
Social Software and the Social Sciences

Keith Dowding and Rineke Verbrugge

Our Philosopher, Political Scientist, Logician and Ethicist meet yet again at a conference after lunch, and are joined by some new discussants: a Computer Programmer and a Cognitive Scientist. This time, there is also a Chair.

Chair: So we are here today at the Lorentz Center in Leiden to discuss precisely what it is that is the subject of this conference *Games, Action and Social Software*. What is social software? And how is it related to the social sciences?

Logician: We have already discussed the question of demarcating social software from several angles in the first few discourses. For the new participants in this discussion, let me give a short reminder. The most obvious place to start to answer that question is to look at the original articles by Rohit Parikh [172; 173] that introduced the term into our language. He suggested, I recall, that the issue of constructing and verifying social procedures in as systematic a manner as verifying computer software is pursued by computer scientists, be called “social software”. And such a process requires theories of program correctness; an analysis of programs and checks to see that different programs do not frustrate each other.

Philosopher: That is all very well but are social processes really like computer programs in that way? Computer software makes applications run; but in social life people make things run.

Logician: As Parikh first suggested people must want to carry out the program. Their aims and the algorithm used must work together. The algorithm must be set in such a way that it somehow conforms to their wishes. We can see that we need to set some optimizing conditions, and then see the algorithm as something that allows that optimizing.

Philosopher: One problem is that people have different sets of interests. And we know from Arrow's Theorem [6], which we will extensively discuss at another occasion (page 77), we cannot have a means of aggregating any set of interests into a social welfare function. One implication is that there is always the possibility that there will be no agreement on an efficient algorithm.

Logician: One direction we can go in is the following. We can agree that in many areas of social life there will be disagreement over the most desirable outcomes. However, it does not follow from the ever-present possibility of disagreement that there will *always* be disagreement. Very often a small group of people, or even a whole society, will agree on what the optimal outcome will be. Under those conditions we can study in a systematic way what the best algorithm might be for attaining their desires.

Cognitive Scientist: We could make distinctions between different sorts of problems. At one level there is no disagreement about what people want. The only difficulty is attaining that outcome. These are pure coordination problems. Dunin-Kępicz and Verbrugge show that once a certain level of collective beliefs, intentions and commitments is established, there is no difficulty in attaining cooperation [62; 73; 74]. It is in everyone's interests to coordinate their activities according to the algorithm.

Political Scientist: Then there are what we might call "collective action problems". Here there might be general agreement about the best outcome so part of the problem is a coordination game, but there are also conflicts of interest. In a pure coordination game for example, everyone might need to act to attain some outcome that is in everyone's interests. In a collective action game, fewer than 100 per cent need to act for the outcome to be achieved. Here there is the possibility of free riders and so each person wants the outcome, but also knows that the outcome can be optimally achieved without their input. For example, the famous analogy in Hume's *Treatise of Human Nature* [121] was that two farmers might agree to meet the following morning to dig a ditch between their fields. They are both needed so both have the incentive to turn up. But a hundred villagers find it harder to all agree to turn up to drain a field, since not all 100 are required. But if all try to free ride then none turn up and the field is not drained.

Philosopher: Yes I see that. But is that not a problem of assuring compliance: with the 100 villagers it is simply that we cannot tell who has turned up or not. We might not see our neighbour and complain to him that he did not

help, but he replies he was there, but in a different corner of the field.

Political Scientist: It might be perceptibility that is the problem, but not necessarily. Think of shortcuts across the grass. If we all take the shortcut the grass becomes worn and mud patches appear. But if only a few take a shortcut occasionally the grass is not affected at all. We might be able to see who takes the shortcut. And as long as it is not always the same people there is no problem. Optimality suggests that people should only take the shortcut if there is an emergency. In other words, there might be good reasons to allow a few free riders. But my point is that for collective action problems, as I am defining them, there is an added problem of conflict as well as cooperation. With pure coordination games the element of conflict is absent.

Logician: I see that.

Cognitive Scientist: Negotiation is a prime example of such a mixed situation of ‘co-opetition’ [184; 38].

Political Scientist: Then we have games which involve conflict over the outcomes as well. Here we might all agree that we should coordinate on some outcome, but we disagree over which. That is, we all want to play an equilibrium strategy, but which equilibrium strategy? Then of course, we have pure conflict games.

Philosopher: Okay, we can see different types of problems, but I thought social software was about finding solutions.

Logician: The point as I see it is that different types of algorithms might be necessary for different types of problems. Economists might be best at constructing answers in some fields, for example about market relations. Logicians are better at solving other problems, for example those concerning belief, knowledge or common knowledge. Political scientists are good at yet other questions, for example about voting procedures. The social software program is not about taking over the social sciences, it is rather an umbrella term to bring people from different disciplines together. Rather as this conference is trying to do.

Computer Programmer: Economists themselves seem to believe that they are best equipped to deal with all domains. One sometimes gets the impression that they want to take over the social sciences. There was a course at Harvard run by two political scientists last year called “The Economists are Coming”.

Philosopher: The big thing is game theory, you know. Logicians have just discovered game theory, and game theorists are just discovering the issues surrounding knowledge and the problems of common knowledge, and they are becoming interested in language interpretation. Each side is aware of the other like the rhino and the elephant in the jungle. Each knows how big they are themselves but can only see a little bit of the other. They are about to come into the open and find out. Will the elephant push the rhino over, or the rhino the elephant?

Political Scientist: Which is the elephant and which the rhino?

Philosopher: We shall see.

Logician: Actually both are grassland animals and don't live in the jungle. And logicians and economists do not have to fight.

Political Scientist: I think we are getting off the point. I am sure that logicians can learn from game theorists and vice versa. My point is that both do social software. They both examine social interactions using algorithmic processes.

Computer Programmer: I am afraid we have gone too quickly for me. We are talking about what social software is, and now I learn it is what logicians and game theorists do. So game theory is social software? Are you saying that everyone who acts in society is a game theorist?

Logician: Of course not everyone is a game *theorist*, but everyone is playing games. Games of coordination, cooperation and conflict. Every social exchange is part of a game, and each game is part of the overall game of life. The point is that we can model those games and see whether they are being played well, or could be played better. We can model the structures and then examine all the strategies available within those structures, and see if there are better strategies we can play; or better structures we can evolve that will help us reach the outcomes we desire.

Computer Programmer: It seems to me that "social software" as you describe it is just a metaphor for modelling. In agent-based social simulations, which can be seen as a part of Artificial Intelligence, we put in all the parameters and the program tells you what to do. This is a micro-level and this is where social software should concentrate. If we find out what is going on in the human mind by modelling it in terms of Artificial Intelligence then that will give us our social software. After all, social software has to exist somewhere, and where better than in the mind. The mind has been programmed to deal

with its environment. We program the cognitively plausible simulated agents, place them in a simulated social environment and model their behaviour. In that way social software is not a metaphor or analogy anymore. It is the real thing. Social software will be software in a programmed agent-based social simulation [224]. We will be modelling the real thing with a real thing.

Philosopher: It will not be exactly the same thing. Artificial Intelligence is programmed on hardware and the human mind is programmed on a biological organism which does not operate in the same way as computers. The brain is not a parallel processor for one thing. I am not putting down Artificial Intelligence, it can teach us much, but it is not yet a good model of the human mind. And what it does is less analogical than the mathematical modelling we are discussing.

Computer Programmer: Is a neural net software? We may not know what is going on, we just have outcomes that we can discuss. Or do we have to be able to understand and represent the action mathematically? A lot of agent-based models are designed to examine situations that we cannot yet model mathematically. The math gets too horrendous so we set up automata and they play a set of games within parameters and we see the sorts of outcomes that occur. Sometimes the same beginning states can lead to very different outcomes. We can of course start to understand why the different outcomes occur, but we do not know in the kind of detail that a computer programmer needs to know to ensure his program does what he wants it to do.

Logician: And I think this misses the point. We do not have to model social software in terms of a representation of the human mind. Rather we can look at processes which occur at a different level.

Computer Programmer: We can model things at all sorts of levels. If we adopt Dennett's intentional stance [60], we do not need to worry about what is "really" going on in actors' minds. We have a set of institutional processes – algorithms – that lead to sets of outcomes and we interpret these in ways that make sense. We interpret behaviour as being rational. Actors do not have to be consciously following the rules that our algorithms model. We explain their behaviour by the algorithms no matter what is "really" going on. Dennett would say, indeed, that if our algorithms are better predictors of behaviour than what the actors say they are doing, then it is the algorithms that are "real".

Logician: We can perhaps predict what people will do in certain situations in

terms of the intentional stance. But to improve their situation we need to set up new algorithmic processes which appeal to their interests and beliefs and lead them to better solutions.

Computer Programmer: That is something I am not sure about. Is social software a positive or a normative exercise? Is it about explaining social processes or is about improving them?

Logician: It could be either. We use our algorithms to explain the processes we see, and then interrogate those processes to see if they can be improved. That is the example of the airport luggage carousel in Parikh's [173, page 193]. Let me quote Parikh:

Then on arrival one has to wait at a moving carousel which brings all pieces of luggage one after another. If we are lucky, ours is among these. However, a curious phenomenon takes place which resembles the problems of the Prisoners' Dilemma, or The Tragedy of the Commons. One gets a better view of the approaching suitcases if one goes closer to the carousel. But by doing this, one inevitably blocks the view of one's neighbour who then also proceeds forward towards the carousel. When this process is finished, all passengers are right at the carousel, blocking each other's view, and every one is worse off than if no one had walked up to the carousel. [...] And the airline does have a solution to the carousel problem. All they need to do is to paint a line about 18 inches from the carousel and post signs saying "Do not cross the line until you actually see your own suitcase approaching". Then social pressure would keep people from crossing the line unnecessarily, and everyone would enjoy a better view of the oncoming luggage. Subways routinely do something similar at platforms to prevent passengers from falling onto the tracks or being hit by an incoming train.

We can see why everyone presses forward to find their bags, but can think of a simple institution—based on social hardware, if you like—that is in everyone's interests to improve the situation for all.

Cognitive Scientist: But if it is about explaining what really happens, then game theory is not a good tool. What really happens in people's minds is not the equilibrium selection strategies of classical game theory. According to some cognitive scientists and behavioural economists, human beings use decision heuristics. They see a situation within one frame of reference and then use the appropriate decision heuristic for that frame. Sometimes what we might see as the objective frame is different from the frame of reference

chosen by the human subject. And that is why we get sub-optimality; or we find that people are inconsistent or contradictory – they are choosing different frames of reference for what is essentially the same decision process.

Logician: I assume you are referring to the work of Cosmides and Tooby [230; 14], and some of the results of experimental psychologists such as Kahneman and Tversky [125]. And you are right, we might understand how evolution has fitted us for working out some kinds of problems within certain contexts and not for doing so within others.

Cognitive Scientist: Such as the Wason selection test?

Logician: : Exactly.

Political Scientist: What is the Wason selection test?

Cognitive Scientist: Experimental subjects are shown the four cards A K 4 7 and are told that each card contains a numeral on one side and a letter on the other. They are asked to evaluate the truth of the statement “If there is a vowel on one side, then there is an even number on the other side” by turning the minimum needed number of cards.

Logician: Based on the truth table for the material implication, they should of course turn around only A and 7.

Cognitive Scientist: Exactly. But in reality, only around 4% of the subjects correctly take these two cards. Around 33% takes only the A, while around 46% take the A and the 4, and around 17% take still other combinations. On the other hand, experimental psychologists have given subjects isomorphic problems in a concrete setting. For example, subjects are shown the four cards: “beer,” “cola,” “16 years,” and “22 years,” and are told that on one side of each card is the name of a drink; on the other side is the age of the drinker. They are asked to put themselves into the shoes of the barman, and are asked what card(s) they must turn over to determine whether the following statement is false: “If a person is drinking beer, then that person must be over 19 years old.” And lo and behold, almost all subjects correctly turn over the cards “16 years” and “beer”. You see, they cannot solve the problem when given as a logical exercise, but they can solve exactly the same problem when framed as a social exercise. Cosmides and Tooby draw the conclusion that humans do not have a general capacity for abstract logic, but must have developed a specialized module for detecting cheaters [50; 230].

Logician: Objection! As Van Lambalgen and Stenning showed, these two problems are not at all isomorphic, let alone “exactly the same” [219; 218; 220]! In the barman version, the subjects are asked to evaluate a *deontic* conditional, containing the word “must” in the consequent, and they can do that *per drinker*. It is highly likely that such a deontic conditional is much easier to process for subjects than a descriptive conditional, such as the one with the vowels and the even numbers, which is to be interpreted as a kind of general statements about all cards. So, in her analysis of the Wason card selection task [50], Cosmides wrongly conflates two logically different kinds of propositions to be of the same logical form but to differ only in content. Still, we can understand why we are able to solve the Wason selection test if it is framed right, but not otherwise. The social software program, in its normative guise, might ask us to frame important issues—where solving the Wason selection really matters—in such a way that people can solve them more precisely.

Political Scientist: Perhaps logicians can solve it when it is posed as a logical problem on cards, but not when they serve behind the bar and try to find out who is an underage drinker.

Logician: Not funny.

Philosopher: But I do not precisely see the point here. Solving the Wason selection test is not exactly an algorithm, is it? I mean, in computing you write a program and if you get it right, it does what you want it to. Here you are trying to set up a problem in such a way that ordinary people can solve it.

Political Scientist: And if they can solve it then you have written the right algorithm. But what interests me is the evolutionary point. It was suggested that what goes on in the mind is not the same as game theory. People do not solve games as game theorists do. They just bumble along somehow. And this is important. One could have two entirely different causal explanations of some outcome, that have the same game-theoretic explanation.

Cognitive Scientist: That’s right! In evolutionary biology, Maynard Smith and Price devised a simple conflict model, where members of a species fight over some resource [146]. Winning the resource is worth 50 points. Individuals have access to two strategies. As a “hawk” they fight over the resource, and as a “dove”, they merely threaten and posture. If both play hawk, they will fight until one is injured (-100 points) and the other gets the resource. The

pay-offs are computed as expected values, so for example for hawk-hawk we have $\frac{1}{2}(50) + \frac{1}{2}(-100) = -25$. If two doves meet, one will eventually get the resource of 50 points, but they will both get -10 for wasted time. Here's the pay-off matrix [222]:

	hawk	dove
hawk	(-25,-25)	(50,0)
dove	(0,50)	(15,15)

It turns out that in populations with pure hawks and doves, if there are almost entirely doves, the doves would expect an average of 15 points per game. In that situation, the hawks are genetically advantaged because they would expect an average of 50 points and their population would rise: the hawk minority would invade the dove population. Similarly, a population of almost entirely hawks would be unstable: then the doves are the guys with an advantage, winning an average of 0 points against -25 points for hawks. Maynard Smith and Price introduced the concept of *evolutionary stable strategy*. It turns out that a population of $\frac{7}{12}$ hawks and $\frac{5}{12}$ doves is stable, and the proportion stays fixed. In such a population, the expected pay-offs for both hawks and doves are $\frac{25}{4}$.

Political Scientist: That's right. But notice that there is an interesting aspect here. The same equilibrium as in the case with the pure strategies would be attained if every individual played a *mixed strategy* of $\frac{7}{12}$ hawk and $\frac{5}{12}$ dove. Or indeed, one could have some mix of strategies in between. All that matters is that the process maps on to the equilibrium mixed strategy at the macro-level. Now if we wanted to write social software to reach that equilibrium we could write it for all to play mixed strategies or for some to play pure dove and some pure hawk.

Computer Programmer: Now it seems that you are saying that social software is something for the brain. Some loaded with pure dove, some with pure hawk and some maybe with a mixed strategy. That is what agent-based modeling can help us solve. We can write programs that map onto cognitive processes and see where that gets us.

Political Scientist: No, I meant the opposite really. We don't need to look into the brains—or minds better—of the players. We can explain the outcomes we get through the macro-level equilibrium. How that equilibrium arises is another question. An interesting one perhaps, but the game-theoretic or social

software explanation does not concern itself with the causal explanation. It simply notes the macro-level algorithms that map onto the outcomes. The analogy is that we don't care how the electric circuits carry our information in computer software (unless there are systematic problems perhaps which can mean that some programs do not run as well as others) and we don't care how people actually behave or think about what they are doing. We just care about the outcomes, and the general processes – the way institutions structure or channel human behaviour into the optimal paths.

Ethicist: I would like to come in here. I've not spoken yet, because I was not sure I could offer anything. But I have been a little concerned about the direction of the conversation from the start. And now I am most concerned. You all seem to think that whilst there might be disagreement about outcomes, once a group of people agree on an outcome, all that matters is coordinating activity so the outcome occurs. Now first, surely even if everyone in a group agreed on an outcome then it might still be wrong. I mean a group of people might all agree that human sacrifice once a year is okay, as long, say, as the victim is chosen by lot. But surely such outcomes are wrong.

Logician: I'm not sure we can say that it is "wrong". Should we judge a community on what they believe if all are in unison?

Ethicist: Okay, look I don't want to get into a debate about objective and subjective values and about cultural relativism, that is not my main point. I just wanted to emphasize that we should not simply accept outcomes as given. It was merely a preface to my main point. I can't accept that once there is agreement on outcomes, the only social software point is the macro-level equilibrium that sustains the outcome. I mean even if we all agree on the outcome, even if the outcome is objectively the right one in some sense, the process by which we get there is important. For example, it matters if some play pure hawk and some play pure dove, rather than all play the same mixed strategy. The one playing pure hawk might be better off than the one playing pure dove. It is all right saying that we can carry some free riders but if some ride free and some do not, then that matters. Some gain and some lose. I mean after all, optimality is one thing, but equality another. We might want an equal society even if it is not socially optimal because some could free ride, say, with no material loss to others. And, after all, there are individual rights as well as socially optimal outcomes in a welfare sense. Can social software model rights? Or another way round, are rights constraints on the kinds of algorithms that are socially acceptable?

Logician: Those are good points. I think we can see rights as constraints on what sorts of algorithms we would want to introduce into a society. I am not sure how rights can be modeled in social software terms. Did Sen not have a proof that welfare maximization and rights are incompatible?

Political Scientist: His impossibility of the Paretian liberal [207]. Sen's result can be seen as viewing society as a social decision function, that is, as a function F that defines possible combinations of individual preference orderings P_1, \dots, P_n to a "social preference" $P = F(P_1, \dots, P_n)$. Whatever the feasibility constraints on the alternatives open to society, F provides the answer to the question of what is best for society by generating the social preference relation R . For any situation the best social state relative to R is chosen from the set of feasible social states.

Logician: Obviously, it is thereby assumed that the social preference relation always contains a best element, right?

Political Scientist: Indeed. Sen defines a condition of minimal liberalism (ML) that is supposed to be a necessary but not a sufficient condition for rights. ML states that there are at least two individuals such that for each of them there is at least one pair of alternatives over which each is decisive. An individual is decisive over a pair of social states (x, y) if it is always the case that if the person strictly prefers x to y , then society should strictly prefer x to y , and conversely, if that person strictly prefers y to x , then society should strictly prefer y to x .

Logician: Okay, so the condition of minimal liberalism, ML , now states that there are at least two individuals who are decisive for some pair of distinct alternatives, that is, there are individuals i and j and distinct pairs of alternatives (x, y) and (z, v) such that $[yP_i x \rightarrow yPx$ and $xP_i y \rightarrow xPy]$ and $[zP_j v \rightarrow zPv$ and $vP_j z \rightarrow vPz]$.

Political Scientist: Indeed. Sen's liberal paradox states that condition ML cannot be satisfied simultaneously with the (weak) Pareto-condition (P) and the condition of Universal Domain (U). According to the Pareto-condition, a strict preference that is shared by all individuals should also be represented in the social preference relation: if all individuals strictly prefer some alternative x to some other alternative y , then society should also strictly prefer x to y . The condition of Universal Domain demands that a social preference relation is generated for any logically possible configuration of individual orderings: no individual preference orderings are excluded a priori. The incompatibility

of the three conditions is easy to demonstrate. Take the simple case in which there are only three alternatives x , y , z and two individuals. Assume the two individuals, i and j , have rights over x and y , and over y and z , respectively. Given U , suppose that individual i strictly prefers y to x and x to z , whilst j strictly prefers x to z and z to y . We then see that i 's rights over x and y implies that yPx and that j 's rights over y and z implies zPy . The Pareto principle yields xPz , which means that we have $xPzPyPx$: the social preference relation is cyclic and hence does not contain a best alternative.

Logician: So what is the upshot of all this?

Political Scientist: Sen's results show that, under the conditions U and P , a person's decisiveness over one pair of alternatives is incompatible with other persons having any rights at all. To see rights in terms of decisiveness over pairs of alternatives means that, in a very strong sense, individual rights are incompatible.

Logician: The incompatibility of rights only follows if the social decision function satisfies U and P . It can easily be shown that if one of these conditions is dropped rights can be made compatible with each other.

Ethicist: I think Dowding and Van Hees [70] showed that for any set of compossible rights the rights are, in their words 'vanishingly small'. Compossibility is the condition that rights can always be satisfied — all rights are "co-possible". But one might see that the acceptance of a set of individual rights is a convention that people abide by, and so the conventional behaviour within that society is composed of algorithms that have rights built into them. We might learn from Sen that respect for such rights might not always be socially optimal in some other welfare sense; but such rights might allow a society to run more smoothly. After all, one cannot even run a market system without some respect for property rights.

Computer Programmer: Can we get back to the levels of analysis discussion? I mean there is nothing incompatible with saying that there are various micro-causal processes that could lead to the same macro-outcome, and then deciding that one is preferable to other. One might even think one preferable to another, but the latter is the one we are stuck with. We simply cannot move to a better equilibrium from the one we are at.

Political Scientist: Kenneth Shepsle, a Harvard political scientist introduced the notion of structure-induced equilibrium into political science [210]. He in-

roduced it in the context of legislatures but it might be generalized. Do you know the McKelvey-Schofield type results [149; 205], where in an n -dimensional issue space with many agents there could be a constant cycling of majority preferred outcomes? Well, Shepsle argued that in those situations one can have equilibria institutionally generated. One just needs to make alterations in the equilibrium costly by having legal constraints and costly actions in order to change them.

Logician: Oh, yes?

Political Scientist: Well, he suggested that the rules by which legislatures are run – the committee systems that are set up; the closure motions on bills, the rules governing how different bills might be bundled together and so on – constitute a set of structures which constrain the space of possible outcomes. Committees for example, can reduce policy space from n dimensions to one dimension, as policy space is carved up into individual issue spaces, so education policy is not bargained over with defense policy. Any bargains struck are struck in the same issue space. He argues that in that way certain equilibria are chosen rather than other ones. There is nothing inherently superior about those equilibria, it is merely that the rules as laid down constrain the possible set [210]. And that is, at least in the short run, of advantage to the whole legislature. Well we might generalize the idea of structure-induced equilibrium to the whole society. And suggest that property rights, welfare rights and so on, are all structures that constrain outcomes.

Philosopher: Didn't Shepsle's mentor Bill Riker point out that the particular rules chosen were probably chosen for the advantage of the majority at the time they were chosen [192]? But as that majority interest changes then so should the rules. You just get a new cycle.

Political Scientist: Yes, but such a cycle might be better in the much longer run; and there are much greater transaction costs of changing rules. I mean, changing a rule like a seniority rule for chair of committees might be against your advantage now, so you want to change it, but you never know it might be to your advantage later.

Logician: Well, it must be to your advantage in the long run, when you get older and more senior!

Political Scientist: Exactly! But even if there is uncertainty over how the rules will operate to your advantage or disadvantage, it might be easier for

you to play by the current rules, than try to change them. You know how the current rules operate. With new rules you might not be aware of all the strategic possibilities.

Philosopher: So again, reasoning about belief and knowledge becomes an important aspect of social software.

Ethicist: And some equilibria might be chosen simply for cultural reasons. I mean some cultures do things one way and others another way.

Logician: But saying it is for cultural reasons does not explain the differences. Saying it is our culture to pick the red pencil and your culture to pick the blue might explain why a particular coordination problem is solved through common knowledge, but it does not explain why those conventions started in the first place [137].

Ethicist: No, but I mean, it does explain the transmission of the coordination. And moreover, it can generate interests. Because it is our culture to pick the red pencils, red pencils are *our* interest, and we will fight you if you say there have to be blue pencils for note-taking at the UN Assembly.

Logician: So both blue and red pencils are provided, and maybe that is more expensive and socially sub-optimal.

Ethicist: But, I mean, such a rule does recognize culturally specific interests and rights. It might be more egalitarian, and that is another consideration. Optimality is not the only thing.

Political Scientist: Sure. Though depending on how you define utility, you might be able to build in those concerns into your social welfare function. But let us not talk about that. I want to get back to the earlier point. Even if everyone knows that some rules might be better than others, getting there could be costly. For example, say property rights are set up so that the members of a small aristocracy own everything. Then it might be advantageous to the vast majority to have a revolution and re-allocate property. But two things. The day after the revolution everyone might be worse off than the day before the revolution. During the re-allocation process, property has been destroyed; crops have been burnt, trade disrupted. Even if the majority are better off in the long- or even medium run, in the short run the revolution has made them worse off. Second, the re-allocation then has to stick. If you have revolutions, too often there is a commitment problem. I won't invest, or lend, or even trade, if I think that next week everything is going to be re-allocated. So

rules are important for governing our behaviour. We might realize the ones we've got are not ideal, but if we are going to change them we had better be careful because the new ones do need to be better; changing them is costly; and we can't keep changing them.

Philosopher: Adam Pzeworski had an argument a bit like that [183]. He wondered why democratic socialist parties never ended up being socialist when they got into power. In the end they always end up supporting capital. He argued that transforming property rights would disrupt capital accumulation in the short run making everyone worse off, even though the vast majority (the working class) would be better off in the long run. The problem for democratic socialist parties is that the short run is longer than the electoral cycle. So if socialist governments pursue such socialist policies they make everyone worse off by the time of the next election, so they lose that election! The leaders then realize this, so in order to stay in power they don't follow socialist policies of re-allocating property rights. Thus social-democratic parties end up supporting capitalism just as much as capitalist parties.

Political Scientist: There is a similar argument in urban political economy over so-called growth coalitions. The argument there is that local politicians always support development for re-election purposes as it gives a buoyant local economy even though a majority of local people would prefer that their communities have stricter zoning laws and not so much development.

Ethicist: These seem to be problems that political scientists have identified. Can social software provide solutions?

Logician: I am not so sure about that. I will not claim that all our problems can be solved by our approach. Just that it can help us examine the issues and evaluate possible solutions.

Chair: Revolutions as the re-allocation of property! I am reminded that when Gordon Tullock was mugged in Rio he described it the next day as "an occasion of the re-allocation of private property." However, this has all been very interesting, but I am afraid we must stop now, or we will miss the presentation of the next paper.

Philosopher: Goodness! That is me, I must dash.