




# Daniel Kahneman

## The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2002

### Autobiography



#### Early years



I was born in Tel Aviv, in what is now Israel, in 1934, while my mother was visiting her extended family there; our regular domicile was in Paris. My parents were Lithuanian Jews, who had immigrated to France in the early 1920s and had done quite well. My father was the chief of research in a large chemical factory. But although my parents loved most things French and had some French friends, their roots in France were shallow, and they never felt completely secure. Of course, whatever vestiges of security they'd had were lost when the Germans swept into France in 1940. What was probably the first graph I ever drew, in 1941, showed my family's fortunes as a function of time - and around 1940 the curve crossed into the negative domain.

I will never know if my vocation as a psychologist was a result of my early exposure to interesting gossip, or whether my interest in gossip was an indication of a budding vocation. Like many other Jews, I suppose, I grew up in a world that consisted exclusively of people and words, and most of the words were about people. Nature barely existed, and I never learned to identify flowers or to appreciate animals. But the people my mother liked to talk about with her friends and with my father were fascinating in their complexity. Some people were better than others, but the best were far from perfect and no one was simply bad. Most of her stories were touched by irony, and they all had two sides or more.

In one experience I remember vividly, there was a rich range of shades. It must have been late 1941 or early 1942. Jews were required to wear the Star of David and to obey a 6 p.m. curfew. I had gone to play with a Christian friend and had stayed too late. I turned my brown sweater inside out to walk the few blocks home. As I was walking down an empty street, I saw a German soldier approaching. He was wearing the black uniform that I had been told to fear more than others - the one worn by specially recruited SS soldiers. As I came closer to him, trying to walk fast, I noticed that he was looking at me intently. Then he beckoned me over, picked me up, and hugged me. I was terrified that he would notice the star inside my sweater. He was speaking to me with great emotion, in German. When he put me down, he opened his wallet, showed me a picture of a boy, and gave me some money. I went home more certain than ever that my mother was right: people were endlessly complicated and interesting.

My father was picked up in the first large-scale sweep for Jews, and was interned for six weeks in Drancy, which had been set up as a way station to the extermination camps. He was released through the intervention of his firm, which was directed (a fact I learned only from an article I read a few years ago) by the financial mainstay of the Fascist anti-Semitic movement in France in the 1930s. The story of my father's release, which I never fully understood, also involved a beautiful woman and a German general who loved her. Soon afterward, we escaped to Vichy France, and stayed on the Riviera in relative safety, until the Germans arrived and we escaped again, to the center of France. My father died of inadequately treated diabetes, in 1944, just six weeks before the D-day he had been waiting for so desperately. Soon my mother, my sister, and I were free, and beginning to hope for the permits that would allow us to join the rest of our family in Palestine.

I had grown up intellectually precocious and physically inept. The ineptitude must have been quite remarkable, because during my last term in a French lycée, in 1946, my eighth-grade physical-education teacher blocked my inclusion in the Tableau d'Honneur - the Honor Roll - on the grounds that even his extreme tolerance had limits. I must also have been quite a pompous child. I had a notebook of essays, with a title that still makes me blush: "What I write of what I think." The first essay, written before I turned eleven, was a discussion of faith. It approvingly quoted Pascal's saying "Faith is God made perceptible to the heart" ("How right this is!"), then went on to point out that this genuine spiritual experience was probably rare and unreliable, and that cathedrals and organ music had been created to generate a more reliable, ersatz version of the thrills of faith. The child who wrote this had some aptitude for psychology, and a great need for a normal

-  Printer Friendly
-  Comments & Questions
-  Tell a Friend

The 2002 Prize in:

Economics

The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2002

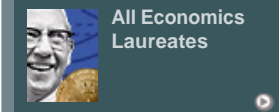
- Press Release
- Advanced Information
- Information for the Public
- Presentation Speech

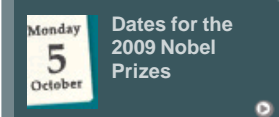
Daniel Kahneman

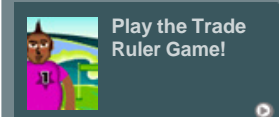
- Autobiography
- Prize Lecture
- Interview
- Diploma
- Photo Gallery
- Prize Presentation
- Other Resources

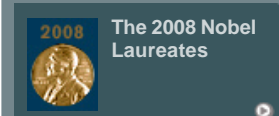
Vernon L. Smith

- Autobiography
- Prize Lecture
- Interview
- Diploma
- Photo Gallery
- Prize Presentation
- Banquet Speech
- Other Resources









life.

### Adolescence

The move to Palestine completely altered my experience of life, partly because I was held back a year and enrolled in the eighth grade for a second time - which meant that I was no longer the youngest or the weakest boy in the class. And I had friends. Within a few months of my arrival, I had found happier ways of passing time than by writing essays to myself. I had much intellectual excitement in high school, but it was induced by great teachers and shared with like-minded peers. It was good for me not to be exceptional anymore.

At age seventeen, I had some decisions to make about my military service. I applied to a unit that would allow me to defer my service until I had completed my first degree; this entailed spending the summers in officer-training school, and part of my military service using my professional skills. By that time I had decided, with some difficulty, that I would be a psychologist. The questions that interested me in my teens were philosophical - the meaning of life, the existence of God, and the reasons not to misbehave. But I was discovering that I was more interested in what made people believe in God than I was in whether God existed, and I was more curious about the origins of people's peculiar convictions about right and wrong than I was about ethics. When I went for vocational guidance, psychology emerged as the top recommendation, with economics not too far behind.

I got my first degree from the Hebrew University in Jerusalem, in two years, with a major in psychology and a minor in mathematics. I was mediocre in math, especially in comparison with some of the people I was studying with - several of whom went on to become world-class mathematicians. But psychology was wonderful. As a first-year student, I encountered the writings of the social psychologist Kurt Lewin and was deeply influenced by his maps of the life space, in which motivation was represented as a force field acting on the individual from the outside, pushing and pulling in various directions. Fifty years later, I still draw on Lewin's analysis of how to induce changes in behavior for my introductory lecture to graduate students at the Woodrow Wilson School of Public Affairs at Princeton. I was also fascinated by my early exposures to neuropsychology. There were the weekly lectures of our revered teacher Yeshayahu Leibowitz - I once went to one of his lectures with a fever of 41 degrees Celsius; they were simply not to be missed. And there was a visit by the German neurosurgeon Kurt Goldstein, who claimed that large wounds to the brain eliminated the capacity for abstraction and turned people into concrete thinkers. Furthermore, and most exciting, as Goldstein described them, the boundaries that separated abstract from concrete were not the ones that philosophers would have set. We now know that there was little substance to Goldstein's assertions, but at the time the idea of basing conceptual distinctions on neurological observations was so thrilling that I seriously considered switching to medicine in order to study neurology. The Chief of Neurosurgery at the Hadassah Hospital, who was a neighbor, wisely talked me out of that plan by pointing out that the study of medicine was too demanding to be undertaken as a means to any goal other than practice.

### The military experience

In 1954, I was drafted as a second lieutenant, and after an eventful year as a platoon leader I was transferred to the Psychology branch of the Israel Defense Forces. There, one of my occasional duties was to participate in the assessment of candidates for officer training. We used methods that had been developed by the British Army in the Second World War. One test involved a leaderless group challenge, in which eight candidates, with all insignia of rank removed and only numbers to identify them, were asked to lift a telephone pole from the ground and were then led to an obstacle, such as a 2.5-meter wall, where they were told to get to the other side of the wall without the pole touching either the ground or the wall, and without any of them touching the wall. If one of these things happened, they had to declare it and start again. Two of us would watch the exercise, which often took half an hour or more. We were looking for manifestations of the candidates' characters, and we saw plenty: true leaders, loyal followers, empty boasters, wimps - there were all kinds. Under the stress of the event, we felt, the soldiers' true nature would reveal itself, and we would be able to tell who would be a good leader and who would not. But the trouble was that, in fact, we could not tell. Every month or so we had a "statistics day," during which we would get feedback from the officer-training school, indicating the accuracy of our ratings of candidates' potential. The story was always the same: our ability to predict performance at the school was negligible. But the next day, there would be another batch of candidates to be taken to the obstacle field, where we would face them with the wall and see their true natures revealed. I was so impressed by the complete lack of connection between the statistical information and the compelling experience of insight that I coined a term for it: "the illusion of validity." Almost twenty years later, this term made it into the technical literature (Kahneman and Tversky, 1973). It was the first cognitive illusion I discovered.

Closely related to the illusion of validity was another feature of our discussions about the candidates we observed: our willingness to make extreme predictions about their future performance on the basis of a small sample of behavior. In fact, the issue of willingness did not arise, because we did not really distinguish predictions from observations. The soldier who took over when the group was in trouble and led the team over the wall was a leader at that moment, and if we asked ourselves how he would perform in officer-training, or on the battlefield, the best bet was simply that he would be as good a leader then as he was now. Any other prediction seemed inconsistent with the evidence. As I understood clearly only when I taught statistics some years later, the idea that predictions

should be less extreme than the information on which they are based is deeply counterintuitive.

The theme of intuitive prediction came up again, when I was given the major assignment for my service in the Unit: to develop a method for interviewing all combat-unit recruits, in order to screen the unfit and help allocate soldiers to specific duties. An interviewing system was already in place, administered by a small cadre of interviewers, mostly young women, themselves recent graduates from good high schools, who had been selected for their outstanding performance in psychometric tests and for their interest in psychology. The interviewers were instructed to form a general impression of a recruit and then to provide some global ratings of how well the recruit was expected to perform in a combat unit. Here again, the statistics of validity were dismal. The interviewers' ratings did not predict with substantial accuracy any of the criteria in which we were interested.

My assignment involved two tasks: first, to figure out whether there were personality dimensions that mattered more in some combat jobs than in others, and then to develop interviewing guidelines that would identify those dimensions. To perform the first task, I visited units of infantry, artillery, armor, and others, and collected global evaluations of the performance of the soldiers in each unit, as well as ratings on several personality dimensions. It was a hopeless task, but I didn't realize that then. Instead, spending weeks and months on complex analyses using a manual Monroe calculator with a rather iffy handle, I invented a statistical technique for the analysis of multi-attribute heteroscedastic data, which I used to produce a complex description of the psychological requirements of the various units. I was capitalizing on chance, but the technique had enough charm for one of my graduate-school teachers, the eminent personnel psychologist Edwin Ghiselli, to write it up in what became my first published article. This was the beginning of a lifelong interest in the statistics of prediction and description.

I had devised personality profiles for a criterion measure, and now I needed to propose a predictive interview. The year was 1955, just after the publication of "Clinical versus statistical prediction" (Meehl, 1954), Paul Meehl's classic book in which he showed that clinical prediction was consistently inferior to actuarial prediction. Someone must have given me the book to read, and it certainly had a big effect on me. I developed a structured interview schedule with a set of questions about various aspects of civilian life, which the interviewers were to use to generate ratings about six different aspects of personality (including, I remember, such things as "masculine pride" and "sense of obligation"). Soon I had a near-mutiny on my hands. The cadre of interviewers, who had taken pride in the exercise of their clinical skills, felt that they were being reduced to unthinking robots, and my confident declarations - "Just make sure that you are reliable, and leave validity to me" - did not satisfy them. So I gave in. I told them that after completing "my" six ratings as instructed, they were free to exercise their clinical judgment by generating a global evaluation of the recruit's potential in any way they pleased. A few months later, we obtained our first validity data, using ratings of the recruits' performance as a criterion. Validity was much higher than it had been. My recollection is that we achieved correlations of close to .30, in contrast to about .10 with the previous methods. The most instructive finding was that the interviewers' global evaluation, produced at the end of a structured interview, was by far the most predictive of all the ratings they made. Trying to be reliable had made them valid. The puzzles with which I struggled at that time were the seed of the paper on the psychology of intuitive prediction that Amos Tversky and I published much later.

The interview system has remained in use, with little modification, for many decades. And if it appears odd that a twenty-one-year-old lieutenant would be asked to set up an interviewing system for an army, one should remember that the state of Israel and its institutions were only seven years old at the time, that improvisation was the norm, and that professionalism did not exist. My immediate supervisor was a man with brilliant analytical skills, who had trained in chemistry but was entirely self-taught in statistics and psychology. And with a B.A. in the appropriate field, I was the best-trained professional psychologist in the military.

#### Graduate school years

I came out of the Army in 1956. The academic planners at the Hebrew University had decided to grant me a fellowship to obtain a PhD abroad, so that I would be able to return and teach in the psychology department. But they wanted me to acquire some additional polish before facing the bigger world. Because the psychology department had temporarily closed, I took some courses in philosophy, did some research, and read psychology on my own for a year. In January of 1958, my wife, Irah, and I landed at the San Francisco airport, where the now famous sociologist Amitai Etzioni was waiting to take us to Berkeley, to the Flamingo Motel on University Avenue, and to the beginning of our graduate careers.

My experience of graduate school was quite different from that of students today. The main landmarks were examinations, including an enormous multiple-choice test that covered all of psychology. (A long list of classic studies preceded by the question "Which of the following is not a study of latent learning?" comes to mind.) There was less emphasis on formal apprenticeship, and virtually no pressure to publish while in school. We took quite a few courses and read broadly. I remember a comment of Professor Rosenweig's on the occasion of my oral exam. I should enjoy my current state, he advised, because I would never again know as much psychology. He was right.

I was an eclectic student. I took a course on subliminal perception from Richard Lazarus, and wrote with him a speculative article on the temporal development of percepts, which was soundly and correctly rejected. From that subject I came to an interest in the more technical aspects of vision and I spent some time learning about optical benches from Tom Cornsweet. I audited the clinical sequence, and learned about personality tests from Jack Block and from Harrison Gough. I took classes on Wittgenstein in the philosophy department. I dabbled in the philosophy of science. There was no particular rhyme or reason to what I was doing, but I was having fun.

My most significant intellectual experience during those years did not occur in graduate school. In the summer of 1958, my wife and I drove across the United States to spend a few months at the Austen Riggs Clinic in Stockbridge, Massachusetts, where I studied with the well-known psychoanalytic theorist David Rapaport, who had befriended me on a visit to Jerusalem a few years earlier. Rapaport believed that psychoanalysis contained the elements of a valid theory of memory and thought. The core ideas of that theory, he argued, were laid out in the seventh chapter of Freud's "Interpretation of Dreams," which sketches a model of mental energy (cathexis). With the other young people in Rapaport's circle, I studied that chapter like a Talmudic text, and tried to derive from it experimental predictions about short-term memory. This was a wonderful experience, and I would have gone back if Rapaport had not died suddenly later that year. I had enormous respect for his fierce mind. Fifteen years after that summer, I published a book entitled "Attention and Effort," which contained a theory of attention as a limited resource. I realized only while writing the acknowledgments for the book that I had revisited the terrain to which Rapaport had first led me.

Austen Riggs was a major intellectual center for psychoanalysis, dedicated primarily to the treatment of dysfunctional descendants of wealthy families. I was allowed into the case conferences, which were normally scheduled on Fridays, usually to evaluate a patient who had spent a month of live-in observation at the clinic. Those attending would have received and read, the night before, a folder with detailed notes from every department about the person in question. There would be a lively exchange of impressions among the staff, which included the fabled Erik Erikson. Then the patient would come in for a group interview, which was followed by a brilliant discussion. On one of those Fridays, the meeting took place and was conducted as usual, despite the fact that the patient had committed suicide during the night. It was a remarkably honest and open discussion, marked by the contradiction between the powerful retrospective sense of the inevitability of the event and the obvious fact that the event had not been foreseen. This was another cognitive illusion to be understood. Many years later, Baruch Fischhoff wrote, under my and Amos Tversky's supervision, a beautiful PhD thesis that illuminated the hindsight effect.

In the spring of 1961, I wrote my dissertation on a statistical and experimental analysis of the relations between adjectives in the semantic differential. This allowed me to engage in two of my favorite pursuits: the analysis of complex correlational structures and FORTRAN programming. One of the programs I wrote would take twenty minutes to run on the university mainframe, and I could tell whether it was working properly by the sequence of movement on the seven tape units that it used. I wrote the thesis in eight days, typing directly on the purple "ditto" sheets that we used for duplication at the time. That was probably the last time I wrote anything without pain. The paper itself, by sharp contrast, was so convoluted and dreary that my teacher, Susan Ervin, memorably described the experience of reading it as "wading through wet mush." I spent the summer of 1961 in the ophthalmology department, doing research on contour interference. And then it was time to go home to Jerusalem, and start teaching in the psychology department at the Hebrew University.

#### Training to become a professional

I loved teaching undergraduates and I was good at it. The experience was consistently gratifying because the students were so good: they were selected on the basis of a highly competitive entrance exam, and most were easily PhD material. I took charge of the basic first-year statistics class and, for some years, taught both that course and the second-year course in research methods, which also included a large dose of statistics. To teach effectively I did a lot of serious thinking about valid intuitions on which I could draw and erroneous intuitions that I should teach students to overcome. I had no idea, of course, but I was laying the foundation for a program of research on judgment under uncertainty. Another course I was teaching concerned the psychology of perception, which also contributed quite directly to the same program.

I had learned a lot in Berkeley, but I felt that I had not been adequately trained to do research. I therefore decided that in order to acquire the basic skills I would need to have a proper laboratory and do regular science - I needed to be a solid short-order cook before I could aspire to become a chef. So I set up a vision lab, and over the next few years I turned out competent work on energy integration in visual acuity. At the same time, I was trying to develop a research program to study affiliative motivation in children, using an approach that I called a "psychology of single questions." My model for this kind of psychology was research reported by Walter Mischel (1961a, 1961b) in which he devised two questions that he posed to samples of children in Caribbean islands: "You can have this (small) lollipop today, or this (large) lollipop tomorrow," and "Now let's pretend that there is a magic man ... who could change you into anything that you would want to be, what you would want to be?" The answer to the latter question was scored 1, if it referred to a profession or to an achievement-related trait, otherwise 0. The responses to these lovely questions turned out to be plausibly correlated with numerous characteristics

of the child and the child's background. I found this inspiring: Mischel had succeeded in creating a link between an important psychological concept and a simple operation to measure it. There was (and still is) almost nothing like it in psychology, where concepts are commonly associated with procedures that can be described only by long lists or by convoluted paragraphs of prose.

I got quite nice results in my one-question studies, but never wrote up any of the work, because I had set myself impossible standards: in order not to pollute the literature, I wanted to report only findings that I had replicated in detail at least once, and the replications were never quite perfect. I realized only gradually that my aspirations demanded more statistical power and therefore much larger samples than I was intuitively inclined to run. This observation also came in handy some time later.

My achievements in research in these early years were quite humdrum, but I was excited by several opportunities to bring psychology to bear on the real world. For these tasks, I teamed up with a colleague and friend, Ozer Schild. Together, we designed a training program for functionaries who were to introduce new immigrants from underdeveloped countries, such as Yemen, to modern farming practices (Kahneman and Schild, 1966). We also developed a training course for instructors in the flight school of the Air Force. Our faith in the usefulness of psychology was great, but we were also well aware of the difficulties of changing behavior without changing institutions and incentives. We may have done some good, and we certainly learned a lot.

I had the most satisfying Eureka experience of my career while attempting to teach flight instructors that praise is more effective than punishment for promoting skill-learning. When I had finished my enthusiastic speech, one of the most seasoned instructors in the audience raised his hand and made his own short speech, which began by conceding that positive reinforcement might be good for the birds, but went on to deny that it was optimal for flight cadets. He said, "On many occasions I have praised flight cadets for clean execution of some aerobatic maneuver, and in general when they try it again, they do worse. On the other hand, I have often screamed at cadets for bad execution, and in general they do better the next time. So please don't tell us that reinforcement works and punishment does not, because the opposite is the case." This was a joyous moment, in which I understood an important truth about the world: because we tend to reward others when they do well and punish them when they do badly, and because there is regression to the mean, it is part of the human condition that we are statistically punished for rewarding others and rewarded for punishing them. I immediately arranged a demonstration in which each participant tossed two coins at a target behind his back, without any feedback. We measured the distances from the target and could see that those who had done best the first time had mostly deteriorated on their second try, and vice versa. But I knew that this demonstration would not undo the effects of lifelong exposure to a perverse contingency.

My first experience of truly successful research came in 1965, when I was on sabbatical leave at the University of Michigan, where I had been invited by Jerry Blum, who had a lab in which volunteer participants performed various cognitive tasks while in the grip of powerful emotional states induced by hypnosis. Dilation of the pupil is one of the manifestations of emotional arousal, and I therefore became interested in the causes and consequences of changes of pupil size. Blum had a graduate student called Jackson Beatty. Using primitive equipment, Beatty and I made a real discovery: when people were exposed to a series of digits they had to remember, their pupils dilated steadily as they listened to the digits, and contracted steadily when they recited the series. A more difficult transformation task (adding 1 to each of a series of four digits) caused a much larger dilation of the pupil. We quickly published these results, and within a year had completed four articles, two of which appeared in *Science*. Mental effort remained the focus of my research during the subsequent year, which I spent at Harvard. During that year, I also heard a brilliant talk on experimental studies of attention by a star English psychologist named Anne Treisman, who would become my wife twelve years later. I was so impressed that I committed myself to write a chapter on attention for a Handbook in Cognitive Psychology. The Handbook was never published, and my chapter eventually became a rather ambitious book. The work on vision that I did that year was also more interesting than the work I had been doing in Jerusalem. When I returned home in 1967, I was, finally, a well-trained research psychologist.

#### The collaboration with Amos Tversky

From 1968 to 1969, I taught a graduate seminar on the applications of psychology to real-world problems. In what turned out to be a life-changing event, I asked my younger colleague Amos Tversky to tell the class about what was going on in his field of judgment and decision-making. Amos told us about the work of his former mentor, Ward Edwards, whose lab was using a research paradigm in which the subject is shown two bookbags filled with poker chips. The bags are said to differ in their composition (e.g., 70:30 or 30:70 white/red). One of them is randomly chosen, and the participant is given an opportunity to sample successively from it, and required to indicate after each trial the probability that it came from the predominantly red bag. Edwards had concluded from the results that people are "conservative Bayesians": they almost always adjust their confidence interval in the proper direction, but rarely far enough. A lively discussion developed around Amos's talk. The idea that people were conservative Bayesian did not seem to fit with the everyday observation of people commonly jumping to conclusions. It also appeared unlikely that the results obtained in the sequential sampling paradigm would extend to the situation, arguably more typical, in which sample evidence is delivered all at once. Finally, the label of 'conservative Bayesian' suggested the implausible image of a

process that gets the correct answer, then adulterates it with a bias. I learned recently that one of Amos's friends met him that day and heard about our conversation, which Amos described as having severely shaken his faith in the neo-Bayesian idea. I do remember that Amos and I decided to meet for lunch to discuss our hunches about the manner in which probabilities are "really" judged. There we exchanged personal accounts of our own recurrent errors of judgment in this domain, and decided to study the statistical intuitions of experts.

I spent the summer of 1969 doing research at the Applied Psychological Research Unit in Cambridge, England. Amos stopped there for a few days on his way to the United States. I had drafted a questionnaire on intuitions about sampling variability and statistical power, which was based largely on my personal experiences of incorrect research planning and unsuccessful replications. The questionnaire consisted of a set of questions, each of which could stand on its own - this was to be another attempt to do psychology with single questions. Amos went off and administered the questionnaire to participants at a meeting of the Mathematical Psychology Association, and a few weeks later we met in Jerusalem to look at the results and write a paper.

The experience was magical. I had enjoyed collaborative work before, but this was something different. Amos was often described by people who knew him as the smartest person they knew. He was also very funny, with an endless supply of jokes appropriate to every nuance of a situation. In his presence, I became funny as well, and the result was that we could spend hours of solid work in continuous mirth. The paper we wrote was deliberately humorous - we described a prevalent belief in the "law of small numbers," according to which the law of large numbers extends to small numbers as well. Although we never wrote another humorous paper, we continued to find amusement in our work - I have probably shared more than half of the laughs of my life with Amos.

And we were not just having fun. I quickly discovered that Amos had a remedy for everything I found difficult about writing. No wet-mush problem for him: he had an uncanny sense of direction. With him, movement was always forward. Progress might be slow, but each of the myriad of successive drafts that we produced was an improvement - this was not something I could take for granted when working on my own. Amos's work was always characterized by confidence and by a crisp elegance, and it was a joy to find those characteristics now attached to my ideas as well. As we were writing our first paper, I was conscious of how much better it was than the more hesitant piece I would have written by myself. I don't know exactly what it was that Amos found to like in our collaboration - we were not in the habit of trading compliments -but clearly he was also having a good time. We were a team, and we remained in that mode for well over a decade. The Nobel Prize was awarded for work that we produced during that period of intense collaboration.

At the beginning of our collaboration, we quickly established a rhythm that we maintained during all our years together. Amos was a night person, and I was a morning person. This made it natural for us to meet for lunch and a long afternoon together, and still have time to do our separate things. We spent hours each day, just talking. When Amos's first son Oren, then fifteen months old, was told that his father was at work, he volunteered the comment "Aba talk Danny." We were not only working, of course - we talked of everything under the sun, and got to know each other's mind almost as well as our own. We could (and often did) finish each other's sentences and complete the joke that the other had wanted to tell, but somehow we also kept surprising each other.

We did almost all the work on our joint projects while physically together, including the drafting of questionnaires and papers. And we avoided any explicit division of labor. Our principle was to discuss every disagreement until it had been resolved to mutual satisfaction, and we had tie-breaking rules for only two topics: whether or not an item should be included in the list of references (Amos had the casting vote), and who should resolve any issue of English grammar (my dominion). We did not initially have a concept of a senior author. We tossed a coin to determine the order of authorship of our first paper, and alternated from then on until the pattern of our collaboration changed in the 1980s.

One consequence of this mode of work was that all our ideas were jointly owned. Our interactions were so frequent and so intense that there was never much point in distinguishing between the discussions that primed an idea, the act of uttering it, and the subsequent elaboration of it. I believe that many scholars have had the experience of discovering that they had expressed (sometimes even published) an idea long before they really understood its significance. It takes time to appreciate and develop a new thought. Some of the greatest joys of our collaboration-and probably much of its success - came from our ability to elaborate each other's nascent thoughts: if I expressed a half-formed idea, I knew that Amos would be there to understand it, probably more clearly than I did, and that if it had merit he would see it. Like most people, I am somewhat cautious about exposing tentative thoughts to others - I must first make sure that they are not idiotic. In the best years of the collaboration, this caution was completely absent. The mutual trust and the complete lack of defensiveness that we achieved were particularly remarkable because both of us - Amos even more than I - were known to be severe critics. Our magic worked only when we were by ourselves. We soon learned that joint collaboration with any third party should be avoided, because we became competitive in a threesome.

Amos and I shared the wonder of together owning a goose that could lay golden eggs - a

joint mind that was better than our separate minds. The statistical record confirms that our joint work was superior, or at least more influential, than the work we did individually (Laibson and Zeckhauser, 1998). Amos and I published eight journal articles during our peak years (1971-1981), of which five had been cited more than a thousand times by the end of 2002. Of our separate works, which in total number about 200, only Amos' theory of similarity (Tversky, 1977) and my book on attention (Kahneman, 1973) exceeded that threshold. The special style of our collaborative work was recognized early by a referee of our first theoretical paper (on representativeness), who caused it to be rejected by *Psychological Review*. The eminent psychologist who wrote that review - his anonymity was betrayed years later - pointed out that he was familiar with the separate lines of work that Amos and I had been pursuing, and considered both quite respectable. However, he added the unusual remark that we seemed to bring out the worst in each other, and certainly should not collaborate. He found most objectionable our method of using multiple single questions as evidence - and he was quite wrong there as well.

#### The Science '74 article and the rationality debate

From 1971 to 1972, Amos and I were at the Oregon Research Institute (ORI) in Eugene, a year that was by far the most productive of my life. We did a considerable amount of research and writing on the availability heuristic, on the psychology of prediction, and on the phenomena of anchoring and overconfidence - thereby fully earning the label "dynamic duo" that our colleagues attached to us. Working evenings and nights, I also completely rewrote my book on *Attention and Effort*, which went to the publisher that year, and remains my most significant independent contribution to psychology.

At ORI, I came into contact for the first time with an exciting community of researchers that Amos had known since his student days at Michigan: Paul Slovic, Sarah Lichtenstein, and Robyn Dawes. Lewis Goldberg was also there, and I learned much from his work on clinical and actuarial judgment, and from Paul Hoffman's ideas about paramorphic modeling. ORI was one of the major centers of judgment research, and I had the occasion to meet quite a few of the significant figures of the field when they came visiting, Ken Hammond among them.

Some time after our return from Eugene, Amos and I settled down to review what we had learned about three heuristics of judgment (representativeness, availability, and anchoring) and about a list of a dozen biases associated with these heuristics. We spent a delightful year in which we did little but work on a single article. On our usual schedule of spending afternoons together, a day in which we advanced by a sentence or two was considered quite productive. Our enjoyment of the process gave us unlimited patience, and we wrote as if the precise choice of every word were a matter of great moment.

We published the article in *Science* because we thought that the prevalence of systematic biases in intuitive assessments and predictions could possibly be of interest to scholars outside psychology. This interest, however, could not be taken for granted, as I learned in an encounter with a well-known American philosopher at a party in Jerusalem. Mutual friends had encouraged us to talk about the research that Amos and I were doing, but almost as soon as I began my story he turned away, saying, "I am not really interested in the psychology of stupidity."

The *Science* article turned out to be a rarity: an empirical psychological article that (some) philosophers and (a few) economists could and did take seriously. What was it that made readers of the article more willing to listen than the philosopher at the party? I attribute the unusual attention at least as much to the medium as to the message. Amos and I had continued to practice the psychology of single questions, and the *Science* article - like others we wrote - incorporated questions that were cited verbatim in the text. These questions, I believe, personally engaged the readers and convinced them that we were concerned not with the stupidity of Joe Public but with a much more interesting issue: the susceptibility to erroneous intuitions of intelligent, sophisticated, and perceptive individuals such as themselves. Whatever the reason, the article soon became a standard reference as an attack on the rational-agent model, and it spawned a large literature in cognitive science, philosophy, and psychology. We had not anticipated that outcome.

I realized only recently how fortunate we were not to have aimed deliberately at the large target we happened to hit. If we had intended the article as a challenge to the rational model, we would have written it differently, and the challenge would have been less effective. An essay on rationality would have required a definition of that concept, a treatment of boundary conditions for the occurrence of biases, and a discussion of many other topics about which we had nothing of interest to say. The result would have been less crisp, less provocative, and ultimately less defensible. As it was, we offered a progress report on our study of judgment under uncertainty, which included much solid evidence. All inferences about human rationality were drawn by the readers themselves.

The conclusions that readers drew were often too strong, mostly because existential quantifiers, as they are prone to do, disappeared in the transmission. Whereas we had shown that (some, not all) judgments about uncertain events are mediated by heuristics, which (sometimes, not always) produce predictable biases, we were often read as having claimed that people cannot think straight. The fact that men had walked on the moon was used more than once as an argument against our position. Because our treatment was mistakenly taken to be inclusive, our silences became significant. For example, the fact that we had written nothing about the role of social factors in judgment was taken as an indication that we thought these factors were unimportant. I suppose that we could have

prevented at least some of the misunderstandings, but the cost of doing so would have been too high.

The interpretation of our work as a broad attack on human rationality - rather than as a critique of the rational-agent model - attracted much opposition, some quite harsh and dismissive. Some of the critiques were normative, arguing that we compared judgments to inappropriate normative standards (Cohen, 1981; Gigerenzer, 1991, 1996). We were also accused of spreading a tendentious and misleading message that exaggerated the flaws of human cognition (Lopes, 1991, and many others). The idea of systematic bias was rejected as unsound on evolutionary grounds (Cosmides & Tooby, 1996). Some authors dismissed the research as a collection of artificial puzzles designed to fool undergraduates. Numerous experiments were conducted over the years, to show that cognitive illusions could "be made to disappear" and that heuristics had been invented to explain "biases that do not exist" (Gigerenzer, 1991). After participating in a few published skirmishes in the early 80's, Amos and I adopted a policy of not criticizing the critiques of our work, although we eventually felt compelled to make an exception (Kahneman and Tversky, 1996).

A young colleague and I recently reviewed the experimental literature, and concluded that the empirical controversy about the reality of cognitive illusions dissolves when viewed in the perspective of a dual-process model (Kahneman and Frederick, 2002). The essence of such a model is that judgments can be produced in two ways (and in various mixtures of the two): a rapid, associative, automatic, and effortless intuitive process (sometimes called System 1), and a slower, rule-governed, deliberate and effortful process (System 2) (Sloman, 1996; Stanovich and West, 1999). System 2 "knows" some of the rules that intuitive reasoning is prone to violate, and sometimes intervenes to correct or replace erroneous intuitive judgments. Thus, errors of intuition occur when two conditions are satisfied: System 1 generates the error and System 2 fails to correct. In this view, the experiments in which cognitive illusions were "made to disappear" did so by facilitating the corrective operations of System 2. They tell us little about the intuitive judgments that are suppressed.

If the controversy is so simply resolved, why was it not resolved in 1971, or in 1974? The answer that Frederick and I proposed refers to the conversational context in which the early work was done:

A comprehensive psychology of intuitive judgment cannot ignore such controlled thinking, because intuition can be overridden or corrected by self-critical operations, and because intuitive answers are not always available. But this sensible position seemed irrelevant in the early days of research on judgment heuristics. The authors of the "law of small numbers" saw no need to examine correct statistical reasoning. They believed that including easy questions in the design would insult the participants and bore the readers. More generally, the early studies of heuristics and biases displayed little interest in the conditions under which intuitive reasoning is preempted or overridden - controlled reasoning leading to correct answers was seen as a default case that needed no explaining. A lack of concern for boundary conditions is typical of "young" research programs, which naturally focus on demonstrating new and unexpected effects, not on making them disappear. (Kahneman and Frederick, 2002, p. 50).

What happened, I suppose, is that because the 1974 paper was influential it altered the context in which it was read in subsequent years. Its being misunderstood was a direct consequence of its being taken seriously. I wonder how often this occurs.

Amos and I always dismissed the criticism that our focus on biases reflected a generally pessimistic view of the human mind. We argued that this criticism confuses the medium of bias research with a message about rationality. This confusion was indeed common. In one of our demonstrations of the availability heuristic, for example, we asked respondents to compare the frequency with which some letters appeared in the first and in the third position in words. We selected letters that in fact appeared more frequently in the third position, and showed that even for these letters the first position was judged more frequent, as would be predicted on the idea that it is easier to search through a mental dictionary by the first letter. The experiment was used by some critics as an example of our own confirmation bias, because we had demonstrated availability only in cases in which this heuristic led to bias. But this criticism assumes that our aim was to demonstrate biases, and misses the point of what we were trying to do. Our aim was to show that the availability heuristic controls frequency estimates *even when that heuristic leads to error* - an argument that cannot be made when the heuristic leads to correct responses, as it often does.

There is no denying, however, that the name of our method and approach created a strong association between heuristics and biases, and thereby contributed to giving heuristics a bad name, which we did not intend. I recently came to realize that the association of heuristics and biases has affected me as well. In the course of an exchange of messages with Ralph Hertwig (no fan of heuristics and biases), I noticed that the phrase "judging by representativeness" was in my mind a label for a cluster of errors in intuitive statistical judgment. Judging probability by representativeness is indeed associated with systematic errors. But a large component of the process is the judgment of representativeness, and that judgment is often subtle and highly skilled. The feat of the



master chess player who instantly recognizes a position as "white mates in three" is an instance of judgment of representativeness. The undergraduate who instantly recognizes that enjoyment of puns is more representative of a computer scientist than of an accountant is also exhibiting high skill in a social and cultural judgment. My long-standing failure to associate specific benefits to the concept of representativeness was a revealing mistake.

What did I learn from the controversy about heuristics and biases? Like most protagonists in debates, I have few memories of having changed my mind under adversarial pressure, but I have certainly learned more than I know. For example, I am now quick to reject any description of our work as demonstrating human irrationality. When the occasion arises, I carefully explain that research on heuristics and biases only refutes an unrealistic conception of rationality, which identifies it as comprehensive coherence. Was I always so careful? Probably not. In my current view, the study of judgment biases requires attention to the interplay between intuitive and reflective thinking, which sometimes allows biased judgments and sometimes overrides or corrects them. Was this always as clear to me as it is now? Probably not. Finally, I am now very impressed by the observation I mentioned earlier, that the most highly skilled cognitive performances are intuitive, and that many complex judgments share the speed, confidence and accuracy of routine perception. This observation is not new to me, but did it always loom as large in my views as it now does? Almost certainly not.

As my obvious struggle with this topic reveals, I thoroughly dislike controversies where it is clear that no minds will be changed. I feel diminished by losing my objectivity when in point-scoring mode, and downright humiliated when I get angry. Indeed, my phobia for professional anger is such that I have allowed myself for many years the luxury of refusing to referee papers that might arouse that emotion: If the tone is snide, or the review of the facts more tendentious than normal, I return the paper back to the editor without commenting on it. I consider myself fortunate not to have had too many of the nasty experiences of professional quarrels, and am grateful for the occasional encounters with open minds across lines of sharp debate (Ayton, 1998; Klein, 2000).

#### Prospect theory

After the publication of our paper on judgment in *Science* in 1974, Amos suggested that we study decision-making together. This was a field in which he was already an established star, and about which I knew very little. For an introduction, he suggested that I read the relevant chapters of the text "Mathematical Psychology," of which he was a co-author (Coombs, Dawes and Tversky, 1970). Utility theory and the paradoxes of Allais and Ellsberg were discussed in the book, along with some of the classic experiments in which major figures in the field had joined in an effort to measure the utility function for money by eliciting choices between simple gambles.

I learned from the book that the name of the game was the construction of a theory that would explain Allais's paradox parsimoniously. As psychological questions go, this was not a difficult one, because Allais's famous problems are, in effect, an elegant way to demonstrate that the subjective response to probability is not linear. The subjective non-linearity is obvious: the difference between probabilities of .10 and .11 is clearly less impressive than the difference between 0 and .01, or between .99 and 1.00. The difficulty and the paradox exist only for decision theorists, because the non-linear response to probability produces preferences that violate compelling axioms of rational choice and are therefore incompatible with standard expected utility theory. The natural response of a decision theorist to the Allais paradox, certainly in 1975 and probably even today, would be to search for a new set of axioms that have normative appeal and yet permit the non-linearity. The natural response of psychologists was to set aside the issue of rationality and to develop a descriptive theory of the preferences that people actually have, regardless of whether or not these preferences can be justified.

The task we set for ourselves was to account for observed preferences in the quaintly restricted universe within which the debate about the theory of choice has traditionally been conducted: monetary gambles with few outcomes (all positive), and definite probabilities. This was an empirical question, and data were needed. Amos and I solved the data collection problem with a method that was both efficient and pleasant. We spent our hours together inventing interesting choices and examining our preferences. If we agreed on the same choice we provisionally assumed that other people would also accept it, and we went on to explore its theoretical implications. This unusual method enabled us to move quickly, and we constructed and discarded models at a dizzying rate. I have a distinct memory of a model that was numbered 37, but cannot vouch for the accuracy of our count.

As was the case in our work on judgment, our central insights were acquired early and, as was the case in our work on judgment, we spent a vast amount of time and effort before publishing a paper that summarized those insights (Kahneman and Tversky, 1979). The first insight came as a result of my naiveté. When reading the mathematical psychology textbook, I was puzzled by the fact that all the choice problems were described in terms of gains and losses (actually, almost always gains), whereas the utility functions that were supposed to explain the choices were drawn with wealth as the abscissa. This seemed unnatural, and psychologically unlikely. We immediately decided to adopt changes and/or differences as carriers of utility. We had no inkling that this obvious move was truly fundamental, or that it would open the path to behavioral economics. Harry Markowitz, who won the Nobel Prize in economics in 1990, had proposed changes of wealth as carriers of utility in 1952, but he did not take this idea very far.

The shifts from wealth to changes of wealth as carriers of utility is significant because of a property of preferences that we later labeled *loss-aversion*: the response to losses is consistently much more intense than the response to corresponding gains, with a sharp kink in the value function at the reference point. Loss aversion is manifest in the extraordinary reluctance to accept risk that is observed when people are offered a gamble on the toss of a coin: most will reject a gamble in which they might lose \$20, unless they are offered more than \$40 if they win. The concept of loss aversion was, I believe, our most useful contribution to the study of decision making. The asymmetry between gains and losses solves quite a few puzzles, including the widely noted and economically irrational distinction that people draw between opportunity costs and 'real' losses. Loss aversion also helps explain why real-estate markets dry up for long periods when prices are down, and it contributes to the explanation of a widespread bias favoring the status quo in decision making. Finally, the asymmetric consideration of gains and losses extends to the domain of moral intuitions, in which imposing losses and failing to share gains are evaluated quite differently. But of course, none of that was visible to Amos and me when we first decided to assume a kinked value function - we needed that kink to account for choices between gambles.

Another set of early insights came when Amos suggested that we flip the signs of outcomes in the problems we had been considering. The result was exciting. We immediately detected a remarkable pattern, which we called "reflection": changing the signs of all outcomes in a pair of gambles almost always caused the preference to change from risk averse to risk seeking, or viceversa. For example, we both preferred a sure gain of \$900 over a .9 probability of gaining \$1,000 (or nothing), but we preferred a gamble with a .9 probability of losing \$1,000 over a sure loss of \$900. We were not the first to observe this pattern. Raiffa (1968) and Williams (1966) knew about the prevalence of risk-seeking in the negative domain. But ours was apparently the first serious attempt to make something of it.

We soon had a draft of a theory of risky choice, which we called "value theory" and presented at a conference in the spring of 1975. We then spent about three years polishing it, until we were ready to submit the article for publication. Our effort during those years was divided between the tasks of exploring interesting implications of our theoretical formulation and developing answers to all plausible objections. To amuse ourselves, we invented the specter of an ambitious graduate student looking for flaws, and we labored to make that student's task as thankless as possible. The most novel idea of prospect theory occurred to us in that defensive context. It came quite late, as we were preparing the final version of the paper. We were concerned with the fact that a straightforward application of our model implied that the value of the prospect (\$100, .01; \$100, .01) is larger than the value of (\$100, .02). The prediction is wrong, of course, because most decision makers will spontaneously transform the former prospect into the latter and treat them as equivalent in subsequent operations of evaluation and choice. To eliminate the problem we proposed that decision-makers, prior to evaluating the prospects, perform an editing operation that collects similar outcomes and adds their probabilities. We went on to propose several other editing operations that provided an explicit and psychologically plausible defense against a variety of superficial counter-examples to the core of the theory. We had succeeded in making life quite difficult for that pedantic graduate student. But we had also made a truly significant advance, by making it explicit that the objects of choice are mental representations, not objective states of the world. This was a large step toward the development of a concept of framing, and eventually toward a new critique of the model of the rational agent.

When we were ready to submit the work for publication, we deliberately chose a meaningless name for our theory: "prospect theory." We reasoned that if the theory ever became well known, having a distinctive label would be an advantage. This was probably wise.

I looked at the 1975 draft recently, and was struck by how similar it is to the paper that was eventually published, and also by how different the two papers are. Most of the key ideas, most of the key examples, and much of the wording were there in the early draft. But that draft lacks the authority that was gained during the years that we spent anticipating objections. "Value theory" would not have survived the close scrutiny that a significant article ultimately gets from generations of scholars and students, who only are obnoxious if you give them a chance.

We published the paper in *Econometrica*. The choice of venue turned out to be important; the identical paper, published in *Psychological Review*, would likely have had little impact on economics. But our decision was not guided by a wish to influence economics. *Econometrica* just happened to be the journal where the best papers on decision-making to date had been published, and we were aspiring to be in that company.

And there was another way in which the impact of prospect theory depended crucially on the medium, as well as the message. Prospect theory was a formal theory, and its formal nature was the key to the impact it had in economics. Every discipline of social science, I believe, has some ritual tests of competence, which must be passed before a piece of work is considered worthy of attention. Such tests are necessary to prevent information overload, and they are also important aspects of the tribal life of the disciplines. In particular, they allow insiders to ignore just about anything that is done by members of other tribes, and to feel no scholarly guilt about doing so. To serve this screening function

efficiently, the experiments usually focus on some aspect of form or method, and have little or nothing to do with substance. Prospect theory passed such a test in economics, and its observations became a legitimate (though optional) part of the scholarly discourse in that discipline. It is a strange and rather arbitrary process that selects some pieces of scientific writing for relatively enduring fame while committing most of what is published to almost immediate oblivion.

#### Framing and mental accounting

Amos and I completed prospect theory during the academic year of 1977 to 1978, which I spent at the Center for Advanced Studies at Stanford, while he was visiting the psychology department there. Around that time, we began work on our next project, which became the study of framing. This was also the year in which the second most important professional friendship in my life - with Richard Thaler - had its start.

A framing effect is demonstrated by constructing two transparently equivalent versions of a given problem, which nevertheless yield predictably different choices. The standard example of a framing problem, which was developed quite early, is the 'lives saved, lives lost' question, which offers a choice between two public-health programs proposed to deal with an epidemic that is threatening 600 lives: one program will save 200 lives, the other has a 1/3 chance of saving all 600 lives and a 2/3 chance of saving none. In this version, people prefer the program that will save 200 lives for sure. In the second version, one program will result in 400 deaths, the other has a 2/3 chance of 600 deaths and a 1/3 chance of no deaths. In this formulation most people prefer the gamble. If the same respondents are given the two problems on separate occasions, many give incompatible responses. When confronted with their inconsistency, people are quite embarrassed. They are also quite helpless to resolve the inconsistency, because there are no moral intuitions to guide a choice between different sizes of a surviving population.

Amos and I began creating pairs of problems that revealed framing effects while working on prospect theory. We used them to show sensitivity to gains and losses (as in the lives example), and to illustrate the inadequacy of a formulation in which the only relevant outcomes are final states. In that article, we also showed that a single-stage gamble could be rearranged as a two-stage gamble in a manner that left the bottom-line probabilities and outcomes unchanged but reversed preferences. Later, we developed examples in which respondents are asked to make simultaneous choices in two problems, A and B. One of the problems involves gains and elicits a risk-averse choice; the other problem involves losses and elicits risk-seeking. A majority of respondents made both these choices. However, the problems were constructed so that the combination of choices that people made was actually dominated by the combination of the options they had rejected.

These are not parlor-game demonstrations of human stupidity. The ease with which framing effects can be demonstrated reveals a fundamental limitation of the human mind. In a rational-agent model, the agent's mind functions just as she would like it to function. Framing effects violate that basic requirement: the respondents who exhibit susceptibility to framing effects wish their minds were able to avoid them. We were able to conceive of only two kinds of mind that would avoid framing effects: (1) If responses to all outcomes and probabilities were strictly linear, the procedures that we used to produce framing effects would fail. (2) If individuals maintained a single canonical and all-inclusive view of their outcomes, truly equivalent problems would be treated equivalently. Both conditions are obviously impossible. Framing effects violate a basic requirement of rationality which we called invariance (Kahneman and Tversky, 1984) and Arrow (1982) called extensionality. It took us a long time and several iterations to develop a forceful statement of this contribution to the rationality debate, which we presented several years after our framing paper (Tversky and Kahneman, 1986).

Another advance that we made in our first framing article was the inclusion of riskless choice problems among our demonstrations of framing. In making that move, we had help from a new friend. Richard Thaler was a young economist, blessed with a sharp and irreverent mind. While still in graduate school, he had trained his ironic eye on his own discipline and had collected a set of pithy anecdotes demonstrating obvious failures of basic tenets of economic theory in the behavior of people in general - and of his very conservative professors in Rochester in particular. One key observation was the endowment effect, which Dick illustrated with the example of the owner of a bottle of old wine, who would refuse to sell it for \$200 but would not pay as much as \$100 to replace it if it broke. Sometime in 1976, a copy of the 1975 draft of prospect theory got into Dick's hands, and that event made a significant difference to our lives. Dick realized that the endowment effect, which is a genuine puzzle in the context of standard economic theory, is readily explained by two assumptions derived from prospect theory. First, the carriers of utility are not states (owning or not owning the wine), but changes - getting the wine or giving it up. And giving up is weighted more than getting, by loss aversion. When Dick learned that Amos and I would be in Stanford in 1977/8, he secured a visiting appointment at the Stanford branch of the National Bureau of Economic Research, which is located on the same hill as the Center for Advanced Studies. We soon became friends, and have ever since had a considerable influence on each other's thinking.

The endowment effect was not the only thing we learned from Dick. He had also developed a list of phenomena of what we now call "mental accounting." Mental accounting describes how people violate rationality by failing to maintain a comprehensive view of outcomes, and by failing to treat money as fungible. Dick showed how people segregate their decisions into separate accounts, then struggle to keep each of these accounts in the black. One of his compelling examples was the couple who drove through a blizzard to a

basketball game because they were paid for the tickets, though they would have stayed at home if the tickets had been free. As this example illustrates, Dick had independently developed the skill of doing "one-question economics." He inspired me to invent another story, in which a person who comes to the theater realizes that he has lost his ticket (in one version), or an amount of cash equal to the ticket value (in another version). People report that they would be very likely still to buy a ticket if they had lost the cash, presumably because the loss has been charged to general revenue. On the other hand, they describe themselves as quite likely to go home if they have lost an already purchased ticket, presumably because they do not want to pay twice to see the same show.

#### Behavioral economics

Our interaction with Thaler eventually proved to be more fruitful than we could have imagined at the time, and it was a major factor in my receiving the Nobel Prize. The committee cited me "for having integrated insights from psychological research into economic science ....". Although I do not wish to renounce any credit for my contribution, I should say that in my view the work of integration was actually done mostly by Thaler and the group of young economists that quickly began to form around him, starting with Colin Camerer and George Loewenstein, and followed by the likes of Matthew Rabin, David Laibson, Terry Odean, and Sendhil Mullainathan. Amos and I provided quite a few of the initial ideas that were eventually integrated into the thinking of some economists, and prospect theory undoubtedly afforded some legitimacy to the enterprise of drawing on psychology as a source of realistic assumptions about economic agents. But the founding text of behavioral economics was the first article in which Thaler (1980) presented a series of vignettes that challenged fundamental tenets of consumer theory. And the respectability that behavioral economics now enjoys within the discipline was secured, I believe, by some important discoveries Dick made in what is now called behavioral finance, and by the series of "Anomalies" columns that he published in every issue of the *Journal of Economic Perspectives* from 1987 to 1990, and has continued to write occasionally since that time.

In 1982, Amos and I attended a meeting of the Cognitive Science Society in Rochester, where we had a drink with Eric Wanner, a psychologist who was then vice-president of the Sloan Foundation. Eric told us that he was interested in promoting the integration of psychology and economics, and asked for our advice on ways to go about it. I have a clear memory of the answer we gave him. We thought that there was no way to "spend a lot of money honestly" on such a project, because interest in interdisciplinary work could not be coerced. We also thought that it was pointless to encourage psychologists to make themselves heard by economists, but that it could be useful to encourage and support the few economists who were interested in listening. Thaler's name surely came up. Soon after that conversation, Wanner became the president of the Russell Sage Foundation, and he brought the psychology/economics project with him. The first grant that he made in that program was for Dick Thaler to spend an academic year (1984-85) visiting me at the University of British Columbia, in Vancouver.

That year was one of the best in my career. We worked as a trio that also included the economist Jack Knetsch, with whom I had already started constructing surveys on a variety of issues, including valuation of the environment and public views about fairness in the marketplace. Jack had done experimental studies of the endowment effect and had seen the implications of that effect for the Coase theorem and for issues of environmental policy. We made a very good team: Jack's wisdom and imperturbable calm withstood the stress of Dick's boisterous temperament and of my perfectionist anxieties and intellectual restlessness.

We did a lot together that year. We conducted a series of market experiments involving real goods (the "mugs" studies), which eventually became a standard in that literature (Kahneman, Knetsch and Thaler, 1990). We also conducted multiple surveys in which we used experimentally varied vignettes to identify the rules of fairness that the public would apply to merchants, landlords, and employers (Kahneman, Knetsch and Thaler, 1986a). Our central observation was that in many contexts the existing situation (e.g., price, rent, or wage) defines a "reference transaction," to which the transactor (consumer, tenant, and employee) has an entitlement - the violation of such entitlements is considered unfair and may evoke retaliation. For example, cutting the wages of an employee merely because he could be replaced by someone who would accept a lower wage is unfair, although paying a lower wage to the replacement of an employee who quit is entirely acceptable. We submitted the paper to the *American Economic Review* and were utterly surprised by the outcome: the paper was accepted without revision. Luckily for us, the editor had asked two economists quite open to our approach to review the paper. We later learned that one of the referees was [George Akerlof](#) and the other was Alan Olmstead, who had studied the failures of markets to clear during an acute gas shortage.

One question that arose during this research was whether people would be willing to pay something to punish another agent who treated them "unfairly", and in some circumstances would share a windfall with a stranger in an effort to be "fair". We decided to investigate these ideas using experiments for real stakes. The games that we invented for this purpose have become known as the ultimatum game and the dictator game. Alas, while writing up our second paper on fairness (Kahneman, Knetsch and Thaler, 1986b) we learned that we had been scooped on the ultimatum game by Werner Guth and his colleagues, who had published experiments using the same design a few years earlier. I remember being quite crestfallen when I learned this. I would have been even more depressed if I had known how important the ultimatum game would eventually become.

Most of the economics I know I learned that year, from Jack and Dick, my two willing teachers, and from what was in fact my first experience of communicating across tribal boundaries. I was also much impressed by an experimental game that Dick Thaler, James Brander, and I invented and called the  $N^*$  game. The game is played by a group of, say, fifteen people. On each trial, a number  $0 < N^* < 15$  is announced. The participants then make simultaneous choices of whether or not to "enter." Those who decide to enter announce their choice simultaneously. The payoff to the  $N$  entrants depends on their number, according to the following formula:  $\$.25(N^* - N)$ . We played the game a few times, once with the faculty of the psychology department at U.B.C. The results, although not surprising to an economist, struck me as magical. Within very few trials, a pattern emerged in which the number of entrants,  $N$ , was within 1 or 2 of  $N^*$ , with no obvious systematic tendency to be higher or lower than  $N^*$ . The group was doing the right thing collectively, although conversations with the participants and the obvious statistical analyses did not reveal any consistent strategies that made sense. It took me some time to realize that the magic we were observing was an equilibrium: the pattern we saw existed because no other pattern could be sustained. This idea had not been in my intellectual bag of tools. We never formally published the  $N^*$  game - I described it informally in Kahneman (1987) - but it has been taken up by others (Erev & Rapoport, 1998).

That was the closest my research ever came to core economics, and since that time I have been mostly cheering Thaler and behavioral economics from the sidelines. There has been much to cheer about. As a mark of the progress that has been made, I recall a seminar in psychology and economics that I co-taught with George Akerlof, after Anne Treisman and I had moved from the University of British Columbia to Berkeley in 1986. I remember being struck by the reverence with which the rationality assumption was treated even by a free thinker such as George, and also by his frequent warnings to the students that they should not let themselves be seduced by the material we were presenting, lest their careers be permanently damaged. His advice to them was to stick to what he called "meat-and-potatoes economics," at least until their careers were secure. This opinion was quite common at the time. When Matthew Rabin joined the Berkeley economics department as a young assistant professor and chose to immerse himself in psychology, many considered the move professional suicide. Some fifteen years later, Rabin had earned the Clark medal, and George Akerlof had delivered a Nobel lecture entitled "Behavioral Macroeconomics."

Eric Wanner and the Russell Sage Foundation continued to support behavioral economics over the years. I was instrumental in the idea of using some of that support to set up a summer school for graduate students and young faculty in that field, and I helped Dick Thaler and Colin Camerer organize the first one, in 1994. When the fifth summer school convened in 2002, David Laibson, who had been a participant in 1994, was tenured at Harvard and was one of the three organizers. Terrance Odean and Sendhil Mullainathan, who had also participated as students, came back to lecture as successful researchers with positions in two of the best universities in the world. It was a remarkable experience to hear Matthew Rabin teach a set of guidelines for developing theories in behavioral economics - including the suggestion that the standard economic model should be a special case of the more complex and general models that were to be constructed. We had come a long way.

Although behavioral economics has enjoyed much more rapid progress and gained more respectability in economics than appeared possible fifteen years ago, it is still a minority approach and its influence on most fields of economics is negligible. Many economists believe that it is a passing fad, and some hope that it will be. The future may prove them right. But many bright young economists are now betting their careers on the expectation that the current trend will last. And such expectations have a way of being self-fulfilling.

#### Later years

Anne Treisman and I married and moved together to U.B.C. in 1978, and Amos and Barbara Tversky settled in Stanford that year. Amos and I were then at the peak of our joint game, and completely committed to our collaboration. For a few years, we managed to maintain it, by spending every second weekend together and by placing multiple phone calls each day, some lasting several hours. We completed the study of framing in that mode, as well as a study of the 'conjunction fallacy' in judgment (Tversky and Kahneman, 1983). But eventually the goose that had laid the golden eggs languished, and our collaboration tapered off. Although this outcome now appears inevitable, it came as a painful surprise to us. We had completely failed to appreciate how critically our successful interaction had depended on our being together at the birth of every significant idea, on our rejection of any formal division of labor, and on the infinite patience that became a luxury when we could meet only periodically. We struggled for years to revive the magic we had lost, but in vain.

We were again trying when Amos died. When he learned in the early months of 1996 that he had only a few months to live, we decided to edit a joint book on decision-making that would cover some of the progress that had been made since we had started working together on the topic more than twenty years before (Kahneman and Tversky, 2000). We planned an ambitious preface as a joint project, but I think we both knew from the beginning that we would not be granted enough time to complete it. The preface I wrote alone was probably my most painful writing experience.

During the intervening years, of course, we had continued to work, sometimes together

sometimes with other collaborators. Amos took the lead in our most important joint piece, an extension of prospect theory to the multiple outcome case in the spirit of rank-dependent models. He also carried out spectacular studies of the role of argument and conflict in decision-making, in collaborations with Eldar Shafir and with Itamar Simonson, as well as influential work on violations of procedural invariance in collaborations with Shmuel Sattath and with Paul Slovic. He engaged in a deep exploration of the mathematical structure of decision theories with Peter Wakker. And, in his last years, Amos was absorbed in the development of support theory, a general approach to thinking under uncertainty that his students have continued to explore. These are only his major programmatic research efforts in the field of decision-making - he did much more.

I, too, kept busy, and also kept moving. Anne Treisman and I moved to UC Berkeley in 1986, and from there to Princeton in 1993, where I happily took a split appointment that located me part-time in the Woodrow Wilson School of Public Affairs. Moving East also made it easier to maintain frequent contacts with friends, children and adored grandchildren in Israel.

Over the years I enjoyed productive collaborations with Dale Miller in the development of a theory of counterfactual thinking (Kahneman and Miller, 1986), and with Anne Treisman, in studies of visual attention and object perception. In addition to the work on fairness and on the endowment effect that we did with Dick Thaler, Jack Knetsch and I carried out studies of the valuation of public goods that became quite controversial and had a great influence on my own thinking. Further studies of that problem with Ilana Ritov eventually led to the idea that the translation of attitudes into dollars involves the almost arbitrary choice of a scale factor, leading some people who have quite similar values to state very different values of their willingness to pay, for no good reason (Kahneman, Ritov and Schkade, 1999). With David Schkade and the famous jurist Cass Sunstein I extended this idea into a program of research on arbitrariness in punitive damage decisions, which may yet have some influence on policy (Sunstein, Kahneman, Schkade and Ritov, 2002).

The focus of my research for the past fifteen years has been the study of various aspects of experienced utility - the measure of the utility of outcomes as people actually live them. The concept of utility in which I am interested was the one that Bentham and Edgeworth had in mind. However, experienced utility largely disappeared from economic discourse in the twentieth century, in favor of a notion that I call decision utility, which is inferred from choices and used to explain choices. The distinction could be of little relevance for fully rational agents, who presumably maximize experienced utility as well as decision utility. But if rationality cannot be assumed, the quality of consequences becomes worth measuring and the maximization of experienced utility becomes a testable proposition. Indeed, my colleagues and I have carried out experiments in which this proposition was falsified. These experiments exploit a simple rule that governs the assignment of remembered utility to past episodes in which an agent is passively exposed to a pleasant or unpleasant experience, such as watching a horrible film or an amusing one (Frederickson and Kahneman, 1993), or undergoing a colonoscopy (Redelmeier and Kahneman, 1993). Remembered utility turns out to be determined largely by the peak intensity of the pleasure or discomfort experienced during the episode, and by the intensity of pleasure or discomfort when the episode ended. The duration of the episode has almost no effect on its remembered utility. In accord with this rule, an episode of 60 seconds during which one hand is immersed in painfully cold water will leave a more aversive memory than a longer episode, in which the same 60 seconds are followed by another 30 seconds during which the temperature rises slightly. Although the extra 30 seconds are painful, they provide an improved end. When experimental participants are exposed to the two episodes, then given a choice of which to repeat, most choose the longer one (Kahneman, Fredrickson, Schreiber and Redelmeier, 1993). In these and in other experiments of the same kind (Schreiber and Kahneman, 2000), people make wrong choices between experiences to which they may be exposed, because they are systematically wrong about their affective memories. Our evidence contradicts the standard rational model, which does not distinguish between experienced utility and decision utility. I have presented it as a new type of challenge to the assumption of rationality (Kahneman, 1994).

Most of my empirical work in recent years has been done in collaboration with my friend David Schkade. The current topic of our research is a study of well-being that builds on my previous research on experienced utility. We have assembled a multi-disciplinary team for an attempt to develop tools for measuring welfare, with the design specification that economists should be willing to take the measurements seriously.

Another major effort went into an essay that attempted to update the notion of judgment heuristics. That work was done in close collaboration with a young colleague, Shane Frederick. In the pains we took in the choice of every word it came close to matching my experiences with Amos (Kahneman and Frederick, 2002). My Nobel lecture is an extension of that essay.

One line of work that I hope may become influential is the development of a procedure of *adversarial collaboration*, which I have championed as a substitute for the format of critique-reply-rejoinder in which debates are currently conducted in the social sciences.<sup>1</sup> Both as a participant and as a reader I have been appalled by the absurdly adversarial nature of these exchanges, in which hardly anyone ever admits an error or acknowledges learning anything from the other. Adversarial collaboration involves a good-faith effort to conduct debates by carrying out joint research - in some cases there may be a need for

an agreed arbiter to lead the project and collect the data. Because there is no expectation of the contestants reaching complete agreement at the end of the exercise, adversarial collaborations will usually lead to an unusual type of joint publication, in which disagreements are laid out as part of a jointly authored paper. I have had three adversarial collaborations, with Tom Gilovich and Victoria Medvec (Gilovich, Medvec and Kahneman, 1998), with Ralph Hertwig (where Barbara Mellers was the agreed arbiter, see Mellers, Hertwig and Kahneman, 2001), and with a group of experimental economists in the UK (Bateman *et al.*, 2003). An appendix in the Mellers *et al.* article proposes a detailed protocol for the conduct of adversarial collaboration. In another case I did not succeed in convincing two colleagues that we should engage in an adversarial collaboration, but we jointly developed another procedure that is also more constructive than the reply-rejoinder format. They wrote a critique of one of my lines of work, but instead of following up with the usual exchange of unpleasant comments we decided to write a joint piece, which started by a statement of what we did agree on, then went on to a series of short debates about issues on which we disagreed (Ariely, Kahneman, & Loewenstein, 2000). I hope that more efficient procedures for the conduct of controversies will be part of my legacy.

## Part 2 - Eulogy for Amos Tversky (June 5, 1996)

People who make a difference do not die alone. Something dies in everyone who was affected by them. Amos made a great deal of difference, and when he died, life was dimmed and diminished for many of us.

There is less intelligence in the world. There is less wit. There are many questions that will never be answered with the same inimitable combination of depth and clarity. There are standards that will not be defended with the same mix of principle and good sense. Life has become poorer.

There is a large Amos-shaped gap in the mosaic, and it will not be filled. It cannot be filled because Amos shaped his own place in the world, he shaped his life, and even his dying. And in shaping his life and his world, he changed the world and the life of many around him.

Amos was the freest person I have known, and he was able to be free because he was also one of the most disciplined.

Some of you may have tried to make Amos do something he did not want to do. I don't think that there are many with successes to recount. Unlike many of us, Amos could not be coerced or embarrassed into chores or empty rituals. In that sense he was free, and the object of envy for many of us. But the other side of freedom is the ability to find joy in what one does, and the ability to adapt creatively to the inevitable. I will say more about the joy later. The supreme test of Amos's ability to accept what cannot be changed came in the last few months. Amos loved living. Death, at a cruelly young age was imposed on him, before his children's lives had fully taken shape, before his work was done. But he managed to die as he had lived - free. He died as he intended. He wanted to work to the last, and he did. He wanted to keep his privacy, and he did. He wanted to help his family through their ordeal, and he did. He wanted to hear the voices of his friends one last time, and he found a way to do that through the letters that he read with pleasure, sadness and pride, to the end.

There are many forms of courage, and Amos had them all. The indomitable serenity of his last few months is one. The civic courage of adopting principled and unpopular positions is another, and he had that too. And then there is the heroic, almost reckless courage, and he had that too.

My first memory of Amos goes back to 1957, when someone pointed out to me a thin and handsome lieutenant, wearing the red beret of the paratroopers, who had just taken the competitive entrance exam to the undergraduate program in Psychology at Hebrew University. The handsome lieutenant looked very pale, I remember. He had been wounded.

The paratrooper unit to which he belonged had been performing an exercise with live fire in front of the general staff of the Israel Defense Forces and all the military attaches. Amos was a platoon commander. He sent one of his soldiers carrying a long metal tube loaded with an explosive charge, which was to be slid under the barbed wire of the position they were attacking, and was to be detonated to create an opening for the attacking troops. The soldier moved forward, placed the explosive charge, and lit the fuse. And then he froze, standing upright in the grip of some unaccountable attack of panic. The fuse was short and the soldier was certainly about to be killed. Amos leapt from behind the rock he was using for cover, ran to the soldier, and managed to jump at him and bring him down just before the charge exploded. This was how he was wounded. Those who have been soldiers will recognize this act as one of almost unbelievable presence of mind and bravery. It was awarded the highest citation available in the Israeli army.

Amos almost never mentioned this incident, but some years ago, in the context of one of our frequent conversations about the importance of memory in our lives, he mentioned it and said that it had greatly affected him. We can probably appreciate what it means for a 20-year old to have passed a supreme test, to have done the impossible. We can understand how one could draw strength from such an event, especially if - as was the

case for Amos - achieving the almost impossible was a once-off thing. Amos achieved the almost impossible many times, in different contexts.

Amos' almost impossible achievements, as you all know, extended to the academic life. Amos derived some quiet pleasure from one aspect of his record: by a large margin, he published more articles in *Psychological Review*, the prestigious theory journal of the discipline, than anyone else in the history of that journal, which goes back more than 100 years. He had two pieces in press in *Psychological Review* when he died.

But other aspects of the record are even more telling than this statistic. The number of gems and enduring classics sets Amos apart even more. His early work on transitivity violations, elimination by aspects, similarity, the work we did together on judgment, prospect theory and framing, the Hot Hand, the beautiful work on the disjunction effect and Argument-Based Choice, and most recently an achievement of which Amos was particularly proud: Support Theory.

How did he do it? There are many stories one could tell. Amos' lifelong habit of working alone at night while others slept surely helped, but that wouldn't quite do it. Then there was that mind - the bright beam of light that would clear out an idea from the fog of other people's words, the inventiveness that could come up with six different ways of doing anything that needed to be done. You might think that having the best mind in the field and the most efficient work style would suffice. But there was more.

Amos had simply perfect taste in choosing problems, and he never wasted much time on anything that was not destined to matter. He also had an unflinching compass that always kept him going forward. I can attest to that from long experience.

It is not uncommon for me to write dozens of drafts of a paper, but I am never quite sure that they are actually improving, and often I wander in circles. Almost everything I wrote with Amos also went through dozens of drafts, but when you worked with Amos you just knew. There would be many drafts, and they would get steadily better.

Amos and I wrote an article in *Science* in 1974. It took us a year. We would meet at the van Leer Institute in Jerusalem for 4-6 hours a day. On a good day we would mark a net advance of a sentence or two. It was worth every minute. And I have never had so much fun. When we started work on Prospect Theory it was 1974, and in about 6 months we had been through 30- odd versions of the theory and had a paper ready for a conference. The paper had about 90% of the ideas of Prospect Theory, and quite properly did not impress anyone. We spent the better part of the following four years debugging it, trying to anticipate every objection.

What kept us at it was a phrase that Amos often used: "Let's do it right". There was never any hurry, any thought of compromising quality for speed. We could do it because Amos said the work was important, and you could trust him when he said that. We could also do it because the process was so intensely enjoyable.

But even that is not all. To understand Amos' genius - not a word I use lightly - you have to consider a phrase that he was using increasingly often in the last few years: "Let us take what the terrain gives". In his growing wisdom Amos believed that Psychology is almost impossible, because there is just not all that much we can say that is both important and demonstrably true. "Let us take what the terrain gives" meant not over-reaching, not believing that setting a problem implies it can be solved.

The unique ability Amos had - no one else I know comes close - was to find the one place where the terrain will yield (for Amos, usually gold) - and then to take it all. This skill in taking it all is what made so many of Amos' papers not only classics, but definitive. What Amos had done did not need redoing.

Whether or not to over-reach was a source of frequent, and frequently productive tension between Amos and me over nearly 30 years. I have always wanted to do more than could be done without risk of error, and have always taken pride in preferring to be approximately right rather than precisely wrong. Amos thought that if you pick the terrain properly you won't have to choose, because you can be precisely right. And time and time again he managed to be precisely right on things that mattered. Wisdom was part of his genius.

Fun was also part of Amos' genius. Solving problems was a lifelong source of intense joy for him, and the fact that he was richly rewarded for his problem solving never undermined that joy.

Much of the joy was social. Almost all of Amos' work was collaborative. He enjoyed working with colleagues and students, and he was supremely good at it. And his joy was infectious. The 12 or 13 years in which most of our work was joint were years of interpersonal and intellectual bliss. Everything was interesting, almost everything was funny, and there was the recurrent joy of seeing an idea take shape. So many times in those years we shared the magical experience of one of us saying something which the other would understand more deeply than the speaker had done. Contrary to the old laws



of information theory, it was common for us to more information was received than had been sent. I have almost never had that experience with anyone else. If you have not had it, you don't know how marvelous collaboration can be ...

## References

Ariely, D., Kahneman, D. & Loewenstein, G. (2000). Joint comment on "When does duration matter in judgment and decision making". *Journal of Experimental Psychology: General*, 129, 524-529.

Arrow, K. J. (1982). Risk perception in psychology and economics. *Economic Inquiry*, 20, 1-9.

Ayton, P. (1998). How bad is human judgment? In *Forecasting with judgment*, G. Wright & P. Goodwin (Eds.). West Sussex, England: John Wiley & Sons.

Bateman, I., Kahneman, D., Munro, A., Starmer, C. & Sugden, R. (2003). Is there loss aversion in buying? An adversarial collaboration. (under review).

Cohen, L.J. (1981). Can human irrationality be experimentally demonstrated? *The Behavioral and Brain Sciences*, 4, 317-331.

Coombs, C.H., Dawes, R.M., Tversky, A. (1970). *Mathematical Psychology: An elementary introduction*. Oxford, England: Prentice-Hall.

Cosmides, L. & Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition*, 58, 1-73.

Erev, I. & Rapoport, A. (1998). Coordination, "magic", and reinforcement learning in a market entry game. *Games and Economic Behavior*, 23, 146-175.

Gigerenzer, G. (1991). How to make cognitive illusions disappear: Beyond 'heuristics and biases'. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology*, (Vol. 2, 83-115). Chichester, England: Wiley.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A rebuttal to Kahneman and Tversky (1996). *Psychological Review*, 103, 592-596.

Gilovich, T., Medvec, V.H., & Kahneman, D. (1998). Varieties of regret: A debate and partial resolution. *Psychological Review*, 105, 602-605.

Kahneman, D., & Schild, E.O. (1966). Training agents of social change in Israel: Definitions of objectives and a training approach. *Human Organization*, 25, 323-327.

Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice-Hall.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237-251.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, 47, 313-327.

Kahneman, D., & Tversky, A. (1984). Choices, values and frames. *American Psychologist*, 39, 341-350.

Kahneman, D., Knetsch, J., & Thaler, R. (1986a). Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review*, 76, 728-741.

Kahneman, D., Knetsch, J., & Thaler, R. (1986b). Fairness and the assumptions of economics. *Journal of Business*, 59, S285-S300.

Kahneman, D., Knetsch, J., & Thaler, R. Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 1990, 98(6), 1325-1348.

Kahneman, D., & Miller, D.T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93, 136-153.

Kahneman, D. (1987). Experimental economics: A psychological perspective. In R. Tietz, W. Albers and R. Selten (Eds.), *Modeling Bounded Rationality*, 11-20.

Kahneman, D., Fredrickson, D.L., Schreiber, C.A., & Redelmeier, D.A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, 4, 401-405.

Kahneman, D. (1994). New challenges to the rationality assumption. *Journal of Institutional and Theoretical Economics*, 150, 18-36. Reprinted as Kahneman, D. New challenges to the rationality assumption. *Legal Theory*, 3, 1997, 105-124.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions: A reply to Gigerenzer's critique. *Psychological Review*, 103, 582-591.

Kahneman, D., Ritov, I., and Schkade, D. (1999). Economic preferences or attitude expressions? An analysis of dollar responses to public issues. *Journal of Risk and Uncertainty*, 19, 220-242. Reprinted as Ch. 36 in Kahneman, D. and Tversky, A. (Eds.), *Choices, Values and Frames*. New York: Cambridge University Press and the Russell Sage Foundation, 2000.

Kahneman, D. and Tversky, A. (Eds.), *Choices, Values and Frames*. New York: Cambridge University Press and the Russell Sage Foundation, 2000.

Kahneman, D., and Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin and D. Kahneman (Eds.) *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press, 2002.

Klein, G. (2000). The fiction of optimization. In *Bounded rationality: The adaptive toolbox*, G. Gigerenzer & R. Selten (Eds.). Cambridge, USA: The MIT Press. 103-121.

Latham, G., Erez, M. & Locke, E. (1988). Resolving Scientific Disputes by the Joint Design of Crucial Experiments by the Antagonists: Application to the Erez-Latham Dispute Regarding Participation in Goal-Setting. *J. of Applied Psychology*, 73, 753-772.

Laibson, D. & Zeckhauser, R. (1998) Amos Tversky and the ascent of behavioral economics. *Journal of Risk and Uncertainty*, 16, 7-47.

Lopes, (1991). The rhetoric of irrationality. *Theory and Psychology*, 1, 65-82.

Meehl, P.E. (1954). Clinical versus statistical prediction: A theoretical analysis and a review of the evidence. Minneapolis: University of Minnesota Press.

Mellers, A., Hertwig, R., and Kahneman, D. (2001). Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science*, 12, 269-275.

Mischel, W. (1961a). Preference for delayed reinforcement and social-responsibility. *Journal of Abnormal and Social Psychology*, 62, 1-15.

Mischel, W. (1961b). Delay of gratification, need for achievement, and acquiescence in another culture. *Journal of Abnormal and Social Psychology*, 62, 543-560.

Raiffa, H. (1968). Decision analysis: Introductory lectures on choices under uncertainty. Reading, MA: Addison-Wesley.

Schreiber, C.A., & Kahneman, D. (2000). Determinants of the remembered utility of aversive sounds, *Journal of Experimental Psychology: General*, 129, 27-42.

Slooman, S.A. 1996. The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.

Stanovich, K. E. (1999). *Who is Rational?: Studies of Individual Differences in Reasoning*. Lawrence Erlbaum. Mahwah, New Jersey.

Sunstein, C., Kahneman, D., Schkade, D., & Ritov, I. (2002). Predictably incoherent judgments. *Standard Law Review*.

Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, 39, 36-90.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.

Tversky, A., & Kahneman, D. (1983). Extensional vs. intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 293-315.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, 59, S251-S278.

Williams, A.C. (1966). Attitudes toward speculative risks as an indicator of attitudes toward pure risks. *Journal of Risk and Insurance*, 33, 577-586.

---

1. Although I did not know it, I was not the first to propose this procedure in psychology: an adversarial collaboration had been reported by Latham, Erez, & Locke (1988).

From *Les Prix Nobel. The Nobel Prizes 2002*, Editor Tore Frängsmyr, [Nobel Foundation], Stockholm, 2003

This autobiography/biography was written at the time of the award and later published in the book series *Les Prix Nobel/Nobel Lectures*. The information is sometimes updated with an addendum submitted by the Laureate. To cite this document, always state the source as shown above.

Copyright © The Nobel Foundation 2002