	1	2,793.659	2,875.891
S3	2	1,512	1,128
S3	3	75.3090	47.2256
S4	1	2,831.390	2,927.985
S4	2	1,610	1,128
S4	3	74.3262	47.2256
S5	1	2,865.769	2,961.709
S5	2	1,617	1,128
S5	3	74.3270	47.2256
S6	1	2,778.286	2,872.713
S6	2	1,569	1,128
S6	3	74.3182	47.2256

Table 3: The table shows the objective function values for each model and each scenario in the 2009 selection process.

not affect the objective function result; therefore, their lower bounds do not vary by scenario—the respective values are 1,128 and 47.2256.

The data in the table reveal that the differences between the ideal theoretical values and the actual values are greater than for the previous years. This implies that the positive discrimination and nonengineer constraints more strongly impact the results. In model 1, the divergence is widest for scenario 4, whereas for models 2 and 3, the gap is greatest in scenarios 5 and 1, respectively.

Under all six scenarios, all the steps of the selection algorithm must be executed to complete the admissions list of 47 candidates. The highest number of selected applicants for which the results of all three models coincided was 44 (of a possible 47), observed in scenarios 3 and 6; the lowest number was 41, which scenario 2 generated.

To select the 47 admitted applicants, we checked which applicants were repeated in all six final-scenario solutions, and found that 43 were in all the solutions; the remaining four candidates were those selected in the greatest number of scenarios subject to the organizers' constraints. This procedure led to a tie between three groups of four. In each, one applicant appeared in five final-scenario solutions, one was in four solutions, and two were in three solutions. The group finally chosen was the one with the highest point total (for each applicant, the total points used were the averages over the six scenarios) weighted by the six scenarios.

Execution time did not exceed five seconds for any of the nine model runs. The entire procedure for each of 2007 and 2008 was completed in about 20 minutes; for 2009, it required slightly more than 90 minutes because of the multiple scenarios involved. The model solutions were generated using CPLEX 10.0 on a 2.0 GHz Pentium IV processor with 1 GB of RAM.

Discussion and Conclusions

In the first part of this section, we present various a posteriori analyses to determine the effects on the results of the various equity-criteria constraints. Lysette Henríquez, the master's program's executive director during the two first-selection processes, explained the significance of this step:

> A key aspect of the model's application is the analyses conducted during the decision processes. Visualizing a solution given a set of constraints is fundamental and practically impossible to do manually, but perhaps even more important is being able to vary the program parameters within a reasonable margin or make minor modifications to the way of calculating the scores of the applicants to examine other interesting elements of the program. A key factor is the ability to appreciate how robust is the presence of certain applicants in the solution, that is, whether or not they appear systematically in the final solution. Having this information allows the decision makers to feel more certain they are making the right admission choices. (L. Henríquez, pers. comm.)

The first of these a posteriori analyses investigates how many of those admitted to the program would not have been admitted without the application of the equity criteria. The results show that 4 of the 53 admitted candidates in 2007 (7.5 percent), 2 of the 51 in 2008 (3.9 percent), and 14 of the 47 in 2009 (29.8 percent) would not have been accepted without this positive discrimination. The decline in the 2008 number reflects that the percentage minima for each equity category were reduced slightly for the non-Santiago region and income-quintile criteria. Although the percentage changes relative to the admissions based solely on ranking are small in 2007 and 2008, the process involves decisions that impact the applicants' personal and professional futures, making it imperative that the criteria we adopt are backed by a fair and transparent mechanism.

The positive-discrimination criteria applied in 2009 strongly impacted the admissions list, demonstrating the need for greater use of mathematical models to ensure a fair admissions process. We stress that this goes further than the simple fact that six scenarios were run for 2009, instead of only one. If, for example, we had used only one scenario in 2009 (i.e., scenario 1), the number of applicants rejected because of the constraints would also have been 14 (with similar results for the other five scenarios). This may be because of the addition of the nonengineer constraint; however, it could also be attributed to the random differences in the qualifications of the applicants in each year.

An analysis of the results for the 2007 process shows that the admitted-applicant numbers exactly fulfilled the minima required by the female and non-Santiago-region equity criteria, but not the minima for the lower-income quintiles. Analyzing the 2008 results, we find that the admitted applicants exactly equaled the non-Santiago region and lower-income quintile minima, whereas the female-category minimum was exceeded by one. Similarly, in 2009, admission numbers equaled the required minima in all categories except women, where the number was one higher than the lower bound.

In the 2007 process, if the female admissions minimum is eliminated, one fewer woman and one more man would be selected. Without a non-Santiago-region minimum, two more Santiago candidates would be admitted. This indicates that removing one of the equity constraints while maintaining the others has no major effect on the final solution.

Turning to the 2008 process, if the regional-origin constraint is eliminated, the algorithm terminates upon completing Step 1 (implying that the three models give the same best solution), after selecting the exact minima for female and lower-income-quintile admissions and one fewer non-Santiago applicant than the required minimum. If the income-quintile constraint (the other active restriction in 2008) is excluded, the algorithm again terminates when Step 1 has been executed after selecting the exact minimum numbers of female and non-Santiago candidates. The percentage of lower-income-quintile candidates is 64.7 percent (three fewer lower-income-quintile applicants than the minimum required when this constraint is included).

Reviewing the 2009 process, if we run the six scenarios without the minimum nonengineer restriction, our final result has nonengineers as only 23.4 percent of the selected group, as opposed to the 30 percent the constraint imposes. If we eliminate the gender constraint in all scenarios, upon combining the solutions, our resulting list of selected applicants also contains only 23.4 percent female instead of 30 percent if we apply the constraint. Removal of the income-quintile constraint results in only 57.44 percent of admitted candidates from the lower quintiles. Finally, if we remove the region-of-origin constraint, the final solution has 55.32 percent (of applicants) from the non-Santiago region, relatively close to the constraint requirement of 60 percent.

The impact of the psychological evaluation on candidate selection interested the program organizers. This test eliminated 42 percent of the applicants who had progressed to the third stage in 2007. The number for 2008 was considerably lower at 19.41 percent, although the test was applied after the English language test, which was not present the previous year. In 2009, the psychological evaluation, which was also administered after the English test, eliminated 19.62 percent of the candidates. Therefore, it had a major impact on applicant selection; had the evaluation not been given, 13 successful candidates (24.52 percent of the total) in 2007, 11 (21.56 percent) in 2008, and 12 (25.53 percent) in 2009 would have been denied admission in favor of others who failed it.

Another general observation is that if we exclude the psychological evaluation, the model 1 results improve, whereas the other models' results either improve or deteriorate because the test reduces the number of applicants. Therefore, model 1 could not produce a better solution than the one it generated without the test because with the same point totals, the list of applicants is a subset of the original list. However, for models 2 and 3, a shortened list does not a priori affect the sum of the applicants' rankings or the ranking of the last-admitted applicant, because the relative positions of those who remain on the list may change in the new scenario. Hence, the impact on the objective function value in both models 2 and 3 depends entirely on which applicants the test eliminates.

According to Ms. Henríquez,

> These *a posteriori* analyses reveal the consequences of applying certain restrictions and enable us to make program policy decisions with full awareness of their impacts. In short, the contribution of the model has been fundamental to ensuring transparency of decisions involving the award of a grant of some USD 75,000 per student for a program that received more than 800 applications from around the country in 2008 and required an investment for the first three years of close to USD 12 million. This is particularly significant considering that the program's purpose is to stimulate the creation of a meritocracy. (L. Henríquez, personal communication)

The 2009 admission process incorporated different scenarios to increase the robustness of the final decision. As Patricio Meller, the program's academic director since its inception, explained,

> having alternative models and scenarios has been fundamental to our ability to select more suitable candidates. No model is perfect, so with various scenarios

> the risk of rejecting a good applicant is lower. That's the real advantage of selecting those who appear in the final solution in the majority of scenarios. (P. Meller, personal communication)

As a general conclusion of the study, we emphasize that it demonstrates how OR and mathematical programming can contribute to social policy issues, and in particular the usefulness of these techniques in identifying the applicants who best fit the desired profile in terms of equity criteria based on regional origin, socioeconomic background, and gender.

It is still too early to conduct a full analysis of the program's impact. Meller offers this preliminary verdict:

> The first graduating class has already completed the program and the 53 students all did extremely well. For students coming from lower-income families or regions distant from Santiago, the program can significantly change their life paths. This is the sort of impact we hope to achieve with positive discrimination. On this point we're convinced that the mathematical models we applied enabled us to choose the most appropriate applicants among those who met the constraints imposed. We cannot imagine decision making for future admission processes without the support of these tools. (P. Meller, pers. comm.)

We conclude by again emphasizing that finding robust solutions to the admissions problem in a matter of minutes using manual techniques would have been impossible. The mathematical tools developed for this task also had the added advantage of bringing transparency to the selection process.

Electronic Companion

An electronic companion to this paper, which includes Appendices A and B, is available as part of the online version that can be found at http://interfaces.pubs.informs.org/ecompanion.html.

Acknowledgments

We express our gratitude to Lysette Henríquez, Patricio Meller, Gastón Held, and the executive and academic personnel in the master's program of the University of Chile, with whom it is a joy to work; to Andrés Weintraub and Kenneth Rivkin for their helpful comments, which helped to improve this study; to Rafael Epstein, who suggested to the authorities of the master's program to apply mathematical models in the selection process; and to the anonymous referees for their valuable suggestions that helped them improve the final version of this work. The first author was partially supported by FONDECYT project number 1080286 (Chile), the Millenium Complex Engineering Systems Institute (Chile), ANPCyT grant PICT-2007-00518 (Argentina), and UBACyT grant X069 (Argentina).

References

Barnett, A. 1996. Building equal opportunity on firmer footing: A case against affirmative action. *OR/MS Today* **23**(4). Accessed April 12, 2010, http://www.lionhrtpub.com/orms/orms-8-96/building.html.

Caulkins, J. 1996. Color-blind policies are not enough: A case in favor of affirmative action. *OR/MS Today* **23**(4). Accessed April 12, 2010, http://www.lionhrtpub.com/orms/orms-8-96/color-blind.html.

Chan, F. 2003. Interactive selection model for supplier selection process: An analytical hierarchy process approach. *Internat. J. Production Res.* **41**(15) 3549–3579.

Epstein, R., L. Henríquez, J. Catalán, G. Y. Weintraub, C. Martínez. 2002. A combinational auction improves school meals in Chile. *Interfaces* **32**(6) 1–14.

Gottlieb, E. 2001. Using integer programming to guide college admissions decisions: A preliminary report. *J. Comput. Sci. Colleges* **17**(2) 271–279.

Grandzol, J. 2005. Improving the faculty selection process in higher education: A case for the analytic hierarchy process. *IR Appl.* **6** 1–13.

Greiner, M. A., J. W. Fowler, D. L. Shunk, W. M. Carlyl, R. T. McNutt. 2003. A hybrid approach using the analytic hierarchy process and integer programming to screen weapon systems projects. *IEEE Trans. Engrg. Management* **50**(2) 192–203.

Horner, P. 1996. Operations researchers cast their analytical eyes on an emotional issue. *OR/MS Today* **23**(4). Accessed April 12, 2010, http://www.lionhrtpub.com/orms/orms-8-96/debatable.html.

Pollack, H. 1996. A view from the ivory tower: Inefficient, unfair affirmative action policies must be reformed, but the costs pale in comparison to the waste of human potential among the disadvantaged. *OR/MS Today* **23**(4). Accessed April 12, 2010, http://www.lionhrtpub.com/orms/orms-8-96/view.html.

Ross, M., R. Nydick. 1992. Selection of licensing candidates in the pharmaceutical industry: An application of the analytic hierarchy process. *J. Health Care Marketing* **12**(2) 60–65.

Saaty, T. 1980. *The Analytic Hierarchy Process*. McGraw-Hill, New York.

Patricio Meller, Academic Director, Master's Program in Globalization Management, University of Chile, Santiago, writes: "I am writing as the academic director of the Master's Program in Globalization Management at the University of Chile, regarding our experience with mathematical models for the selection of applicants based on equity criteria that were developed at our request by the University's Center for Operations Management (CGO).

"The program in question was launched in 2007 by the Faculty of Physical and Mathematical Sciences in alliance with one of Chile's major mining companies. Its goal is to address the challenges currently facing the country in the development of human and social capital through the training of young professionals. The length of the program is 18 months, during which time the successful applicants complete a number of courses in Chile before taking up internships at universities abroad (Australia, UK, and the U.S.A.) To be admitted, applicants must meet a series of requirements relating to age, work experience and educational background.

"One of the program's central objectives is to begin the process of reversing the long-standing centralization of Chile's highly trained human resources among men from the Santiago (capital) region of the country in the top income quintile. To accomplish this and also help promote greater equality of opportunities, a decision was made to apply "positive discrimination" criteria to the admissions process based on gender, socioeconomic background and region of origin.

"The dilemma we faced in choosing the first entering class was how to select the applicants who best fit the qualifications profile of the program (high academic background) while also satisfying the equity criteria minima we set for female, lower-income quintile and non-Santiago region candidates.

"We turned to a group of operations research specialists at the CGO in the hope they could design robust mathematical programming models that were able to rapidly and efficiently execute a selection process incorporating these conditions. Based on our experience with the first two entering classes, the models they designed for the task were extremely successful. In particular, the solutions generated were robust in the sense that they did not vary excessively with small changes in the criteria. Performing the equivalent process manually in a short period of time would have been simply impossible. Furthermore, the use of the mathematical tool made the process more transparent.

"Finally, as a more general conclusion on the CGO's models I would emphasize that in defining the list of selected candidates in accordance with the above-mentioned equity criteria, the techniques of operations research and mathematical programming have also made an important contribution to social policy."