Is There a Logic of Society?

Rohit Parikh
City University of New York[1]
365 Fifth Avenue
New York, NY 10016
rparikh@gc.cuny.edu

Abstract: *The modern phase in Logic begins in the middle of the 19th century with the work of Frege, Boole, de Morgan and Peirce. A second phase began in the early 20th century with Russell's discovery of his paradox, followed by Hilbert's program to try (unsuccessfully) to get around it. Finally, the thirties saw fundamental developments by Gödel, Turing, Post, Church and Tarski.*

By this time, the following questions looked like they were settled.

- What is a valid proof?
- What is an effective procedure?

But what are the corresponding notions for everyday life? Even if not all of ordinary life is logical or systematic, it has large logical and systematic elements in it. Can these be formalized at least to some extent, and how can we profit from this formalization?

1 Introduction

Continuing the discussion above, we came to identify 'valid proof' with 'proof in first order logic,' proved complete by Gödel. And we came to identify effective procedures with procedures which could be carried out (or somehow simulated) by Turing machines.

But reasoning precedes first order logic; in India and Greece, as well as in China, Logic has existed for several thousand years. And the notion of algorithm is implicit in so many things which happen in everyday life. We humans are tool-making creatures (as are chimps to a somewhat smaller extent) and both individual and social life are over-run with routines, from cooking recipes to elections.

So it would seem that here is another area, dealing with people rather than computers, crying out for development and study. Such a study seems to have been neglected until recently, but the work of [14, 16, 19, 11] has constituted a start.

[1]Brooklyn College and CUNY Graduate Center, research supported in part by a grant from PSC-CUNY

But let us start first with logic.

2 Logic vs Rationality

When we come to the question of why we want an argument to be logically rigorous, one answer is that a logically correct argument leads from true premises to true conclusions. This criterion may be modified somewhat and we may require, as Adams [1] does, that an argument should lead from premisses true with high probability to conclusions also true with high probability. But the point stands.

But what is important in social contexts is not just that one should be logical but also that one should be rational. The two requirements are related, but of course not identical. And the requirement of rationality enters into Grice's defense of the material conditional.

The material conditional $A \supset B$ of two formulas A and B is defined to be equivalent to $\neg A \vee B$ or $\neg(A \wedge \neg B)$. The material conditional is truth functional in that the truth value of $A \supset B$ depends *only* on the truth values of A and of B. It is true if A is false, or if B is true. No connection between A and B is needed.

But is $A \supset B$ the *same* as "If A then B"?

Dorothy Edgington shows that *if* "If P then Q" is truth-functional, then it must coincide with the material conditional. But then she gives a nice argument against interpreting "if A then B" as $A \supset B$. For consider the following statement, $X =$ *If God does not exist, then it is not the case that if I pray, my prayers will be answered.* Surely this seems right. Now there are two 'if's in X. If we formalize both the 'if's as material conditionals, we get $\neg G \supset \neg(P \supset A)$. Suppose now that I don't pray. Then P is false, and hence $P \supset A$ is true. Thus $\neg(P \supset A)$ is false. Thus if X is true, which we granted, $\neg G$ must be false, and hence G is true. I can prove the existence of God simply by not praying!

But surely the existence of God cannot be proved so simply and it would seem that "if A then B" cannot always be represented as $A \supset B$.

Can the material conditional be defended in spite of such examples? The best known defense of material conditionals is due to Paul Grice [6] who introduced his notion of *implicature* to this end.

Grice brings in rationality considerations in his celebrated defense of the material conditional. The material conditional $A \supset B$ is true if A is false or B is true. But then why do we resist that interpretation for *if A then B*?

Grice's answer is that it is a social convention that when making a statement one should supply the maximum relevant information. If I know $\neg A$ then it is *misleading* to say $A \supset B$ because the person who hears me will assume that I did not know $\neg A$; otherwise why not just say *that* rather than the longer formula $A \supset B$? Similarly, if I already know B, then it is *true* according to Grice, that "if A then B". But it is misleading to say that.

Coming back to our God example, suppose I say to someone, "Every time I pray to God, my prayers are answered." This is true if I do not pray. But the hearer assumes that this simple sentence "I don't pray" could not be what I intended to say because I have said something more complex instead. So he assumes that it was necessary to make the more complex statement, that I do pray, and that there is a connection between my praying and my prayers being answered. If I don't pray, then my statement, "Every time I pray to God, my prayers are answered" is true, but misleading.

Thus Grice proposes to defend the material conditional by saying that when it appears to differ in sense from "if A then B", that is because of the conventions of conversation.

Grice's analysis of conditionals, although brilliant, is not generally accepted because there are other problems (which we shall not discuss here) with his analysis. But the notion of *implicature* which he introduced has proved to be of great value. The implicature of a statement is something which is conveyed by the act of making that statement and goes beyond the mere truth of it. If a child says, "I am hungry," and I say, "Your mother will be back in a few minutes," I have implicated but not said that the mother will bring food. The child infers this from the fact that if the mother was not going to bring food, then my remark, even if true, is pointless.

This view of Grice has led to new developments in our understanding of conversation, that it does not consist merely of exchanging true assertions, but that it is an activity with a social purpose which is usually helpful (but sometimes not so). If a professor is asked for a letter of recommendation for a student and the professor says, "She always came regularly to class," and nothing else, then the professor has not *said* but he has implicated that there is nothing good about the student which is of relevance to the application.

We have seen how the logic which goes on within a social context is different from the logic of textbooks. But there is also a logic *about* social situations to which I shall now turn.

3 Plans

Suppose that I am going from my apartment to a hotel in Chicago. Then my schedule (or program) will consist of several steps. Taking a cab to Laguardia airport, checking in, boarding the plane, and then in Chicago, taking a cab to the hotel. Each step needs to be tested for correctness. For instance, if I do not have a picture ID, I cannot board the plane, and it is irrelevant that the other steps would have gone well.

The schedule I described is a simple *straight line* program. There could also be decision points like, *take a taxi if available, otherwise take a bus.* A plan with many such decision points would look like a tree rather than a line. But the entire schedule does have to be checked for correctness – something we do informally.

But other social procedures like organizing a conference or a wedding can be far more complex. We can invite aunt Betsy only if her ex-husband Eric has said he cannot come. Formal instruments can come in useful when the complexity surpasses some modest level. And we have not developed such instruments for *social* contexts. A logic is needed.

The paradigm of that logic comes (in part) from the logic of computer programs. A social procedure, to be useful, should serve some purpose, and it should be designed so that in normal circumstances the purpose is indeed fulfilled. The same holds for computer programs, and there is a well developed field devoted to proving that they work.

3.1 Proving Computer Programs Correct:

Let us look quickly at an example of the use of Hoare logic for proving a program correct [8].

Consider the program α_g for computing the greatest common divisor $gcd(u, v)$ of two positive integers u, v. We want this program to have the property

$$\{u > 0, v > 0\}\ \alpha_g\ \{x = gcd(u, v)\}$$

This property has three components: the *pre-condition* $u > 0, v > 0$, the desired *post-condition* $\{x = gcd(u, v)\}$ and the program α_g itself. We want that if we start with two integers u, v which are both positive, then the program α_g when it ends will have set x to the *gcd* of u, v. One such program α_g, *Euclid's algorithm*, is given by

$$(x := u); (y := v); (\text{while } (x \neq y \text{ do } (\text{if } x < y \text{ then } y := y - x \text{ else } x := x - y))).$$

The program sets x to u and y to v respectively, and then repeatedly subtracts the smaller of x, y from the larger until the two numbers become equal. If u, v are positive integers then this program terminates with $x = gcd(u,v)$, the greatest common divisor of u, v.

The reason is that after the initial $(x{:=}u);(y{:=}v)$, clearly $gcd(x, y) = gcd(u, v)$.

Now it is easy to see that both the instructions $x := x - y$ and $y := y - x$ leave the *gcd* of x, y unchanged. Thus if B is $gcd(x, y) = gcd(u, v)$, and β is $(if\ x < y\ then\ y := y - x\ else\ x := x - y)$, then $\{B\}\beta\{B\}$ holds. β *preserves* B.

Thus if the program α_g terminates, then by Hoare's rule for "*while*," $x \neq y$ will be false, i.e. $x = y$ will hold, and moreover B will hold. Since $gcd(x, x) = x$, we will now have $gcd(u, v) = gcd(x, y) = gcd(x, x) = x$.

Hoare's rules allow us to derive the properties of complex programs from those of simpler ones. Given program β let $\alpha =$"while A do β". A, B are first order formulas. Then the Hoare rule says that if we know $\{B\}\beta\{B\}$ i.e. that β preserves the truth of B,[2] then we can conclude $\{B\}\alpha\{B \wedge \neg A\}$. I.e. that α will also preserve B and will moreover falsify A.

$$\frac{\{B\}\beta\{B\}}{\{B\}while\ A\ do\ \beta\{B \wedge \neg A\}}$$

In other words, provided that β preserves the truth of B and that B holds when α begins, then B will still hold when α ends and moreover, A will be false. This rule allows us to predict that when the *gcd* program terminates, $gcd(u, v) = gcd(x, y)$ will still hold and $x \neq y$ will be false, i.e. x will equal y.

A social algorithm may have a similar proof of correctness, which essentially amounts to showing that provided the steps are performed correctly, the result will also be correct. Since social algorithms consist of a series of steps organized in some logical sequence, logical methods similar to those used for computer programs will sometimes work. [14] shows how formal methods can be used to prove the correctness of the Banach-Knaster algorithm for fairly dividing a cake.

However, this sort of logic works only for sequential computer programs. When it comes to people taking part in social algorithms, four other factors enter. They are

- **Communication and Knowledge**
- **Preferences and Incentives**
- **Co-ordination and Conflict**
- **Culture and Tradition**

[2] Indeed it is sufficient if β, applied once, yields B under the precondition that both A and B hold.

Of these four, the first enters already in computer science, i.e., in distributed computing. When *several* processors are co-operating on a task, then each needs to know what the others are doing so that actions can be co-ordinated [15]. [17] provides a model of knowledge acquired in the course of activity and relies on a Kripke model based on histories to develop a theory of the knowledge of each agent, whether a human or a computer.

4 Communication and Knowledge

Here is a discussion from [13]. The following examples illustrate the type of situations we have in mind.

Example 1: Uma is a physician whose neighbour is ill. Uma does not know and has not been informed. Uma has no obligation (as yet) to treat the neighbour.

Example 2: Uma is a physician whose neighbour Sam is ill. The neighbour's daughter Ann comes to Uma's house and tells her. Now Uma does have an obligation to treat Sam, or perhaps call in an ambulance or a specialist.

The difference between Uma's responsibilities in examples 1 and 2 is that in the second one she has knowledge of a situation which requires action on her part. In the first case, none of us would expect her to address a problem whose existence she does not know of. Thus any decent social algorithm must allow for the provision of requisite knowledge. However, in the example 3 below, it is the agent's own responsibility to *acquire* the proper knowledge.

Example 3: Mary is a patient in St. Gibson's hospital. Mary is having a heart attack. The caveat which applied in case 1) does not apply here. The hospital cannot plead ignorance, but rather it has an obligation to *be aware* of Mary's condition at all times and to provide emergency treatment as appropriate.

In all the cases we mentioned above, the issue of an obligation arises. This obligation is circumstantial in the sense that in other circumstances, the obligation might not apply. Moreover, the circumstances may not be fully known. In such a situation, there may still be enough information about the circumstances to decide on the proper course of action. If Sam is ill, Uma needs to know that he is ill, and the nature of the illness, but not where Sam went to school.

To see that such matters can have serious political consequences, consider the following from the *New York Times*.

By CLAUDIO GATTI and JAD MOUAWAD
Published: May 8, 2007

> Chevron, the second-largest American oil company, is preparing to acknowledge that it should have known kickbacks were being paid to Saddam Hussein on oil it bought from Iraq as part of a defunct United Nations program, according to investigators.
>
> The admission is part of a settlement being negotiated with United States prosecutors and includes fines totaling $25 million to $30 million, according to the investigators, who declined to be identified because the settlement was not yet public.
>
> The penalty, which is still being negotiated, would be the largest so far in the United States in connection with investigations of companies involved in the oil-for-food scandal.

[13] contains a detailed discussion of such cases relying on the semantics from [17] which is used to calculate the obligations.

Such knowledge issues arise all the time in real life. Suppose a shy student is in your office and you wonder if it is time for your next appointment. If you look at your watch, then you will know the time, but the student will also realize that you *wanted* to know the time, and may, being shy, leave even though he need not.

But we turn now to the other two aspects.

5 Incentives and Preferences

This aspect is of course what decision theory is all about. In decision theory, an agent is faced with some uncertainty as to how things are, and needs to make a decision. Suppose for instance that a die is to be tossed and the agent has to choose between (A) getting $30 if a six shows up, and (B) getting $12 if a six does *not* show up. Now the probability of a six is of course 1/6 (assuming that the die is fair) and that of a non-six is 5/6. Thus the expected value of the first choice is $30/6 or $5. That of the second choice is $12(5/6) or $10. It seems that the second choice is better and the theory says that that is what the agent should choose.

This situation was clear enough, but when an agent needs to choose between two jobs or two men who want to marry her, then numbers in the form of value or probabilities in the form of reals in [0,1] may be hard to come by. Still, it can be argued that some of the theory which we applied with the die will continue to apply. Savage [21] following on ideas of Ramsey and de Finetti shows that if the choices made by an agent (both in simple situation and in probabilistic bets) satisfy certain axioms, then there exists a utility function for that agent, *and* a subjective probability over the outcome space, such that the agent's actions can be seen as seeking to maximize the expected utility of an action.

We saw how an agent who believed that the die was fair and whose utility was linearly proportional to the amount of money would choose option (B) above. If the agent chooses (A) instead, then we would conclude that either the agent's subjective probability for a six was more than 1/6, i.e., the agent thought the die was not fair, or perhaps the utility of \$30 to the agent was more than 2.5 times the utility of \$12. The latter could happen if the agent needed to get to her home town, and the busfare was \$27. Then a 1/6 chance of getting \$30 would be better than a 5/6 chance of getting \$12. Yet another possibility is that the agent did not know probability theory [18].

Sen [23] points out that our reliance on utility theory is a bit naive, for people want many things which may not be comparable.

> I argue in favour of focusing on the capability to function, i.e., what a person can *do* or can *be* and argue against the more standard concentration on *opulence* (as in 'real income' estimates) or on *utility* (as in traditional 'welfare economic' formulations). Insofar as opulence and utility have roles (and they certainly do), these can be seen in terms of their indirect connections with well-being and advantage, in particular, (1) the *causal* importance of opulence, and (2) the *evidential* importance of utility (in its various forms, such as happiness, desire-fulfilment and choice)

Thus Sen emphasizes *abilities* rather than utilities. However, of course he did not have (I assume) access to the relevant material from computer science like analysis of algorithms, or logic of programs. Naturally he does not offer a detailed theory of abilities.

But such a detailed theory is crucial. For instance, when a bus line from town A to town B going through towns C,D,E is established, then the presence of the bus line increases the abilities of people living in these towns. Or alternately, when cellphones are invented and services established, then that too increases the abilities of people. Such things enter into the GNP, but only in an indirect way, via what people are willing to pay for a cellphone and not in terms of how it actually improves the abilities of people. This is why an expensive treatment for cancer or an expensive divorce turns out to have enormous (positive) implications for the GNP – not what we would intuitively expect.

Now decision theory is a theory of a *single* agent facing uncertainty.

6 Co-ordination and Conflict

Game theory extends this to a theory of an agent trying to maximize his utility in a situation where others are also trying to maximize theirs. Since the agent wants to maximize his own utility which depends on the actions of others, he needs a theory of what *they* want. The earliest major contributions here are due to von Neumann and Morgenstern, and to John Nash [25, 10].

A Nash equilibrium (for two players) is a pair of choices (a, b) by them which has the property that given that one agent is choosing a, the other's best bet is to choose b and vice versa.

Suppose for instance that a husband and wife want to go to a music concert and she prefers Mozart, whereas he prefers Stravinsky. Still each would rather go with the other than go alone.

Thus, putting the wife's location (or choice) first in the pairs below, their orderings are,
for the wife: $(M, M) > (S, S) > (M, S) > (S, M)$, and
for the husband, $(S, S) > (M, M) > (M, S) > (S, M)$.

Then there will be two Nash equilibria, one where they both go to Mozart and are both happy but the wife is happier. She gets her top choice and the husband gets his second choice. But if she *is* going to Mozart, then M is his best choice, For now he can only choose between (M, M) and (M, S) and the second possibility, (M, S) is worse for him. The other Nash equilibrium is where they both go to Stravinsky, and this time the husband is the happier one.

In so called zero sum games, a profit to one agent is a loss to the other, and each tries to outsmart the other. In such a case, pure strategy Nash equilibria like (M, M) do not usually exist, but mixed strategies do exist which are equilibria.

Suppose Jack and Ann are playing the game of matching pennies and Jack is the matcher and Ann is the mismatcher. Thus Jack gets both pennies if the pennies match and Ann gets them both if they do not. Then Jack wants to show heads if Ann does and tails if she shows tails. Ann wants the opposite and so each tries to guess what the other is going to do. There is no (pure) equilibrium here as there was in the Mozart-Stravinsky scenario.

The following story of either Indian or Iranian origin is very instructive of the game theoretic situation.

6.1 The two horsemen

Suppose we want to find out which of two horses is faster. This is easy, we race them against each other. The horse which reaches the goal first is the faster horse. And surely this method should also tell us which horse is *slower*, it is the other one. However, there is a complication which will be instructive.

Two horsemen are on a forest path chatting about something. A passerby M, the mischief maker, comes along and having plenty of time and a desire for amusement, suggests that they race against each other to a tree a short distance away and he will give a prize of $100. However, there is an interesting twist. He will give the $100 to the owner of the *slower* horse. Let us call the two horsemen Bill and Joe. Joe's horse can go at 35 miles per hour, whereas Bill's horse can only go 30 miles per hour. Since Bill has the slower horse, he should get the $100.

The two horsemen start, but soon realize that there is a problem. Each one is trying to go slower than the other and it is obvious that the race is not going to finish. There is a broad smile on the canny passerby's face as he sees that he is having some amusement at no cost. Figure I, below, explains the difficulty. Here Bill is the row player and Joe is the column player. Each horseman can make his horse go at any speed upto its maximum. But he has no reason to use the maximum. And in figure I, the left columns are dominant (yield a better payoff) for Joe and the top rows are dominant for Bill. Thus they end up in the top left hand corner, with both horses 'going' at 0 miles per hour.

	0	10	20	30	35
0	0, 0	100, 0	100, 0	100, 0	100, 0
10	0, 100	0, 0	100, 0	100, 0	100, 0
20	0, 100	0, 100	0, 0	100, 0	100, 0
30	0, 100	0, 100	0, 100	0, 0	100, 0

Figure I

However, along comes another passerby, let us call her S, the problem solver, and the situation is explained to her. She turns out to have a clever solution. She advises the two men to switch horses. Now each man has an incentive to go fast, because by making his competitor's horse go faster, he is helping his own horse to win! Figure II shows how the dominant strategies have changed. Now Joe (playing row) is better off to the bottom, and Bill playing column is better off to the right – they are both urging the horse they are riding (their opponents' horse) as fast as the horse can go. Thus they end up in the bottom right corner of figure II. Joe's horse, ridden by Bill comes first and Bill gets the $100 as he should.

	0	10	20	30	35
0	0, 0	0, 100	0, 100	0, 100	0, 100
10	100, 0	0, 0	0, 100	0, 100	0, 100
20	100, 0	100, 0	0, 0	0, 100	0, 100
30	100, 0	100, 0	100, 0	0, 0	0, 100

Figure II

7 The Free-rider Problem

When many agents participate in an activity where the benefit is to all, the so called *free rider problem* can arise. If all households are asked to use less water to avoid a shortage, then someone who surreptitiously used more water will benefit. Whether there is water shortage or not will depend on the actions of thousands of others, and his own action will not decide whether he has to endure a water shortage. On the other hand by taking long showers he will get more pleasure, and if the others save, he will *also* get the benefit of more water in the future.

Some readers may be familiar with the **Tragedy of the Commons**, In [7], Hardin introduces a hypothetical example of a pasture shared by local herders. The herders are assumed to wish to maximize their yield, and so will increase their herd size whenever possible. The utility of each additional animal has both a positive and negative component:

- **Positive:** the herder receives all of the proceeds from each additional animal
- **Negative:** the pasture is slightly degraded by each additional animal

Crucially, the division of these components is unequal: the individual herder gains all of the advantage, but the disadvantage is shared between all herders using the pasture. Consequently, for an individual herder weighing up these utilities, the rational course of action is to add an extra animal. And another, and another. However, since all herders reach the same conclusion, overgrazing and degradation of the pasture is its long-term fate.

7.1 Akbar and Birbal

There is a wonderful Indian story about the Mughal emperor Akbar[3] and his minister Birbal [2] about the way in which incentives (and knowledge) affect a social algorithm. Birbal had asserted to the emperor that all wise people think alike and had been challenged.

Then at Birbal's instance, all men in Agra, the capital, were ordered to come at night to the palace grounds, and pour one potful of milk into the pool on the palace grounds. The punishment for not doing so was severe, so one by one, all the residents came at night and poured a potful into the pool which was covered by a white sheet.

When the sheet was removed in the morning, it turned out that the pool was entirely full of water! Birbal explained, "Your majesty, each man thought that if he, and he alone were to pour a potful of water into the pool, it would not make much difference, and no one would notice. So, since all wise men think alike, and of course all of your subjects are wise, they all did the same thing and the pool is full of water."

Akbar was a emperor of India during the second half of the 16th century, so this fable predates William Forster Lloyd's 1833 parable of the Tragedy of the Commons by some 300 years. However, Aristotle, who said, "That which is common to the greatest number has the least care bestowed upon it," predates Birbal by almost two thousand years!

There is a similar situation which we might call the *Tragedy of the Beer*. Suppose that ten people go out to eat dinner and there is a convention that the total bill is to be split equally. Many restaurants in fact encourage this convention by refusing to give itemized bills. Suppose that a beer costs $7 which is a bit high. But if you, as a diner, are thinking

[3]Akbar was the grandfather of Shah Jehan who built the Taj Mahal in memory of his wife Mumtaz Mahal

whether to have that extra beer, its cost to *you* will only be 70 cents (plus tax and tip). Again, the game theoretic situation is the same as before, the benefit is to the individual and the cost is to the whole group; an individual who acts irresponsibly will gain.

8 Culture and Tradition

Finally we come to culture, which is our shorthand for inherited algorithms, so that we are using the word *culture* in a fairly wide sense.

One interesting example of the influence of culture is the ultimatum game, what the theory predicts about it and what actually happens in practice. Suppose I offer a hundred rupees to Laxmi and Ram in the following way. First Laxmi decides who is to get how much of the hundred each of them will get. Then Ram has to decide if he accepts her division. If he does, then the money is allocated as Laxmi decided. If he rejects, neither gets anything.

Suppose Laxmi decides that she gets 90 and Ram gets 10. Now it is Ram's turn to decide. According to conventional game theory, 10 is better than nothing which is what he will get if he rejects. So he should accept and take the 10. In practice, it is quite unlikely that Ram will accept less than 30. Moreover, in some cultures, Laxmi may offer 40:60 with 40 for herself and 60 for Ram, and Ram may reject this division as being excessively generous.

So actual behavior differs from theoretical prediction and seems to follow some pre-existing cultural pattern.

Essentially our view is that in game theoretic situations, a complex calculation needs to be made and the correctness of the calculation depends on the other party making a corresponding calculation.

Theoretical considerations cannot work out these scenarios and in practice people use ready made solutions which are part of culture or tradition.

Thus both Lewis and Schelling [9, 22] refer to the notion of saliency. If two people in New York city get separated, how will they find each other? One solution is to say that they will make their way to the giant clock in Grand Central station which is a *salient* point.

But what makes that a salient point? Surely that is part of the culture. If the two people are Chinese they may go to the Buddhist temple in Chinatown instead!

Group membership is crucial here. People tend to use algorithms which they have learned or inherited from their group, which may be a group defined by nationality, or ethnicity, or religion, or even by some subfield of some area. Thus the same kind of object, a subset of some space W will be called a *proposition* by a philosopher and an *event* by an economist. When a philosopher talks to an economist, they may use different conventions, and there could be confusion.

in [3] Michael Bacharach suggests that identifying oneself as a member of a group is crucial to making choices when co-ordinated action is needed.

> What produces team reasoning? Team reasoning — in any of its versions — is unmysterious; it is an algorithm, or routine, for arriving at decisions. As many things could get people to execute it, as could get them to do long division. It could be done as a game, or as a mathematical exercise, or because someone in control of a group of people found it in her interest that they should. (p. 135)

Recently, when attending a logic conference at IIT Mumbai, I met the young graduate student Tithi Bhatnagar who was doing research in what was important to human beings. She has found that regardless of age, gender or class, the most important thing for people is relationship. This aspect of human beings, that we value relationships above most everything, is crucial to understanding human beings who have, for far too long, been thought of primarily as selfish individuals looking out only for ourselves.

In sum, we have suggested that a study of society and its social algorithms requires logic (especially the logic of programs), but also requires attention to the other three legs of knowledge, incentive, and culture. With all these things in hand a good theory can be worked out.

Acknowledgements: Thanks to Johan van Benthem, Jongjin Kim and Eric Pacuit for comments.

References

[1] Ernest Adams, "Probability and the logic of conditionals," in *Aspects of Inductive Logic*, edited by Suppes and Hintikka, (1968), North Holland, 265-316.

[2] Amita Sarin, *Akbar and Birbal*, Penguin India (2005).

[3] Michael Bacharach, *Beyond Individual Choice: teams and frames in game theory*, edited by Natalie Gold and Robert Sugden, Princeton University Press, 2006.

[4] Johan Benthem, Jan van Eijck and Barteld Kooi, "Common Knowledge in Update Logics," in Ron der Meyden (ed.), *Theoretical Aspects of Rationality and Knowledge*, Proceedings of the Tenth Conference, National University of Singapore, 2005, pp. 253–261.

[5] Kurt Gödel, "Die Vollständigkeit der Axiome des logischen Funktionen-kalküls", Monatshefte für Mathematik und Physik **37** (1930), 349-360.

[6] Paul Grice, *Studies in the Way of Words*, Harvard University Press, 1989.

[7] Garrett Hardin, "The tragedy of the commons," *Science*, **162** (1968) 1243-1248.

[8] Dexter Kozen, and J. Tiuryn, "Logics of Programs", in *Handbook of Theoretical Computer Science*, vol B, edited by J.V. Leeuwen, North Holland 1990, pp. 789-840.

[9] David Lewis, *Convention: A philosophical study*, Harvard University Press (1969).

[10] John Nash, "Equilibrium points in N-person games," *Proceedings of the National Academy of Sciences, USA*, **36** (1950) pp. 48-49.

[11] Eric Pacuit, *Topics in Social Software: Information in Strategic Situations* Doctoral dissertaion, City University of New York (2005).

[12] Eric Pacuit and Rohit Parikh, "Social Interaction, Knowledge, and Social Software," In *Interactive Computation: The New Paradigm,* Springer-Verlag 2006. Editors Dina Goldin, Scott Smolka and Peter Wegner.

[13] Eric Pacuit, Rohit Parikh and Eva Cogan, "The logic of knowledge based obligation," presented at Society of Exact Philosophy meeting in Maryland, and at *DALT 2004. Synthese*, **149** (2006) 311-341. Also in in *Knowledge, Rationality & Action* (2006), 57-87.

[14] Rohit Parikh, "The Logic of games and its applications" *Annals of Discrete Math.*, 24 (1985) 111-140.

[15] Rohit Parikh, "Knowledge based computation (Extended abstract)" in *Proceedings of AMAST-95* Montreal, July 1995, Edited by Alagar and Nivat, Springer Lecture Notes in CS no. 936, 127-42.

[16] Rohit Parikh, Social Software, *Synthese*, **132**, Sep 2002, 187-211.

[17] Rohit Parikh and R. Ramanujam, "A Knowledge based Semantics of Messages," *J. Logic, Language and Information* **12** 2003, 453-467

[18] Rohit Parikh, "WHAT do we know and what do WE know?", the proceedings of *Theoretical Aspects of Rationality and Knowledge*, Ed. Ron van der Meyden, June 2005, University of Singapore.

[19] Marc Pauly, *Logic for Social Software*, Ph.D. Thesis, University of Amsterdam. ILLC Dissertation Series 2001-10, ISBN: 90-6196-510-1.

[20] Ramsey, F.P., ‘Truth and probability’, in *The Foundations of Mathematics*, Routledge and Kegan Paul (1931), 156-198.

[21] L. J. Savage, *The Foundations of Statistics*, Wiley 1954.

[22] Thomas Schelling, *The Strategy of Confict*, Harvard University press (1960).

[23] Amartya Sen, *Commodities and Capabilities*, Elsevier Science (1985).

[24] Turing, A. M., 1936-7, “On computable numbers, with an application to the Entscheidungsproblem”, *Proc. London Maths. Soc.*, ser. 2, 42: 230-265;

[25] John von Neumann and Oscar Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, 1944.