Perceptions of Fairness in Interpersonal and Individual Choice Situations

Max H. Bazerman, Sally Blount White, and George F. Loewenstein

You are graduating with a Ph.D. from a good psychology program. After a few interviews, a university that you are very interested in makes you an offer of an assistant professorship at $40,000 a year. The offer is not negotiable. You like the people. You like the job. You like the location. However, right before you are about to accept the offer, you find out that the same university is offering another new assistant professor $42,000. You do not see any characteristics that make the other individual more qualified than you. Will you still accept the offer?

Many people would be bothered by the difference between their starting salary and the starting salary of the other person, even if they were certain that this difference would not affect how they would be treated by the university in the future. That is because people care about fairness; it affects their decisions and how they feel about their lives.

Recent research has sought to understand how people judge particular situations as fair or unfair, and how fairness influences decision making. The findings from this research present a mixed picture. On the one hand, fairness judgments have been found to be remarkably nuanced, responding systematically to a wide range of situational and individual-level factors. On the other hand, there are pervasive biases in the way that people judge fairness, and inconsistencies in the importance they place on fairness in different contexts.

This article reviews research on fairness perceptions and on the impact of fairness on decision making. We concentrate on distributive, as opposed to procedural, notions of fairness—that is, on how people evaluate the fairness of a resource distribution, not the fairness of the procedure used to create the distribution. First, we briefly review some of the past research on the determinants of fairness perceptions and on the impact of fairness on decision making. Then, we discuss our own current research, which examines inconsistencies in the impact of fairness perceptions on decision making.

Max H. Bazerman is the J. Jay Gerber Distinguished Professor of Dispute Resolution at the J.L. Kellogg Graduate School of Management, Northwestern University. **Sally Blount White** is Assistant Professor of Behavioral Science at the Graduate School of Business, University of Chicago. **George F. Loewenstein** is Professor of Economics in the Department of Social and Decision Sciences, Carnegie Mellon University. Address correspondence to Max H. Bazerman, J.L. Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Rd., Evanston, IL 60208; e-mail: mbazer@casbah.acns.nwu.edu.

THE ROLE OF FAIRNESS IN DECISION MAKING

Research reported in the organizational and social psychological literatures has shown that both the degree of concern for other people's outcomes and the nature of that concern depend on a variety of factors. For example, Deutsch[1] identified three principles that underlie many judgments of distributive fairness: equity, equality, and need. Under equity, outcomes are rewarded in direct proportion to earned rights or inputs. Under equality, outcomes are distributed equally. Need differs from equality in that resources are distributed first to those in most need until equality of general circumstances is obtained. Deutsch suggested, and subsequent research has confirmed, that the particular principle people tend to invoke depends on the goals prevailing in a particular situation. Equity concerns tend to prevail in relationships focusing on economic productivity, equality concerns in relations that emphasize enjoyable social relations, and needs in relations that foster personal welfare as a primary goal.

Yaari and Bar-Hillel[2] found, in a two-person distributional context, that the needs, tastes, and beliefs of the two parties were important determinants of fairness judgments. They asked subjects to judge the fairness of alternative distributions of commodities between two parties who differed from one another in their needs, tastes, and beliefs. They found that fairness judgments were most influenced by the parties' needs, that tastes were secondary, and that beliefs had the smallest impact. Overall, subjects' preferred divisions reflected a mix of equality and the desire to provide the resources to the party that valued the resources more (social efficiency). For example, if A liked a commodity slightly more than B, subjects judged outcomes in which B was allocated slightly more of the commodity as being fair. Although A valued each unit more, B received a few extra units to compensate for the lower value that B received from each unit. However, if B placed very little value on the commodity, subjects thought it was fair to give most or all of it to A, thereby maximizing social efficiency. Thus, the equalization of subjective payoffs was perceived as most fair when it was relatively easy to achieve; however, when equality

was more difficult to implement, the goal of maximizing social efficiency became more prominent.

Kahneman, Knetsch, and Thaler[3] found that the preexisting expectations of parties influenced fairness judgments. They conducted a series of telephone interviews in which they asked subjects to judge the fairness of various practices. For example, subjects were asked to evaluate the fairness of a hardware store that raises the price of snow shovels from $15 to $20 the morning after a snowstorm. The economically rational action is to raise the price if doing so increases the store's profit because, from an efficiency standpoint, raising the price ensures that those people who value the shovels the most highly will obtain them.

Despite the economic rationality of raising the snow shovels' price, 82% of respondents considered this action unfair. Kahneman et al. found that actions which made some party worse off than the status quo were generally viewed as unfair, unless it was unavoidable that someone would be hurt. For example, respondents thought it unfair for a profitable firm to decrease wages even though other firms were doing so, but if the firm were losing money, this action was judged to be fair. Likewise, dramatically reducing a new worker's wages or raising a new tenant's rent was viewed as fair, but these measures were not viewed as fair when applied to existing workers or tenants. Such a "don't make anyone worse off" principle seems to underlie the Supreme Court's ruling that a firm whose past hiring practices discriminated against minority workers can hire minority replacements for workers who retire or leave the firm voluntarily, but cannot fire existing majority workers to achieve parity.[4]

Although the work of Yaari and Bar-Hillel and of Kahneman et al. suggests that people make remarkably sophisticated judgments of fairness, taking systematic account of factors such as needs, wants, beliefs, and prior expectations, other research presents a less sophisticated picture of fairness judgments. Several studies have found that parties involved in disputes tend to arrive at judgments of fair settlement points that favor themselves. For example, in one study,[5] subjects and a legal expert reviewed the case materials from a lawsuit in Texas. The expert reached a judgment in the form of an award to the plaintiff. Subjects, who had been assigned the role of plaintiff or defendant, predicted the expert's decision, determined their own judgment of a fair outcome, and attempted to negotiate a settlement. Although subjects in both roles read identical materials and were paid for predicting the decision accurately, predictions of the legal expert's settlement and assessments of what a fair settlement would be were significantly biased by the subject's role. On average, defendants guessed that the expert's award would be about $20,000 lower than did plaintiffs—a huge difference given that the actual award amount was about $30,000. Furthermore, the greater the difference between the two parties' predictions of the expert's decision and between their assessments of what a fair outcome would be, the less likely they were able to negotiate a settlement.

In sum, existing research suggests that fairness judgments are responsive to a wide range of factors, including the goals prevailing in a relationship and the needs, wants, and beliefs of the individuals. Fairness judgments seem to reflect basic principles that have wide acceptance in economics, such as "promote equity unless it undermines efficiency unacceptably" and "don't make anyone worse off than before unless it's unavoidable." However, other research has found that what is considered fair in allocation situations is often biased by self-interest. Moreover, the "don't make people worse off" notion of fairness has conservative implications for public policy because it favors the status quo, however inequitable that may be. Most changes in policy, including those that reduce preexisting inequality, make some people worse off than they were before and are thus likely to be viewed as unfair.

INCONSISTENCY IN THE TENDENCY TO VALUE COMPARATIVE OUTCOMES

Our own recent research explores how people in multiparty transactions weigh their own payoffs against fairness considerations. Consider the ultimatum bargaining game, now ubiquitous in behavioral economics and social psychology. The game involves two players, one assigned the role of divider and the other assigned the role of chooser. The divider is presented with a sum of money (e.g., $10) and proposes a split of the money, which the chooser can either accept or reject. If the chooser accepts the proposed division, then the parties receive the allocation proposed by the divider. If the chooser rejects the proposed division, then neither party gets anything.

Numerous experiments involving minor variants of the ultimatum game have found that individuals do not behave in the way that economic models predict. The economic solution predicts that Player 1 will offer Player 2 the smallest possible amount above zero and keep the rest for himself or herself. Further, Player 2 will accept this offer because it is greater than zero, which is the alternative outcome. However, not surprisingly to psychologists, dividers in fact typically offer choosers more than a trivial amount, favoring most often an equal split (e.g., $5/$5). Furthermore, choosers who receive offers that favor the proposing subject often reject them, preferring to get nothing at all.

As this example shows, when evaluating multiparty transactions, people seem to compare their own outcomes with other people's. Although social psychologists and sociologists have written about this tendency since the 1950s,[6] this insight has only recently been formalized in decision-making models through the introduction of the concept of social utility. Social utility functions represent an individual's level of satisfaction (utility) with an outcome as a function of the payoff to himself or herself and the payoffs to other interdependent parties.

We[7] have found that individuals' social utility functions are nonlinear and that the form of the utility function depends on a number of factors, such as the quality of the relationship between the parties. In these studies, we placed subjects in a situation characterized by a strong equality norm and asked them to assess their satisfaction with different monetary outcomes for themselves and for another person in a number of fictional disputes. A social utility function for each subject was then computed by simultaneously regressing the subject's satisfaction with each potential outcome against (a) the straight monetary value of that outcome to the subject and (b) the difference between the monetary value of the subject's outcome and the other person's outcome. In general, we found that interpersonal comparisons overwhelmed concern for personal outcomes. For example, most individuals rated the outcome of $500 for oneself and $500 for the other person as more satisfactory than the outcome of $600 for oneself and $800 for the other person. Thus, they preferred equal outcomes in which they received less money over unequal outcomes in which they received more. Yet, from a social benefit perspective, the second outcome is clearly the better one, because both parties receive more money.

In more recent research,[8] we built upon these results and found that another factor that affects how people weigh comparative payoffs is the informational context within which the outcome is viewed. Specifically, when individuals evaluate outcomes independently (i.e., on a case-by-case basis), relative payoffs tend to be given enormous weight. However, when asked to choose between alternative outcomes (e.g., "Would you prefer $500 for yourself and $500 for the other party, or $600 for yourself and $700 for the other party?"), people place greater weight on their own payoff than on the disparity between the two payoffs. For example, 70% of subjects evaluated Outcome A ($400 for self, $400 for other) as more acceptable than Outcome B ($500 for self, $700 for other) when they rated these outcomes sequentially, but only 22% selected Outcome A over Outcome B when presented with a simultaneous choice between the two. This basic pattern of preference reversal was consistent across a number of comparisons in a variety of contexts.

In ongoing research, we extended this result to a situation involving real payoffs: In recruiting subjects for a colleague's experiment, we made different offers to different groups of potential subjects. One group was offered $7 to participate in a 40-min experiment and was told that all subjects would be receiving $7. A second group was offered $8 to participate in a 40-min experiment and was told that some subjects would be paid $10 for the same work (based on the last digit of their social security number). A third group was given an opportunity to participate in a 40-min experiment in which everyone would be paid $7; to participate in a 40-min experiment in which some subjects, including themselves, would receive $8 and others would receive $10; or not to participate at all. Although significantly more subjects in the first group chose to participate (72%) than in the second group (55%), the majority of subjects in the third group chose to participate in the experiment that gave them $8 while some others were given $10 (56% chose the $8/$10 experiment; 16% chose the $7/$7 experiment; 28% chose not to participate in either). Thus, in the decision whether to participate in one experiment, the outcomes of other potential subjects were critical. However, when multiple opportunities were available, the outcomes of others became less important than simply maximizing one's own pay for 40 min of work.

In another context, 2nd-year M.B.A. students at Kellogg (in the fall of 1991) reviewed hypothetical job offers and were asked whether they would accept or reject them.[9] Some subjects were presented with one offer at a time. Others were always given two job offers at a time. In both cases, they were told to imagine that it was January 15, 1992, and that the job offer or offers expired today. Subjects given one offer were then asked whether they would accept the offer or reject the offer and remain on the market. Subjects given two offers were asked if they would accept one of the jobs and, if so, which one. Two of the jobs that were assessed were described as follows:

Job A: The offer is from Company 4 for $75,000 a year. It is widely known that this firm pays all starting MBAs from top schools $75,000. (additional descriptive information about the firm was then provided)

Job B: The offer is from Company 9 for $85,000 a year. It is widely known that this firm is paying some other graduating Kellogg students $95,000 a year. (additional descriptive information about the firm was then provided)

As expected, individuals examining one job offer at a time were more likely to accept Job A than to accept Job B, but individuals examining Job A and Job B simultaneously were more likely to accept Job B than Job

A. This pattern of outcomes held across several different job comparisons.

Finally, our ongoing research applied the insights from this research back to the ultimatum game. We conducted two versions of the game, one in which choosers stated a minimum amount that they would accept (with the consequence of not accepting clearly explained) from a $10 pool, and another in which they made a series of choices in which they decided explicitly between various splits of the $10 pie or both parties getting nothing. The first condition was intended to approximate one-at-a-time evaluation and the second comparative evaluation; thus, we expected fairness to be weighted more heavily in the former. As hypothesized, the mean minimum acceptable amount was significantly higher in the one-at-a-time evaluation condition ($4.00), in which people stated minimum amounts, than in the choice condition ($2.33), in which they made explicit trade-offs.

Together, our studies suggest that when evaluating outcomes in isolation, people tend to be more concerned with interpersonal comparison of outcomes than with maximizing personal outcomes. This concern is often justified by calling upon norms of fairness. However, when evaluating more than one outcome at a time, people are less concerned with comparative payoffs than with maximizing their own payoff. These results imply that if people make policy decisions on a case-by-case basis, they may have a tendency to base these decisions on perceptions of fairness that are suboptimal for themselves and for society as a whole.

CONCLUDING COMMENTS

Judgments of fairness present human decision making at its best and worst. On the one hand, such judgments display a systematic and nuanced responsiveness to numerous factors, only a fraction of which we have been able to review in this essay. On the other hand, these judgments appear to be biased in a self-serving manner and to respond to factors that are difficult to justify on normative grounds. Our own recent research has focused on one particularly dramatic example—the tendency for relative comparisons to receive much greater weight in one-at-a-time evaluation than in comparative choice.

Because decision problems vary in their form of presentation, the discrepancy between one-at-a-time and comparative evaluation is likely to have significant ramifications in the real world. For example, job searchers often do make sequential take-it-or-leave-it decisions about jobs but, at other times, choose explicitly between jobs. Our research indicates that people are likely to arrive at different choices in these two situations, a discrepancy that is difficult to justify normatively.

In our current research, we are exploring possible causes of the discrepancy between one-at-a-time and comparative evaluation. Simultaneously with our own work, many researchers have been documenting preference reversals between independent and comparative evaluation in situations that do not involve fairness or social comparison, suggesting that the phenomenon is not unique to the interpersonal domain. Instead, a variety of conditions seem to cause such reversals. For example, attributes that are difficult to evaluate in isolation are given greater weight in choice, emotional attributes are given greater weight in independent judgment than in choice, and attributes that are easy to justify taking into account are weighted more heavily in choice than in independent judgment.[10] The preference reversals we have observed may be particularly strong because all of these factors are at play. Payoffs to self are generally difficult to evaluate out of context, payoff inequity tends to evoke powerful emotions, and it is easier to justify maximizing one's own payoff than to justify maximizing fairness. All of these factors contribute to the greater weighting of one's own payment in choice and of fairness considerations when alternatives are evaluated independently.

Ample evidence supports the argument that people use fairness and social comparison information to interpret their world; judgments of fairness permeate life. The assistant professor in the story at the beginning of this article cannot simply ignore the comparison information that has been provided. However, the strength of that concern and its impact on decision making is by no means predetermined. Our hope is that an understanding of the variable influence of fairness on decision making could contribute to decisions that are made more coherently and thus are more consistent with long-term interests.

Acknowledgments—This article was supported by the Dispute Resolution Research Center at Northwestern University and benefited from the helpful comments of Sandra Scarr and two anonymous reviewers.

Notes

1. M. Deutsch, Equity, equality, and need: What determines which value will be used as the basis of distributive justice, *Journal of Social Issues, 31*, 137–149 (1975)

2. M E. Yaari and M. Bar-Hillel. On dividing justly, *Social Choice and Welfare. 1*, 1–24 (1984)

3. D. Kahneman, J L. Knetsch, and R H. Thaler, Fairness and the assumptions of economics, in *Rational Choice: The Contrast Between Economics and Psychology*. R.W. Hogarth and M W. Reder, Eds. (University of Chicago Press, Chicago, 1986)

4. S. Issacharoff. When substance mandates procedure: *Martin v. Wilks* and the rights of vested incumbents in civil rights consent decrees. *Cornell Law Review. 77*, 189–253 (1992)

5. G. Loewenstein, S. Issacharoff, C. Camerer, and L. Babcock. Self-serving assessments of fairness and pretrial bargaining. *Journal of Legal Studies. XXII*, 135–159 (1993).

6. L. Festinger. *A Theory of Cognitive Dissonance* (Row. Peterson. Evanston. IL. 1957)

7. G F. Loewenstein. L L. Thompson, and M H. Bazerman. Social utility and decision making in in-

terpersonal contexts. *Journal of Personality and Social Psychology, 57,* 426–441 (1989)

8 M H. Bazerman, G F Loewenstein, and S B White, Psychological determinants of utility in competitive contexts: The impact of elicitation procedure, *Administrative Science Quarterly. 37.* 220–240 (1992)

9 M H. Bazerman. H Schroth, P. Pradhan, K. Diekmann, and A Tenbrunsel, Inconsistent preferences in job acceptance: The role of social comparison processes and procedural justice, *Organizational Behavior and Human Decision Processes* (in press).

10. These three effects can be found in the following three papers. respectively: D Lowenthal, *Reversals of preference in candidate choice.* unpublished manuscript, Carnegie-Mellon University. Pittsburgh (1993); D Kahneman and I Ritov, Determinants of stated willingness to pay for public goods: A study of the headline method. *Journal of Risk and Uncertainty, 9,* 5–38 (1994); and C K Hsee, *Elastic justification: Affective influences on monetary decision making.* unpublished manuscript, University of Chicago, Chicago (1993)

Motion Perception as a Partnership: Exogenous and Endogenous Contributions

Robert Sekuler

William James portrayed perception as a partnership: "Whilst part of what we perceive comes through our senses from the object before us, another part (and it may be the larger part) always comes . . . out of our own head."[1] On this view, visual experience is constructed not only from exogenous resources such as the raw material provided by the eye, but also from endogenous resources, including knowledge, memory, and expectation. This review shows that visual motion, like other forms of perception, emerges from a Jamesian partnership between endogenous and exogenous influences.

LIVING WITH UNCERTAINTY

For one study of motion detection, Karlene Ball and I devised conditions that would maximize or minimize the validity of the observer's expectations, a key endogenous influence.[2] In the first of these conditions, the target's direction was constant over trials so that an observer could be certain about the motion that had to be detected; in the second condition, randomization kept the observer uncertain about what direction to expect. We reasoned that certainty might allow an observer to attend to the one appropriate direction. Our stimuli, spatially random dots moving in the same direction at a common, constant speed, covered the entire display of 8° diameter. As a result, there was no advantage in looking at or attending to any particular location. Instead, the observer could attend to the expected direction by monitoring signals within neurons that are tuned to the expected direction.[3]

Each trial comprised two 600-ms intervals. Selected at random, one interval contained low-contrast, spatially random, moving dots; the other interval contained only a uniform, veiling luminance. The experiment interleaved two types of blocks. In certainty blocks, dots always moved in the same, predictable direction, upward. In uncertainty blocks, the direction was unpredictable, alternating randomly between upward and rightward. The subject was not asked to judge direction of motion, but merely to identify which interval contained the stimulus.

In the certainty condition, the 3 subjects averaged approximately 75% correct detection (Fig. 1). With uncertainty, detection of the upward-moving dots declined to only slightly above chance, 53%. So, without altering any property of the stimulus, but reducing the value of endogenous information, we rendered a moving stimulus virtually invisible.

Thus, expectation of a particular direction of motion modulates the detectability of a moving target. In another study, to learn how rapidly such expectations can be constructed or changed, on each trial, we gave the subject a hint about the motion that might be presented.[4] The direction in which the dots moved varied randomly from trial to trial, over a range of 360°. On some trials, a cue to the upcoming direction was presented for 50 ms. This cue, which appeared at various times before or after the test interval, was a line radiating out from the center of the display, and its orientation corresponded to the dots' direction of motion. Results with the cue were compared with results from two control conditions without a cue. In one control condition, the direction of motion varied randomly from trial to trial; in the other control condition, the direction of motion was fixed (upward), allowing the subject to be certain about which direction would be presented. Figure 2 illustrates the sequence of events in the trials.

On half the trials, low-contrast,

Robert Sekuler is Frances and Louis Salvage Professor of Psychology at Brandeis University, and is in the Departments of Biomedical Engineering and Cognitive & Neural Systems at Boston University. Address correspondence to Robert Sekuler, Volen Center for Complex Systems, Brandeis University, Waltham, MA 02254; e-mail: sekuler@binah.cc.brandeis.edu.