

Toward a Theory of Bargaining: An Experimental Study in Economics

Alvin E. Roth

A growing number of economists have begun to conduct controlled laboratory experiments to test economic theories or generate new hypotheses. While the roots of this activity go back a number of years (1), controlled experimentation still constitutes a very small part of economic research (2). Different groups of researchers have used experimental methods to attack a wide variety of prob-

For simplicity, the term bargaining will be reserved here for situations in which exactly two agents negotiate over a set of possible agreements. While much of economic theory concerns the idealized case of "perfect competition," involving economic environments with sufficiently many agents so that any one can have only a negligible impact, bargaining is perhaps the economic phe-

Summary. Contemporary economic theories of bargaining depend on aspects of the bargainers' preferences that are difficult to observe. This makes these theories difficult to test in natural environments. It has proved possible, however, to design experiments to test these theories in a controlled, laboratory environment. The results of these experiments reveal shortcomings and incompleteness in the descriptive power of currently available theories of bargaining. However, these results also suggest important regularities in bargaining behavior. Together with recent theoretical developments, these results suggest some directions in which a more descriptively powerful theory of bargaining might be developed.

lems and, in particular, a number of market institutions have been the subject of penetrating experimental investigation (3). This article summarizes several experiments concerning bargaining behavior (4-7), and considers what these experiments, together with recent theoretical developments, suggest about future theoretical and empirical research.

Theories of Bargaining

Some of the most powerful explanatory hypotheses in economics concern the different readiness to tolerate risk (or risk posture) shown by different individuals, households, or firms. For example, the decision to buy or sell a particular insurance or commodity futures contract can be explained in terms of the risk posture of the agents involved, as can different investment decisions (8). It is thus not surprising that the principal economic theories of bargaining should predict that the outcome of bargaining depends on the risk posture of the bargainers.

nomenon furthest removed from perfect competition. The formal mathematical tools for the study of economic interactions that are not perfectly competitive constitute the field known as game theory, which is divided into two principal traditions: cooperative and noncooperative game theory. Cooperative game theory concentrates on what agreements agents are likely to reach, while noncooperative game theory concentrates on what strategies are likely to be adopted by the agents. Both of these traditions have produced models of bargaining (9, 10) in which each agent's risk posture is conveyed by comparing his preferences for risky and riskless alternatives.

If an individual has preferences that exhibit certain regularities (11), then his choices over (possibly) risky alternatives are the same as if he were maximizing the expected value of some real-valued function called his expected utility function, which is uniquely defined only up to an interval scale; that is, only up to the arbitrary choice of origin and unit. If a set of alternatives contains elements a , b , and c such that the individual prefers a

to b and b to c , then a utility function u representing this individual's preferences has $u(a) > u(b) > u(c)$. Since the choice of unit and origin is arbitrary, we may take $u(a) = 1$ and $u(c) = 0$. If $L(p)$ denotes the lottery that with probability p yields alternative a and with probability $(1 - p)$ yields alternative c , then the utility of participating in the lottery $L(p)$ is its expected utility, $u[L(p)] = pu(a) + (1 - p)u(c) = p$. If p^* is the value of p such that the individual is indifferent between b and $L(p^*)$, then $u(b) = u[L(p^*)] = p^*$.

Consider, for example, an individual faced with a choice of receiving \$500,000 for certain or participating in a lottery $L(p)$ that yields \$1,000,000 with probability p and otherwise yields \$0. If we set $u(\$1,000,000) = 1$ and $u(\$0) = 0$, determining $u(\$500,000)$ means determining the probability p^* that leaves this individual indifferent between receiving \$500,000 for certain or participating in the lottery $L(p^*)$. An individual who is indifferent when $p^* = 1/2$ would be called risk neutral, but most of us would require p to be considerably greater than $1/2$ before we would consider taking the lottery instead of the sure \$500,000, and this reflects our risk aversion. [A risk-preferring individual would be prepared to choose the lottery $L(p)$ over the sure \$500,000 for some value of p less than $1/2$]. Different individuals would switch from choosing the sure \$500,000 to choosing $L(p)$ at different values of p^* , reflecting differences in the way they compare risky and riskless alternatives. The higher the probability p^* , the less risk the individual is willing to bear, and the more risk averse he is said to be.

Economists have long been aware that choices made by individuals over various domains may not be sufficiently regular to be descriptively modeled as utility maximization (12). However, much of the use of utility maximization as a model of individual choice behavior rests on the manner in which it captures attitudes toward risk, and there is reason to believe that this aspect of the utility model may be quite robust (13). Nevertheless, a great deal of information is required to assess an individual's risk posture over a domain of any complexity. This makes it difficult to design compelling empirical tests, in uncontrolled natural environments, of theories whose predictions depend on agents' risk posture.

For reasons of theoretical parsimony, individual agents in most economic mod-

The author is the A. W. Mellon Professor of Economics and a member of the Center for Philosophy of Science at the University of Pittsburgh, Pittsburgh, Pennsylvania 15260.

els are represented entirely by their utility functions over the set of possible outcomes. Indeed, the simplest class of bargaining models are those in which the agents themselves are presumed to be able to assess one another's preferences and risk posture, and these models are said to presume a state of complete information. Thus an implicit assumption of the classical models of bargaining is that the only information about agents needed to predict the outcome of bargaining is the information about preferences and risk posture contained in their utility functions.

Specifically, bargaining theories in the tradition of cooperative game theory (10) all take the set of feasible utility payoffs to the bargainers as the data for predicting the outcome of bargaining. Formally, a bargaining game is defined as a pair (S, d) , where S is a subset of the plane and d is a point in S . The interpretation is that each element x of S corresponds to the utility payoffs to the bargainers corresponding to some feasible agreement a that they might reach, that is, $x = [u_1(a), u_2(a)]$ where u_1 and u_2 are the two bargainers' utility functions, and d corresponds to the utility payoffs resulting from disagreement. Let B denote some class of bargaining games [such as the class of all pairs (S, d) with S compact and convex]. Then the outcome predicted by any particular theory of bargaining that uses the pair (S, d) as data can be embodied in a function (called a solution) $f: B \rightarrow R^2$, such that $f(S, d)$ is an element of S . Thus any bargaining solution selects a particular outcome $f(S, d)$ of a bargaining game (S, d) , where all outcomes are represented in terms of the utility functions of the bargainers.

Solutions defined on some class B in this way (10) are theories of bargaining that predict that the outcome of bargaining depends only on the pair (S, d) , so that two bargaining situations that give rise to the same representation in terms of the bargainers' utility functions are predicted to yield the same utility payoffs to the bargainers. The first experiment reported below was designed to test this hypothesis, and the subsequent experiments were designed to investigate further the anomalies (from the point of view of existing theory) observed in the first experiment.

Methods and General Design

These experiments were designed to allow the expected utility of the participants to be determined. Participants bargained over the probability that they

would receive some monetary prize, possibly a different prize for each bargainer. Specifically, they bargained over how to distribute "lottery tickets" to determine the probability that each player would win his personal lottery (that is, a player who received 40 percent of the lottery tickets would have a 40 percent chance of winning his monetary prize and a 60 percent chance of winning nothing). The rules specified which distributions of lottery tickets were allowable. If no agreement was reached in the allotted time, each player received nothing. Bargaining situations of this kind, in which each bargainer has only two possible monetary payoffs, are called binary lottery games. It is straightforward to show that, on this simple domain of alternatives, the choice behavior of any individual who chooses high probabilities over low probabilities can be represented by a utility function.

To interpret the outcomes of a binary lottery game in terms of each bargainer's utility for money, recall that if each bargainer's utility function is normalized so the utility for receiving his prize is 1, and the utility for receiving nothing is 0, then the bargainer's utility for any lottery between these two alternatives is the probability of winning the lottery. The set of feasible utility payoffs in such a game equals the set of allowable divisions of lottery tickets. This simply reflects the fact that a utility function captures an individual's willingness to substitute any given alternative for a particular kind of lottery. In a bargaining game in which the agreements are themselves this kind of lottery, no substitution is involved in going from the agreement to the lottery, so the identification of p^* with the probability p of winning the lottery is immediate.

In all the experiments, each participant was seated at a visually isolated computer terminal, through which all instructions were presented, and all communication between bargainers conducted. The allowable forms of communication differed in different experiments, but the anonymity of the bargainers was always strictly preserved. The detailed methodology of each experiment is reported elsewhere (4-7).

Experiments 1 and 2

Note that the set of feasible utility payoffs in a binary lottery game does not depend on the prizes, so, if the players know the allowable division of lottery tickets, the game is one of complete information, regardless of whether each

player also knows the other's prize. Models of bargaining that depend only on the set of feasible utility payoffs to the bargainers thus predict that the outcome of a binary lottery game will not depend on whether the players know their opponent's prize.

The first experiment (4) was designed to test this hypothesis, among others. Participants played binary lottery games under either full or partial information. Under full information, each player knew his own and his opponent's potential prizes. Under partial information, each player knew only his own prize. In both conditions, the bargainers were free to send each other any messages they wished, subject only to the limitations that they could not identify themselves, or, in the partial information condition, discuss the monetary value of their prizes. Contrary to the predictions of the classical models, the outcomes observed in the two information conditions exhibited dramatic differences: under partial information, outcomes tended to be very close to an equal division of the lottery tickets, while under full information, outcomes showed a pronounced shift toward equal expected payoffs; that is, when the bargainers had full information and unequal prizes, the observed agreements gave a significantly higher probability of winning to the bargainer with the smaller prize. The difference observed between the two conditions thus indicates that theories which depend only on the pair (S, d) are insufficiently powerful to model this kind of bargaining.

Game-theoretic models in the noncooperative tradition, however, describe situations in more detail. The strategic form of a game describes not only the feasible utility payoffs, but also the strategies available to the players. (Formally, a pure strategy is a function from an agent's information to his potential actions; that is, it is a rule that tells an agent what to do based on what he knows.) In the bargaining games of the first experiment, strategy choices concern the choice of messages and proposals in the course of the bargaining. Since the strategies available to the bargainers depend on their information, the results of the first experiment might be due to the different strategies available to the bargainers in the two information conditions.

The second experiment (5) was designed to address this question. It employed binary lottery games with prizes stated in terms of an intermediate commodity, "chips," having monetary value. Bargaining was conducted under either high, intermediate, or low informa-

tion. In each condition, each bargainer knew the number and value of chips in his prize, but a bargainer's information about his opponent's prize varied with the information condition. Under high information, each bargainer knew the number and value of chips in his opponent's prize. Under intermediate information, each bargainer knew the number of chips in his opponent's prize, but not their value. Under low information, each bargainer knew neither the number of chips in his opponent's prize, nor their value. In the last two conditions, bargainers were prevented from communicating the missing information about the prizes, but were otherwise free to send each other any messages that preserved anonymity.

Two bargaining games in strategic form are strategically equivalent if there is an isomorphism between their strategy sets that preserves the utility payoffs of the bargainers. Theories in the noncooperative tradition depend on the strategic form of a game, and make the same predictions for any two strategically equivalent games.

This experiment took advantage of two kinds of strategic equivalence relations. Binary lottery games with prizes expressed in both chips and money, in the low information condition of this experiment, are strategically equivalent to binary lottery games with the same monetary prizes expressed in money alone, in the partial information conditions of the previous experiment, because any legal message in one kind of game would be legal in the other. So the strategy sets are the same for both kinds of games, as are the utility functions and the underlying alternatives. Also, games expressed in both chips and money in the intermediate information condition of this experiment are strategically equivalent to games expressed in money alone in the full information condition of the previous experiment. This is because any legal message in one kind of game could be transformed into a legal message in the other by substituting references to chips for references to money (or vice versa).

If the observed difference between partial and full information in the previous experiment was due to the different strategy sets in the two conditions, a similar difference should be observed between the low and intermediate information conditions of this experiment. The observed results did not support this hypothesis. The low and high information conditions replicated the partial and full information conditions of the previous experiment, but in the intermediate

Table 1. Mean outcomes to the \$20 and \$5 players in each information-common knowledge condition when agreements were reached (disagreements excluded). Outcomes are the mean lottery percentages obtained by the \$20 player (expressed first) and the \$5 player when they reached agreement.

Information	Common knowledge		Non-common knowledge	
	\$20 Player	\$5 Player	\$20 Player	\$5 Player
Neither player knows both prizes	48.8	51.2	47.5	52.5
Only the \$20 player knows both prizes	43.6	56.4	49.1	50.9
Only the \$5 player knows both prizes	33.6	66.4	37.2	62.8
Both players know both prizes	32.6	67.4	34.3	65.7

information condition the observed agreements tended to give both players equal probabilities, regardless of the size of their prize in chips. Information about the artificial commodity, chips, did not affect the outcomes in the same way as did strategically equivalent information about money.

Experiment 3

Experiments 1 and 2 uncovered an effect of information that cannot be accounted for within the framework of the classical models. The next experiment (6) was conducted to separate this effect into components that can be attributed to the possession of specific information by each of the bargainers.

In the two earlier experiments, either both bargainers knew their opponent's prize or neither did, and it was common knowledge whether the bargainers knew one another's prizes. Information is common knowledge between two individuals if it is known to both of them, and if each knows that the other knows, and each knows the other knows that *he* knows, and so forth (14). Two individuals can be thought of as having common knowledge about an event if it occurs when both of them are present to see it, so that they also see each other seeing it, and so on. In these experiments, a set of instructions provides common knowledge to the bargainers if it contains the information that both of them are receiving exactly the same instructions.

Each game of the third experiment was a binary lottery game in which one player had a \$20 prize and the other a \$5 prize. In each of the eight conditions of the experiment, each player knew at least his own prize. The experiment used a 4 (information) $\times$ 2 (common knowledge) factorial design. The information conditions were: (i) *Neither knows* his opponent's prize; (ii) the *\$20 player knows* both prizes, but the \$5 player knows only his own prize; (iii) the \$5

player knows both prizes, but the \$20 player knows only his own prize; and (iv) *Both players know* both prizes. The second factor made this information common knowledge for half the bargaining pairs, and not common knowledge for the other half. In the common knowledge conditions, the instructions stated that the bargainers were both reading the same instructions, and that certain private information would be presented at the end of the instructions. For example, in the \$20 player knows-common knowledge condition, both bargainers were instructed that one bargainer's private information would include both prizes while the other's would include only his own prize. In the non-common knowledge conditions, the instructions concerning the private information stated that each bargainer's private information might or might not include his opponent's prize. Bargainers were free to send any messages that preserved their anonymity, and in all conditions could make any statements they wished about the prizes.

Tables 1 to 3 summarize the results of this experiment (15). Table 1 indicates that the effect of information on the agreements reached is primarily a function of whether the bargainer with the smaller prize knows both prizes, and Table 2 shows that the frequency of disagreement is influenced by whether the bargainers' information about the prizes is common knowledge. Also, in the non-common knowledge conditions there is a tradeoff between the higher payoffs demanded by the \$5 player when he knows both prizes (as reflected in the mean agreements in Table 1) and the number of agreements actually reached (as reflected in the frequency of disagreements in Table 2).

Closer analysis of this tradeoff permits us to consider how accurately the bargainers were able to assess the minimum payoff that would ultimately be acceptable to their opponents. Most economic theories assume that agents are able to

Table 2. Frequency of disagreements. The m/n values indicate m disagreements out of n games played.

Information	Common knowledge		Non-common knowledge	
	m/n	Percentage	m/n	Percentage
Neither player knows both prizes	4/27	14	3/36	8
Only the \$20 player knows both prizes	6/30	20	4/24	17
Only the \$5 player knows both prizes	5/26	19	18/55	33
Both players know both prizes	5/30	17	9/35	26

Table 3. Mean outcomes to the \$20 and \$5 players in each information–common knowledge condition over all interactions (disagreements are included as zero outcomes). Within a column, means with common subscripts are *not* significantly different from one another using the Mann-Whitney U test ($\alpha = .01$); none were significantly different in the non-common knowledge conditions for the \$5 player.

Information	Common knowledge		Non-common knowledge	
	\$20 Player	\$5 Player	\$20 Player	\$5 Player
Neither player knows both prizes	41.6 _{ab}	43.3 _c	43.5 _a	48.2
Only the \$20 player knows both prizes	34.9 _{bc}	45.1 _{bc}	40.9 _a	42.4
Only the \$5 player knows both prizes	27.2 _c	53.6 _{ab}	25.0 _b	42.0
Both players know both prizes	27.2 _c	56.4 _a	25.5 _b	48.8

form correct assessments of their situation, the underlying hypothesis being that the behavior of such “rational” agents will be a good approximation of the behavior of observable economic agents. In strategic models, this assumption of rationality takes the form of assuming that agents’ strategies will be in equilibrium. An equilibrium is a pair of strategies, one for each bargainer, such that each bargainer’s strategy is the best response he could make to his opponent’s strategy. Thus strategies are in equilibrium when neither bargainer has reason to regret his strategy, and out of equilibrium when one of the bargainers could have responded better to his opponent.

Table 3 combines the data from Tables 1 and 2 to show the mean payoffs to the bargainers in each condition. The extent to which the observed strategies were in equilibrium can be assessed by considering whether the behavior observed in any condition could have profitably been substituted for the behavior observed in any other condition.

In the four common knowledge conditions, the fact that what kind of information each bargainer possesses is common knowledge often makes observed behavior infeasible in all but one condition. It is thus difficult to draw firm conclusions about equilibrium behavior in these conditions (16). However, in the four non-common knowledge conditions, since neither bargainer knows if his opponent

knows both prizes, a bargainer who knows both prizes is always free to behave precisely as if he knew only his own prize. (A bargainer who does not know his opponent’s prize may not be able to behave precisely as if he did, since, for instance, he cannot state his opponent’s prize.) If the observed behavior in these conditions is in equilibrium, bargainers who do not know their opponent’s prize must not do better on average than those who do, since, otherwise, a bargainer who knew both prizes could profit from adopting the strategy he would have used if he knew only his own prize. (And, to the extent that a bargainer who does not know his opponent’s prize can behave as if he did, equilibrium requires that bargainers who do know their opponent’s prize do no better on average than those who do not.)

First let us consider the \$20 players in the non-common knowledge conditions. A \$20 player whose opponent knew both prizes received a mean overall payoff of 25.5 if he knew his opponent’s prize and 25.0 if he didn’t, which are not significantly different. A \$20 player whose opponent knew only his own prize received a mean overall payoff of 40.9 if he knew his opponent’s prize and 43.5 if he did not, which also do not differ significantly from one another. Thus a \$20 player who managed to find out whether his opponent knew both prizes (which was often the case, since the transcripts show that \$5 players who knew both prizes fre-

quently mentioned the \$20 prize in their messages) could not improve his overall payoff by acting as he would have if his own information about his opponent’s prize were different. And a \$20 player who thought it equally likely that his opponent did or did not know his prize faced a 50–50 gamble of receiving 25.5 or 40.9 if he knew the \$5 player’s prize, or a 50–50 gamble between 25.0 or 43.5 if he did not, and, since these two gambles do not significantly differ, he also could not improve his expected overall payoff by acting as he would have if his own information about his opponent’s prize were different.

The situation facing the \$5 players in the non-common knowledge conditions was slightly different, since \$20 players virtually never revealed when they knew their opponent’s prize. If a \$5 player thought it equally likely that his opponent did or did not know his prize, then he faced a 50–50 gamble between 48.8 or 42.0 if he knew the \$20 player’s prize, or a 50–50 gamble between 42.4 or 48.2 if he did not. Since the expected values of these two gambles do not significantly differ, the \$5 player also had no opportunity to improve his expected overall payoff by acting as he would have if his information about his opponent’s prize had been different. Thus in the non-common knowledge conditions, the observed outcomes appear to conform to equilibrium behavior.

Discussion

These results, which cannot be explained within the framework of existing theory, open up several directions of further investigation. First, since the experimental results show that the customary specification of what constitutes “complete information” is inadequate for descriptive models of bargaining, even in situations that allow the bargainers’ utility for agreements to be determined, it would be desirable to know what information about the bargainers would be sufficient to account for the observed effects. Second, since the equilibrium pattern of results in experiment 3 suggest that the observed effects are consistent with the rationality assumptions embodied in strategic equilibrium, it would be desirable to know what features of existing bargaining models might be preserved in more descriptive models. Finally, it is necessary to consider the nature of the interaction between theoretical and experimental results in a field in which great theoretical progress

The Mysterious Expulsion of Steven Mosher

Stanford's anthropology department says it has its reasons, but Mosher asserts he's being muzzled for reporting unsavory truths about China

Steven Westley Mosher was regarded as a very bright and highly ambitious graduate student at Stanford University's department of anthropology. When China first opened its door in 1979 to field researchers, Mosher was one of the privileged few to be allowed in, beating out other more senior China scholars, who had been awaiting such an opportunity for decades. By every indication, he was on his way to an impressive career.

But in February, 11 members of Stanford's anthropology department voted unanimously to oust Mosher, 34, a fourth-year graduate student. The full explanation for Mosher's expulsion, a very rare action, has not been made public. The department will only say that, while in China, Mosher engaged in "illegal and seriously unethical conduct . . . and in doing so he endangered his research subjects."

The matter is not yet settled. On 25 April, Mosher filed a grievance with the dean of the School of Humanities and Sciences. Dean Norman Wessells has 30 days to decide whether to uphold the department's ruling. If he rules in the department's favor, Mosher would still have two more levels of appeal within the university administration.

Mosher contends that Stanford is trying to punish him for reporting on abortion, a subject highly sensitive to the Chinese. News accounts have given credence to his contention. But, according to people highly knowledgeable about the Stanford report, the real reasons for his expulsion are otherwise and that has created a difficult dilemma for the department. Sources say the report asserts that Mosher was involved in smuggling and in giving extraordinary gifts to collect his data. But the "persuasive evidence" of his misconduct, the department says, "cannot be revealed without endangering innocent persons," which include Chinese peasants. Faculty members say that, by disclosing the report, they would be violating a principle which they criticize Mosher for abusing—a research subject's right to privacy. Thus, Stanford is in the difficult position of asking the academic community to believe that it acted appropriately while not revealing its evidence.

Although neither Stanford nor Mosher will discuss that evidence, it is apparent the Mosher's style as a researcher displeased officials in the Chinese Academy of Social Sciences, as well as his mentors at the university.

In May 1981, an article written by Mosher appeared in a popular weekly magazine in Taiwan, detailing third-trimester abortions in a Chinese village where he had conducted research. The article, which criticized China's family-planning program, included photographs of Chinese women undergoing the abortions; their faces were not masked.

Stanford says making the report public could endanger innocent people.

The article incensed the Chinese and further strained negotiations over the social science exchange program between the two countries. According to Kenneth Prewitt, president of the Social Science Research Council, one of the main groups that coordinates the exchanges, Chinese officials warned him that the program would suffer "negative consequences" and "could be harmed" if Mosher were not dealt with.

Mosher contends that, faced with the possible demise of the entire exchange program, Stanford knuckled under. Indeed, since Mosher's article appeared, the Chinese no longer permit social science researchers to study the country's rural areas. But American specialists on China argue that Mosher did not cause the cancellation of the program. An official of the National Academy of Sciences' Committee on Scholarly Communication with the People's Republic of China says that the Chinese placed restrictions on researchers from the outset and that securing Chinese permission to place scientists in rural areas has always been difficult.

The anthropology department maintains that the article—however offensive for its violations of privacy and its lack of political judgment—was not the root cause of Mosher's ouster. Although the department refuses to release the 47-

page report of its two-year investigation, interviews with department faculty members and others outside Stanford provide a sketch of Mosher's purported behavior in China. According to these sources, Mosher tried to take antiquities out of the country without proper authorization and gave remarkable gifts to his research subjects. They refused to specify the exact nature of the antiquities or the gifts. Mosher, for his part, has denied any improper behavior. He said in an interview, "This is a typical vilification campaign by the Chinese. I didn't do anything wrong." Mosher is not handing out the report either. In his opinion, its release would damage his case in future legal proceedings.

According to several people, who agreed to speak on condition that they remain anonymous, allegations of Mosher's misconduct arose well before the Taiwan article was published. Chinese officials complained to U.S. social scientists studying in China as early as February 1980, more than a year before the Taiwan article appeared. But American China scholars there and in the United States initially disregarded their allegations. They believed that the Chinese, leery of letting field researchers into the country in the first place, were having second thoughts about the exchange program. After all, Mosher had impeccable credentials.

Mosher came to Stanford after serving as a naval officer in Japan and Hong Kong. He finished the master's program in Asian studies at Stanford and then entered the doctoral program.

When China finally agreed to permit American field researchers into the country, Mosher easily won approval, thanks to the backing of eminent China scholar G. William Skinner and another Stanford China specialist, Arthur Wolf. Mosher went to China for 5 weeks in the spring of 1979, returned to Taiwan—where he had already been conducting research for his dissertation—and then went back to China from September 1979 to June 1980. It was during the 9-month trip that trouble with Mosher began brewing.

Mosher, who was then married to a Chinese-American, chose to conduct his research in the village of his wife's rela-