


3

Policy



Professor Karma Howell not only had a unique first name, she had a smile warm enough to melt steel. Her students loved her, and so did Ralph, in a proper brother-in-lawish sort of way.

Ralph and Betty sat in Karma's cramped office listening to her describe the research on decision making. They had told her that they wanted to clone their best salespeople. She thought that this was a novel way of putting it, but that the technology existed to do it—well, not genetic cloning, of course, but identifying the salient attributes of good salespeople and formulating a policy for using those attributes to guide decisions about hiring or training salespeople. She explained that the technology had its weaknesses, but that it could at least introduce order and reason to the selection process. Whatever its faults, the technology probably would lead to better overall results than the company's present methods. By the time they left, Ralph and Betty knew the next step in solving their sales problem.



The purpose of this chapter is to examine the use of policy in decision making. This follows directly from our previous discussion of framing.

32 THE PSYCHOLOGY OF DECISION MAKING

When a situation is framed, the decision maker can make decisions about it in one of three ways:

1. Recognition. The situation is so similar to one that he or she has encountered before that behavior that worked before can be used again—or, at least, a variation of what worked before can be used.
2. Inference. The situation is familiar enough that the decision maker can make an educated guess about what to do.
3. Choice. The situation is sufficiently unique that neither recognition nor inference provides adequate guidance, and the decision maker must explore his or her options and choose the most promising option.

Recognition and inference both involve policy, the subject of both this chapter and the one that follows.

❖ RECOGNITION

When a decision situation is encountered, the decision maker uses its salient features to probe his or her memory. If the probe locates a contextual memory that has features that are virtually the same as those of the current situation, the latter is said to be recognized. The advantage of recognition is that the decision maker can draw upon his or her knowledge about the previously encountered situation to guide behavior in this situation. In decision making, old behaviors that are used in new but similar situations are called *policies*. In the psychology of learning, they are called *habits*. In social psychology they are called *scripts*. In all cases the point is to add efficiency to the process by using precedent rather than engaging in choice for every situation.

One of the most thorough explorations of policy for decision making was done by Gary Klein (1993) and his associates. Their studies focused on crucial, real-life decision tasks such as firefighting, Army tank platoon command, and design engineering. The key is that the decision makers are highly familiar with the situations in question and that they have extensive training and experience in dealing with those situations. Upon encountering a situation of the same type, they draw

upon their training and experience to act appropriately. Klein calls this "recognition-primed decision making," and his description of it is called the RPD model.

There are three levels to the RPD model. The most basic level is the simple match, in which the situation is recognized and what has been done in the past is done again. The second level involves more evaluation; the decision maker performs mental simulations of variations on past behaviors and what might happen if one or another variation were used here. The third level applies when there are apparent flaws in the variations, requiring major modifications and more complex mental simulations. That is, recognition primes the decision about what to do, but it does not wholly determine what to do.

The mental simulations are an important key to recognition-primed decision making. Klein has gathered descriptions of what decision makers think about when making decisions and shows that they imagine "what might happen if they did this," and "what might be the result of doing that," where both "this" and "that" are drawn from past experience. Past behaviors have not all been successful, and past failures inform current decision making as much as, or more than, past successes. Thus when they think about whether some act might successfully deal with a situation, decision makers can use past failures to help them see flaws in the action under consideration, and to guide modifications (or rejection) of the action. The point is recognition of situations does not necessarily result in blind application of a policy based on past experience; it often prompts consideration of modifications and revisions while providing the general strategy for approaching the decision.

❖ INFERENCE

The RPD model strikes a responsive chord in most people who hear about it. The lowest of the three levels is closely related to learning and memory research in psychology. The middle level allows for minor modification of learned behavior, introducing a cognitive flavor and accounting for the fact that behavior seldom is exactly the same from one time to another, just as situations are never identical. However, the RPD model's account of the highest of the three levels is too vague to be of much use either in studying decision behavior or in helping decision makers do their jobs. For the necessary specificity, we must turn to an


older model, the Lens Model (Brunswik, 1947). While the Lens Model does not directly address all of the issues raised by the RPD model's third level, it provides a more systematic treatment of the processes involved when it is not wholly clear what to do in a familiar situation.

It all began with Egon Brunswik, an Austrian psychologist who came to the United States just before the outbreak of World War II and who specialized in the study of perception. Specifically, he studied how people use perceptual cues to make decisions (inferences) about the state of their surrounding environment. Let us take a moment to explore the underlying logic of Brunswik's position because that logic is the basis of all that follows in this chapter.

It is widely accepted by perceptual psychologists that one's mind does not directly experience the objects and events in the world around one. Instead, sensory information permits the mind to construct mental representations of those objects and events (e.g., Prinzmetal, 1995; Yantis, 1995). There are many ways to demonstrate that what seems to be direct experience of the external world actually is a mental representation built from sensory information, but let us concentrate on just one:

Consider a table that has a rectangular top. When you look at the top what do you see—a square, a rectangle, a trapezoid? In fact, the image of the table's rectangular top forms a trapezoid on the retina of your eye, but you perceive that table top to be rectangular. That is, in order to make the table top look realistic, an artist would have to draw it as a trapezoid; the nearer edge of the table would be longer than the farther edge and the two sides would form acute angles with the nearer edge and obtuse angles with the farther edge. When you looked at the artist's drawing you would see a table with a rectangular top.

Look around you for a table. Then look at it analytically, as an artist would, and you will see that its rectangular top is in fact presented to your eye as a trapezoid. Moreover, if you walk around the table you will note that the trapezoid changes shape as you move. But your perceptual system is never tricked into thinking that the top changes shape as you walk around it. This is called "perceptual constancy," and it has been studied for a hundred years (James, 1983). For our purposes, it serves as evidence that what you experience (a solid rectangular table top that does not change as you move around) and what is presented to your senses (a trapezoid that changes shape as you move) are not identical. In fact, your perceptual system must use the dynamic sensory information about the table in order to make inferences that tell you there is a

Figure 3.1 Brunswik's Lens Model

stable object out there—otherwise the contents of your experience would be in constant flux.

Brunswik (1947) summarized his idea about perceptual inference in the form of the Lens Model, diagrammed in Figure 3.1. To make it simple, imagine that your perceptual system is called upon to identify some target object (Y_e). Light and air surround the target and convey information to you through your vision, hearing, and other senses. These are cues (C_1 , C_2 , C_3) that convey information about the target. On the basis of these cues you must make a decision (Y_s) about the nature of the target.

Suppose the target's outline on your retina indicates that it is small. Your color receptors sense that it is white. And because of the way in which its surface reflects light you infer that it is covered with something soft, perhaps fur. You smell alfalfa on its breath, you can feel the radiance of its body heat, and your hearing tells you that it is nearly silent when moves. You see that the pointy part of the target (its nose?) is constantly twitching, and even more telling, there are two long white appendages that project upward from just above and slightly behind the pointy part (ears?). And the round things between the pointy part and the (assumed) ears are pink (eyes?). This must be a rabbit!

Actually, all of this happens virtually instantaneously and with little effort on your part—your perceptual system is designed to make inferences and it does so without much help from your conscious mind. This is especially true if the situation is framed so that there are a limited number of possibilities to begin with. If we had told you that we were

36 THE PSYCHOLOGY OF DECISION MAKING

at a country fair it would have limited the possibilities so much that the system would not have had to work very hard. However, if we told you that we were in a pet store, you might have expected the target to be a dog or cat, which would have required your perceptual system to work a little harder to avoid error.

Identifying targets is but one function of the perceptual system. It also serves to make comparisons among targets. Suppose we framed the situation as a county fair and told you that your job was to judge the rabbits. Now there are lots of rabbits and your perceptual system is called upon to make comparisons among them. This requires some additional effort, and your conscious mind now will be called upon to augment your automatic perceptual processes. Now you must be more discriminating; the cues will again be size, coloring, texture of the fur, and so on, together with whatever else livestock judges use when judging rabbits. This time, instead of merely inferring that the target is a rabbit (which you now take as given), you must use the cues consistently (a *policy*) when deciding about each of the rabbits so you can tell which are winners and which are not. In short, you must evaluate each rabbit's location on an underlying scale, where the best rabbits are on the high end and the worst are on the low end. Then you must select the winner based upon the rabbits' locations on this scale—presumably you would select the one that was highest on the scale.

❖ APPLYING THE LENS MODEL

Enough about rabbits. The idea is that Brunswik's (1947) apparently simple description of how the perceptual system identifies targets has hidden within it a great deal of sophistication. It applies both to basic perceptual decisions and to higher-order decisions that permit ordering of the targets on an underlying scale, usually related to preference. And it is here that we return to Ralph and Betty and their cloning problem.

When Betty and Ralph talk about cloning their good salespeople, they mean that they would like to know how to look at a job applicant (target) and make a good guess (decision) about whether he or she will turn out to be a successful salesperson. The first step is to turn the inference process around and find out what makes successful salespeople successful. The second step is to use this knowledge to devise a policy

for using information about potential job applicants in order to decide how well they are going to do (which is not unlike judging which rabbits are at the winning end of the scale). The third step is getting the HR folks to use this policy to select new salespeople or to design a program for retraining failing salespeople. Let us look at each of these steps in turn.

The Left Side of the Lens

To simplify things, let us assume that you have been hired to do the work for Betty's company. Let us assume further that the company has a pretty good performance evaluation scheme in place. The first step in discovering what makes a good salesperson (Y_1) would be to sit down and read the performance evaluations of both the good performers and the poor performers, concentrating on what differentiates them. What you learn should be augmented by discussions with supervisors and with the salespeople themselves. Eventually you should get a fair idea of some of the differences between the two groups.

Unfortunately, much of what you learn will not be very valuable. That is, much of what you learn will be about how good salespeople behave in the selling situation, but you seldom have direct information about selling behavior when you are judging the applicant's suitability for hiring. This is the crux of the problem: You usually have to make your evaluation on the basis of only tangential information—and this could have been one reason the HR people told Betty that hiring was hardly better than a crap shoot.

On the other hand, things may be better than they appear. For one thing, you know who the good and bad salespeople are, and you have their original job applications somewhere in their personnel folders. We can think of the information on job applications as potential cues for predicting whether an applicant will turn out to be successful or unsuccessful. Except that now you can turn things around—by knowing who the good salespeople are, who the so-so salespeople are, and who the poor salespeople are, you can look back at the cues on their applications and see which cues would have been the best predictors of later performance.

The procedure is fairly straightforward. First you have supervisors rate each salesperson for performance, perhaps on a scale from 1 (poor) to 7 (excellent). Then you go to each person's original job

application and code the answer to each question on the application so that it can be entered into a computer (i.e., you must convert qualitative information to quantitative information, coding yes/no answers as 1 or 0, and assigning numbers to different levels of more complicated answers.)

For each salesperson, you now have a set of data consisting of his or her performance rating and one coded answer for each of the n questions on the application. Then across all salespeople you do a multiple regression analysis using the performance ratings as the dependent variable and the coded answers to each of the n questions on the application as the n independent variables. This analysis will yield an equation that describes the structure of the relationship between the cues (the independent variables) and performance (the dependent variable) for the sales force as a whole:

$$Y_e = a + B_1C_1 + B_2C_2 + B_3C_3 + \cdots + B_nC_n. \quad (\text{eq. 1})$$

This regression equation is a statement of the optimal policy for using the cues to predict performance. Y_e stands for the supervisors' ratings of performance, C stands for a coded answer to a question on the application (questions 1, 2, 3 and so on up to n questions), and B stands for the standardized regression weight for each question, called the *beta weight*. The a is merely a scaling constant that has no particular importance for us at the moment. A cue (question) that has a large beta weight can be regarded as contributing more information to predicting performance than a cue with a small weight.

In addition, the analysis yields an indication of how well this equation fits the data, expressed as the square of the multiple correlation coefficient, R^2 . This coefficient can be interpreted as the degree to which the equation is capable of predicting the performance ratings. Of course, because the equation is derived from data across all of the salespeople, because the ratings are done by different supervisors, and because there is bound to be noise (error) in the data, the equation can never perfectly predict the performance ratings for every one of the salespeople. But if R^2 is large, it means that the cues from the application permit reasonably accurate prediction across the group as a whole. If it is small, you must assume that the cues are not very diagnostic—you have the wrong cues, they are not measured accurately, or it simply is not possible to predict performance in this situation.

The Right Side of the Lens

The use of multiple regression in the way just described is a pretty standard use. Moreover, it does not have a lot to do with human decision making. So, turning to the right side of the lens in Figure 3.1, let us consider the role of the HR people in the hiring process.

As noted above, the purpose of an application is to provide information for an HR employment officer (let us call him Hank) to make a decision about how well an applicant will perform. Use of equation 1 might appear to eliminate Hank from the process; a clerk could merely obtain the required information and put it into the equation. However attractive (and cheap) this substitution might appear to Ralph and Betty, it is unlikely to work. The equation is mechanical and inflexible. It treats every applicant the same, acknowledging no mitigating circumstances and making no exceptions. As such, its blind application probably is an invitation to a lawsuit.

Human intellect clearly has its limitations, but it excels in its ability to detect exceptions to rules (anomalies). Indeed, it tends to go to sleep when things are constant, but when something unique comes along it is quick to seize on it. Therefore, the intellect of the people in HR (Hank in particular) should be used, not subjugated to the inflexible application of equation 1. You need the HR people to make sure the equation does not blindly commit injustices or do stupid things. For example, if you have no cue in the equation for some kind of very diagnostic information, perhaps because it has never arisen before, you could overlook the best applicant you have ever had. Let us say that the applicant has owned his own highly successful company in the same field as Betty's company, but his physician told him that the stress of running it was going to kill him. He decides therefore to sell his company and to go to work as a salesperson for a similar company—he knows the business inside and out, he was his own best salesperson, and he wants to work for Betty. Because this has never happened before, there is no way to put this information in equation 1, so somebody has to override the system and hire him before he goes to a competitor.

Even for more mundane cases, Hank has to be active in the hiring process—laws must be observed and information has to be interpreted. Equation 1 is merely a policy guide. So, the question becomes one of helping Hank align his hiring policies with the policy described by equation 1, while retaining his discretion in unique cases.

The first step is to find out what Hank is doing now—what policy characterizes his use of information about the applicant to evaluate acceptability. This involves returning to the Lens Model. Your evaluation of how the cues could be used to predict good sales performance involved the left side of the lens in Figure 3.1. Your evaluation of how the cues are used by Hank involves the right side. In the past, Hank read the applications, talked with the applicants, and made a decision about how acceptable the person was, that is, placed him or her on a scale of preference for hiring.

The question of interest is about Hank's policy for using the cues on the application to make his decisions. You can find this out by performing an analysis very like the one you did before, except that this time the resulting equation will describe Hank's policy (equation 2) rather than the optimal policy (equation 1). First, you present Hank with the applications for each of the people on the sales force (with identifying information blanked out). Do not tell him that these are people who have already been hired so he will treat them as new applicants and rate their acceptability (on a scale from 1 to 7). For each salesperson, you now have the coded answers on their original application (which you already coded for the earlier analysis) and Hank's acceptability rating. Once again you do a multiple regression analysis, this time using Hank's rating as the dependent variable (Y_s) and the same coded cues (C) you used before. The analyses will yield an equation of the following form:

$$Y_s = a + B_1C_1 + B_2C_2 + B_3C_3 + \cdots + B_nC_n. \quad (\text{eq. 2})$$

Note that the form is the same as equation 1, but some of the components will be different. The things that are different are Y_s , which is Hank's acceptability rating; a , which again is an uninteresting constant; and the beta weight, B , for each cue, which indicates the relative influence (importance) of the cue to Hank's ratings. Notice that the coded cues, C , in equation 2 are the same as in equation 1. That is, you used the same information from the applications in the analysis for equation 2 as you used for equation 1. This means that if Hank's policy for using the cues (equation 2) is different from the optimal policy (equation 1), it will show up as differences between the beta weights in the two equations.

In addition, having Hank's acceptability ratings allows you to see how good a decision maker he is; you merely correlate his ratings with

the ratings the supervisors made for the salespeople. This correlation is called the *achievement coefficient*, and it indicates how well Hank's judgments of acceptability (which presumably reflects his prediction of the applicant's performance) correspond to the supervisors' performance ratings. However, a low achievement coefficient is not necessarily an indictment of Hank. Let us look more closely at this, and then let us look at how we might help Hank do a better job if he is not doing too well.

Consider Hank's dilemma. He has the answers to the questions on the application, which may not be very good predictors of performance in the first place. So, if he does poorly, it could be for either of two reasons: (1) nobody could do a good job with the information on the application, or (2) another person might do a good job but Hank cannot.

The difficulty of Hank's task is revealed by the R^2 for equation 1. If it is low, the cues are not good predictors of rated performance, and we should not expect Hank to do very well either. The low R^2 puts an upper limit on Hank's achievement coefficient (actually on the square of his achievement coefficient), so he will be unable to do better than that no matter how hard he tries.

On the other hand, if the R^2 for equation 1 is high but Hank's achievement coefficient is low, there is only one place the problem can lie—he must be weighting the cues in a way that lowers his ability to predict accurately. As we observed above, this will be revealed by differences between the beta weights in his equation and the beta weights in equation 1. The remedy is to train Hank to weight the information in the appropriate manner.

If Hank is poor at predicting performance, it does not necessarily mean that the R^2 will be low for equation 2. This is because the R^2 for equation 2 indicates how well equation 2 is able to account for *his* ratings on the seven-point scale, not how accurate these ratings are (which would be achievement). He may be using the cues all wrong, but as long as he uses them in a consistent manner the equation will be able to account for his ratings and the R^2 will be high. To the degree that he is inconsistent in his use of the cues, it becomes more difficult to account for his ratings and, therefore, the R^2 is reduced.

Making a Recommendation

Now you have equations for both the left and right sides of the Lens Model, and you know that you are not simply going to replace

Hank with a clerk who knows how to use equation 1 because you need Hank's intellect and experience to deal with unique cases.

What, then, might you recommend to Betty and Ralph about how to clone their best salespeople? One recommendation might be to go ahead and have a clerk code the data from each job seeker's application and enter it into a computer that is programmed to apply equation 1 and yield a score (the predicted Y_j). Except instead of replacing him with the equation, Hank could be given this score together with other information that is unique to the applicant and asked to make an overall decision about the qualifications of the applicant. This would mean that Hank would not have to learn to use the cues the way equation 1 uses them; he simply would use the output of equation 1 as information upon which his decision could be based. This would leave him in control, but it also would make sure that the information on the application was presented to him in a form that was both concise and valid. Of course, you would have to do a follow-up study to see if this new arrangement improved the sales force, which, after all, is the point.

❖ PROBLEMS WITH THE LENS MODEL

While the logic of the Lens Model has its attractions, it is rather simplistic, and the appropriateness of the multiple regression analyses can be questioned on at least two grounds.

First, equations 1 and 2 presume that the cues are additive—that the information provided by each cue simply piles on top of the other information, and as the pile increases, the applicant moves higher on the preference scale. However, it does not take much imagination to understand that information conveyed by some cues sometimes amplifies what is conveyed by other cues, rather than just adding to it, so the applicant would move up the scale in ever-increasing steps. Amplification is described mathematically as a multiplicative combination of information instead of the additive combination assumed by the model. That is, the cues may be non-independent (multiplicative), but the mathematics of multiple regression requires them to be independent (additive) if the equation is to make sense. The problem is important because even if the cues are not independent, the analysis will blindly impose additivity on them, with the result that the optimal policy (equation 1) or the decision

maker's policy (equation 2) will appear to be much simpler than it really is.

Second, the equation assumes that each cue is linearly related to the dependent variable. For example, the larger the cue value the higher the applicant will be on the preference scale. However, some cues have a curvilinear relationship to preference, for example, low and high values of the cue are associated with low preference and medium cue values are associated with high preference. Unless special steps are taken, the analysis will blindly impose linearity on the cues, and the resulting equations will suggest a greater simplicity than actually exists.

Those who defend the Lens Model point out that multiple regression is very robust. For example, it is possible to include non-independence in the equation by combining coded cue values before they are included in the analysis. However, they argue that from a practical standpoint there often is little gained by doing so. This was demonstrated very early on by Kort (1968), who showed that unless cues are very highly dependent (in which case one or the other might best be eliminated as redundant), the R^2 for the equation that uses combined coding usually is not much different from the R^2 for the equation that treats the cues as independent.

The argument continues that the equations do not have to be absolutely accurate representations of the environment or of the decision maker's policy to be valuable. That is, they can get you in the ball park even though they might not be right on the mark. As such, it usually is cheaper and less work to assume independence (or drop one of the dependent cues, or combine the dependent cues into a single measure), go with the simple equation, and then temper its use with a little common sense.

These same defenders of multiple regression use a similar argument in regard to curvilinear cues. Unless the curvilinear relationship is especially crucial, the analysis's imposition of linearity may not do much damage. By recoding the cue (so now both low and high values of the cue are recoded as low numbers to use in the equation, and medium values are recoded as high numbers), it is possible to derive a more accurate equation, but it often does not buy much for the effort; the R^2 may stay pretty much the same. Of course, there are other forms of non-linearity than curvilinearity, but much the same argument is made for them.

❖ THE LENS MODEL IN RESEARCH

The Lens Model has been used in many studies of decision making. In one of the earliest, Frederick Todd (reported in Hammond, 1955) examined 10 clinical psychologists' ability to use patients' responses to Rorschach cards (ink blots) to predict the IQ of each of 78 patients. IQ tests had previously been given to the patients, and their Rorschach responses had previously been coded using a standard coding system. The results of each patient's IQ test was Y_e and his or her coded Rorschach responses were the cues. Each clinician was given the coded Rorschach responses for each patient and asked to make a decision (Y_s) about each patient's IQ. Using multiple regression to obtain the optimal policy (across patients) yielded an R^2 of only .23, which means that it objectively is very difficult to predict IQ from Rorschach responses. The 10 clinicians' median achievement coefficient (squared) was .22, which means that their achievement was about as good as could be expected given that IQ cannot be reliably predicted from Rorschach responses. The median R^2 for their policy equations was .72, which means that even in the face of this unpredictability they tended to use the cues consistently—which is the best strategy in a low-predictability situation and which probably accounts for their achievement being as high as it was.

Among the more interesting results of early policy research is, for example, Rorer, Hoffman, Dickman, and Slovic's (1967) finding of no consistency among the policies of the attendants in a mental hospital in judging whether patients should have weekend passes; imagine how confusing that must have been for the patients. Similarly, Slovic, Rorer, and Hoffman (1971) found four different policies being used by various radiologists for judging malignancy of ulcers. Dawes (1971) modeled the decision policies of the admission committee for a doctoral program and found that replacing the committee with its own policy equation (called "bootstrapping") resulted in better predictions of applicant success than using the committee itself, presumably because the committee's membership changed over time and it therefore was not wholly consistent. Roose and Doherty (1976) studied how agency managers for an insurance company used application information to hire new agents and found that they relied too highly on a non-diagnostic cue; training them to use the optimal policy promised a substantial increase in successful hires.

Somewhat later, Dougherty, Ebert, and Callender (1986) studied how applicant information is used by employment interviewers in a large corporation; they had very consistent policies, but it was not clear that they were the best policies. Over the years there have been many studies done in the context of businesses or other organizations, but for proprietary reasons or because there is no payoff for the investigators to do so, little of this has been published in the scientific literature.

Social Judgment Theory

Since the 1970s, the point of view represented by the Lens Model and its applications has come to be known as Social Judgment Theory (Hammond, Rohrbaugh, Mumpower, & Adelman, 1977; Hammond, Stewart, Brehmer, & Steinman, 1975). In addition to examining interesting practical problems, such as the selection of bullets by the Denver Police Department (Hammond & Adelman, 1976, and see chapters in Brehmer & Joyce, 1988, for other applications), fundamental psychological issues also have been addressed. Among these is how people learn to use cues appropriately (Klayman, 1988). It is found that people learn most quickly when the cues have a simple, linear relationship to events (Y_e) and when the task content gives rise to reasonable hypotheses about weighting, but learning often is surprisingly slow. It turns out that having to learn by trial and error is very inefficient, but efficiency can be increased through provision of "cognitive feedback" (Hammond, 1971) consisting of information about the optimal policy. There also has been Social Judgment Theory research on decision processes in small groups, focusing largely on conflicts resulting from differences among the group members' various decision policies and how these differences can be reconciled (Rohrbaugh, 1988). An article by Hammond, Harvey, and Hastie (1992) provides a Social Judgment Theory analysis of social policy formation.

❖ SUMMARY

In this chapter we have examined how the Lens Model can be used to study both the decision maker's policy and the optimal policy. Moreover, the degree to which these policies account for Y_e and Y_{s_e} , as well as the degree to which the latter are correlated (achievement),



46 THE PSYCHOLOGY OF DECISION MAKING

permits us to prescribe different courses of action—replacement of the decision maker with mechanical application of the optimal policy or training the decision maker to use the optimal policy. As it has been developed in Social Judgment Theory, the logic of the Lens Model is applicable to a broad range of non-laboratory, socially interesting areas of decision.

