New Issues in Human Rights Statistics

David L. Banks, ISDS, Duke University; banks@stat.duke.edu

Yasmin Said, George Mason University; ysaid99@hotmail.com

Abs: The traditional western conception of human rights has focused upon political and civil liberties. Recent trends are bringing social and economic rights to the fore. This paper considers the kinds of contributions that statisticians might make in this area, and how issues of equity might be framed from a mathematical perspective.

Key Words: discounting; economic rights; social rights; veil of ignorance

1 Introduction

When the statistics of human rights began, the cold war was hot. Most attention focused upon civil and political rights—the right to vote, the right to freedom of religion, the right to citizenship. But there has long been another thread, holding that people have a right to food, education, and medicine. This paper focuses upon the statistical aspects of these social and economic rights.

Obviously, statisticians can play and have played a major role in tracking trends in the standard of living over time. We have participated in model-building for economic input-output models that speak to the mechanisms of improvement. We have helped to invent quality of life measures that are pertinent to this area, and we have documented the impact of war on civilians. And there are new challenges that are arising in this arena—the G8 summit is trying to rescue Africa from poverty, the Millennium Development Fund wants statistical documentation to support applications, and new models of philanthropy require metrics for progress.

Additionally, there are other kinds of contributions that statisticians may be able to make. Some involve decision-theoretic aspects involved in the philosophical foundation of human rights; others involve adjudication of grievances and balancing of competing interests. This paper addresses a subset of those possibilities.

The members of the American Statistical Association are diverse in their backgrounds and broad in their thought. These new opportunities in the social and economic sphere provide the kind of challenge that our profession should welcome.

2 A Brief History

The evolution of human rights seems to move from the general to the specific. The *Declaration of Independence* suggested some basic principles: life, liberty, pursuit of happiness. This perspective got lightly codified in the *Constitution*, and then detailed more particularly in the *Bill of Rights*. Subsequent court decisions have created a large body of highly specific rights.

The U.S. law applies fully only to citizens; foreign nationals and prisoners of war have different status. Such exceptions seem contrary to the initial broad vision in the *Declaration of Independence*, and that discomfort with narrowness was mirrored in philosophical writings by Jeremy Bentham and John Stuart Mill. The slow evolution of their thinking, combined with the world's reactions to the horrors of the Holocaust and a burgeoning internationalism, eventually crystallized in the *Universal Declaration of Human Rights*. This document declared that citizens in all countries that were signatory to the United Nations signatory countries enjoyed common freedoms. Among the thirty enumerated freedoms were political freedoms, such as:

Article 13 (1) Everyone has the right to freedom of movement and residence within the borders of each state. (2) Everyone has the right to leave any country, including his own, and to return to his country.

Article 15 (1) Everyone has the right to a nationality. (2)No one shall be arbitrarily deprived of his nationality nor denied the right to change his nationality.

There were civil freedoms, such as:

Article 4 No one shall be held in slavery or servitude; slavery and the slave trade shall be prohibited in all their forms.

Article 5 No one shall be subjected to torture or to cruel, inhuman or degrading treatment or punishment.

Article 20 (1) Everyone has the right to freedom of peaceful assembly and association. (2) No one may be compelled to belong to an association.

There were economic freedoms, including:

Article 17 (1) Everyone has the right to own property alone as well as in association with others. (2) No one shall be arbitrarily deprived of his property.

Article 23 (1) Everyone has the right to work, to free choice of employment, to just and favorable conditions of work and to protection against unemployment. (2) Everyone, without any discrimination, has the right to equal pay for equal work. (3) Everyone who works has the right to just and favorable remuneration ensuring for himself and his family an existence worthy of human dignity, and supplemented, if necessary, by other means of social protection. (4) Everyone has the right to form and to join trade unions for the protection of his interests.

And there were social freedoms, such as:

Article 1 All human beings are born free and equal in dignity and rights.They are endowed with reason and conscience and should act towards one another in a spirit of brotherhood.

Article 6 Everyone has the right to recognition everywhere as a person before the law.

Article 25 (1) Everyone has the right to a standard of living adequate for the health and well-being of himself and of his family, including food, clothing, housing and medical care and necessary social services, and the right to security in the event of unemployment, sickness, disability, widowhood, old age or other lack of livelihood in circumstances beyond his control. (2) Motherhood and childhood are entitled to special care and assistance. All children, whether born in or out of wedlock, shall enjoy the same social protection.

Article 26 (1) Everyone has the right to education. Education shall be free, at least in the elementary and fundamental stages. Elementary education shall be compulsory. Technical and professional education shall be made generally available and higher education shall be equally accessible to all on the basis of merit. (2) Education shall be directed to the full development of the human personality and to the strengthening of respect for human rights and fundamental freedoms. It shall promote understanding, tolerance and friendship among all nations, racial or religious groups, and shall further the activities of the United Nations for the maintenance of peace. (3) Parents have a prior right to choose the kind of education that shall be given to their children.

Note that in many examples the wording is vague (e.g., Articles 1, 6, and 15). And even when the intent is clear, as with the insistence on the equality of women, this agreement has had little force in most countries.

Western democracies frequently criticized countries that failed to meet these principles, and were criticized in turn for their own failures. During the cold war, the U.S. emphasized political and civil rights; communist nations such as Cuba and the U.S.S.R. were chided on Articles 13, 15, and 20, while Iran and Vietnam were taken to task on Articles 5 and 20, among others. Reciprocally, the United States and other wealthy democracies were condemned for minimizing social and economic rights, such as are found in Articles 23, 25, and 26. Communist nations felt their ethos was more responsive to economic security and social equality.

The ASA Committee on Scientific Freedom and Human Rights was founded in this cold war climate, and its context was framed by human rights efforts led by President Carter, Amnesty International, and Human Rights Watch. The Committee's charge stressed case work on behalf of statisticians. It held letter-writing campaigns to support statisticians and allied scientists who were imprisoned, or denied access to foreign conferences, or who were prevented from emigration. Necessarily, the clientele were highly educated people, and thus tended to live in or be from developed countries. Thus, a great deal of the early attention was upon Soviet statisticians, and the Committee's efforts were congruent with the prevailing emphasis of western democracies upon political and civil freedoms.

Over time, this changed. The casework began to include professionals in Africa, South America, and the Far East. The axis of discussion moved away from communism, and the Committee members sought to define a research role for statisticians in human rights measurement. This led to a fresh view of the role of statistics, and there were long debates on what fell within the purview of the Committee. Questions that arose included whether the systematic exclusion of women from mathematics education was within the scope of the charge; whether providing the public with official estimates of the civilian cost of war was a scientific freedom issue; whether clinical trials conducted in countries where life was cheap posed an abuse of rights.

This evolution led to one of the great research successes of the Committee. Several members and fellow-travelers developed theory and applications for multiple systems estimates of civilian casualties. These methods combined information from different human rights monitors, or gravesite censuses, or other records to create a reconciled estimate of incidents. This enabled estimates and uncertainties for the total number of deaths and the total number of relocations. This evidence from Khosovo was used in the Hague during the trial of Slobodan Milosevic, and it has also supported investigations in Guatemala and East Timor.

This new direction is one of several that relate to social and economic concepts of human rights. There are new opportunities for statistical research, and perhaps even for public influence. The remainder of the paper lays out some of these topics.

3 Agent-Based Theories of Justice

The *Veil of Ignorance* is a gedanken experiment introduced by the philosopher John Rawls (1971). It imagines that if one wants to decide which of two social systems is more just, then one pretends that one's identity is going to be randomly assigned among all the citizens. Thus one cannot predict whether one will be male or female, straight or gay, black or white, rich or poor, sane or mentally ill, healthy or sick, young or old. A sufficiently impartial thinker would then have an objective criterion to compare the societies in terms of the expected quality of life.

Rawls hoped that this device would enable theorists displace themselves from ideology in making cultural judgments. He also hoped that this approach would undermine Utilitarianism, which puts no bound on the degree to which some people's interests can be subordinated to the advantage of the majority, where the advantage is calculated vaguely as some kind of expectation. Therefore, Utilitarianism might permit slavery, provided that the slaves are sufficiently few and that their work for the rest is sufficiently valuable.

From a statistical standpoint, Rawls' approach raises some questions:

- Rawls felt that the exchangeability of identities created a principle (the "Liberty Principle") that a fair social contract must provide minimum guarantees to the least well-off. But a gambler would disagree; the framers of the social contract might risk slavery if the odds were good.

- Rawls deduced that people should have equal assets, with the exception (the "Difference Principle") that some inequality was allowed if it enhanced the welfare of all (e.g., a doctor could make more money than others, since medicine helps everyone). But this appears to permit significant differences in compensation, perhaps comparable to those found in raw capitalism.

- Rawls maintained that this construct led to societies that maximized liberty, creating the largest possible benefit for all. But this is not obvious—perhaps the society is a local mode, and a different set of principles achieves a larger maximum.

Rawls was a philosopher, not a statistician. And we are statisticians, not philosophers, so in the spirit of exchangeability we cannot criticize this work too much. It has been a hotly debated perspective for years, and the arguments have grown more subtle and refined than the synopsis here suggests. See Wellbank (1982) or Baynes (1992) for more details.

Instead of constructing principles from this exchangeability, perhaps a more experimental approach is possible. A number of economists have used agent-based simulations to study the evolution of market behavior (cf. Jackson, 2002; Sallans et al., 2003). Epstein and Axtell (1996) have shown that simple agent rules can create complex behaviors that mimic societies. We believe that this approach can be applied to Rawls' formulation. One could devise agent-based models in which alternative rule sets for model societies play out, and examine the distribution in the quality of life among these.

More specifically, assume that the agents are initially placed at random in a p-dimensional space of latent psychosocial variables. These variables might be interpreted as, say, tolerance, criminality, political activism, and so forth. Each actor would then form friendship links (or enmity links) with other actors according to a probabilistic mechanism; Hoff, Raftery, and Handcock (2002) propose

$$\log \frac{p_{ij}}{1-p_{ij}} = \mu_0 + \alpha_i + \alpha_j + \boldsymbol{\beta}^T \boldsymbol{x}_{ij} - d(\boldsymbol{z}_i - \boldsymbol{z}_j)$$

where p_{ij} is the probability of forming a link from agent i to agent j, α_i is the propensity for agent i to send a link, α_j is the propensity for agent j to receive a link, and $\boldsymbol{\beta}$ weights interaction factors between agents i and j that are captured in the vector $\boldsymbol{x}_{ij}$ (e.g., the first component of $\boldsymbol{x}_{ij}$ might equal one if both agents have the same ethnicity, and be zero otherwise; the second might measure their degree of kinship; the third might measure their difference in social class; and so forth). The $\boldsymbol{z}_i$ and $\boldsymbol{z}_j$ are the locations of the agents in the latent space.

At the first time step, suppose all of the actors form links according this model. Those who are close in the latent space are more likely to become friends than those who are distant, since increasing distance lowers the logit value. At the second time step, the rule is that each agent moves towards the mean of their friends. Then the process repeats. Similar rules would enable agents to disconnect friendship links, with probabilities that depend on distance and history.

Realistic implementation would require a more complex set of rules; most people drop friends less quickly than they add them, and most people only have time for a certain number of close friendships. If a similar model describes enmity ties, then with a little tuning one could get stable groups to form, and to move away from each other in the psychosocial space. Some agents might be highly charismatic, with large values of $\boldsymbol{\alpha}_i$; if they are in the intolerant region of space, they might attract a large group. And when there is an isolated group in a quadrant that is intolerant, criminal, and activist, then one has the social ingredients for oppression.

One nice aspect of the agent-based simulation is that one can allow time-varying covariates, to capture the effects of changing status, education, or aging (for example, older agents can be made less likely to move towards criminality, and slower to change in general). Another nice feature is that can give each agent a slightly different "personality" and thus alter the probabilities of certain kinds of interactions. A third nice feature is that some covariates may depend upon the location in space. Finally, a fourth feature is that different agents can value different goals; this leads to a more complicated model for how the agents move through the latent space. But a key problem is to develop a sufficiently rich class of probability models to describe this complexity.

3.1 Hierarchical Probability Models

The agent model is built around social factors that affect human rights violations. One key element is to divide the population into relatively homogeneous subpopulations which have relatively similar behaviors (or rule sets). These subpopulations would surely reflect social class, gender, and perhaps religion or race, depending upon the application of interest. Even so, it is desirable to allow variability among the members of each subpopulation.

Let D be the domain representing the possible states of an individual; this includes the location in the social space, as well as history and other covariates. As discussed previously, each simulated agent moves randomly in D according to agent-dependent, time-varying transition probabilities. Thus

$$P_a(t) = P_a[\boldsymbol{x}(t) \mid \boldsymbol{X}(s),\ s \in [0, t)]$$

is the probability distribution for agent a at time t on domain D given the entire history of agent a and all other agents up to, but not including, time t. Ideally, one would like to create a full model of the behavior of every agent. However, such a modeling effort entails estimation of far too many parameters relative to available data.

Nonetheless, variability among agents is essential. As a compromise, some agent-based simulations might employ Bayesian hierarchical models. Specifically, one defines a manageably small number of subpopulations of individuals representing different socio-demographic and personality characteristics; for example, one might have a subpopulations of lower-class white males, and then allow agents within that group to have a range of incomes, employment status, and attitudes towards racism, according to some plausible hierarchical model. Individuals within a subpopulation have similar forms for their transition probabilities $P_a(t)$, but will evolve along trajectories that are more different than just random evolution would suggest because their initial covariates are also randomly different.

In many applications the domain D is taken to be discrete; this can be done by coarsening the continuous space in reasonable ways (e.g., counting income to the nearest \$1,000, or having a ten-point scale for attitudes on racism). Similarly, one can coarsen the time domain so that transitions occur at discrete (integer) times. This simplification allows much faster simulation, and there is a natural class of Dirichlet models for the transition probabilities.

To see some of the advantages of this discretization, ignore the full time dependence and focus upon the transition matrix $\{P_a[\boldsymbol{x}(t)|\boldsymbol{X}(t-1)]\}$; its entries are the probabilities that agent a moves to state $\boldsymbol{x}(t)$ when the ensemble of agents is at $\boldsymbol{X}(t-1)$ in the previous time step, for all possible values of the discretized domain. Under the simplified model, this matrix can be chosen randomly from a distribution over all transition matrices. For both modeling flexibility and technical reasons, it is natural to use class-dependent Dirichlet prior distributions for each row of the transition matrix. That is, with $\boldsymbol{X}(t-1)$ fixed, the transition probabilities for $\boldsymbol{x}(t)$ are sampled from the Dirichlet distribution

$$f(\boldsymbol{p}|\boldsymbol{\theta}) = c \prod_{i=1}^{k-1} p_i^{\theta_i}$$

where c is a normalizing constant, k is the number of states in the discretized domain D, $\boldsymbol{p}$ is the row of the transition matrix, and $\boldsymbol{\theta}$ contains the parameters that depend upon the past history of all of the agents. Rapid calculation is possible because the Dirichlet is the conjugate prior for the multinomial; intuitive modeling is possible because researchers have experience with how the $\boldsymbol{\theta}$ terms carry the information.

An extreme alternative to this hierarchical approach is to assume that all individuals in a class have exactly the same transition probabilities. This alternative does not reduce the number of parameters to be estimated, but does reduce significantly the richness of the simulation.

The hierarchical model for discretized transition probabilities requires $k \times T \times n$ parameters, where k is the number of states, T is the number of time steps, and n is the number of subpopulations. This grows rapidly. But there are strategies that make this manageable; one can make assumptions that some states have probabilities that do not depend upon other actors, or that some states have probabilities that depend upon covariates in a simple way, or that some states have probability zero for given values $\boldsymbol{X}(t-1)$. All of these assumptions are plausible for the kinds of human rights simulations under discussion.

3.2 Simulation

In order to understand the qualitative features of agent-based models of human rights behavior vis-a-vis clique formation, hierarchies, polarization, disenfranchisement, and so forth, one needs to run the model many times and summarize the results. The principal steps in such a study are:

1. For each individual agent a, assign its subpopulation. This may be done either deterministically or stochastically.

2. For each agent in a specified subpopulation, stochastically generate its initial location in psychosocial space, its covariates, and and perhaps an initial set of social network links.

3. At each time step, agents form ties (of friendship, enmity, political alliance, religious affiliation, or all of these) according to an appropriate generalization of the model in Hoff, Raftery, and Handcock (2002). Ties from previous steps are unlikely to be broken or changed, but new ties may still form; the instantiation of this kind of complex behavior in agent rules is more of an art than a science and would require experimentation and judgment.

4. Between the time steps, all agents move in the psychosocial space towards those agents with whom they have positive links, and away from agents with whom they have negative links. (More complex rules for movement might ensure that cliques never degenerate into points, that rivalry enters into close partnerships, and so forth.)

5. Also between time steps, all agents move in the covariate subdomain within D (i.e., that portion of the domain that is not the psychosocial space) according to transition probabilities derived from hierarchical models.

6. Repeat steps 3 through 5, and track the behavior of interest (so the formation of an oppressed minority).

The computational effort is, to first approximation, linear in the number of agents and quadratic in the number of states in the discretized domain D. The effort is also linear, obviously, in the number of replications of the simulation.

From a practical standpoint, one would start with small societies and simple rules for the agents. The hierarchical model for transition probabilities can also be made more transparent by making additional convenient assumptions about how probabilities are affected by covariates. For example, one might consider discrete choice models in which an agent moves at time t from state $\boldsymbol{x}(t-1)$ to one of a number of possible states according to a multinomial logistic model with probabilities determined by history and covariates. This leads to transition probabilities of the form

$$P_a[\boldsymbol{x}_j(t)|\boldsymbol{X}(t-1)] = \frac{\exp[\beta' \boldsymbol{x}_j(t)]}{\sum_{i=1}^k \exp[\beta' \boldsymbol{x}_i(t)]}$$

where the $\boldsymbol{\beta}$ is vector of weights on the covariates (and recall that k is the number of possible states). This may oversimplify, but it seems flexible enough to capture many qualitative features of rights behavior; it also allows one to incorporate categorical features, such as race or gender, in a natural way.

4 Investment

Many modern philanthropies want to kicks-tart sustainable growth, rather than making donations that can only prolong intractable situations of appalling misery. This leads to the possibility of comparative cost-benefit analysis of social programs.

One-off cost-benefit analysis can be straightforward. For pennies a day one can vaccinate children. A simple calculation makes it clear that the benefit outweighs the cost. But a more complicated situation arises when one can invest in feeding children, or in educating them, or in vaccinating them. This makes it into a portfolio problem, and the prudent investor must do more mathematics to decide what is best. To make things more difficult, one often has much longer lists of choices, with complex interactions between them (educating a hungry child is not very efficient—they do not learn as well). Also, one has to decide the time frame over which the analysis is done: for a few months, food is probably best; for a decade, vaccination is best; and in the long run, the multiplier effect of education gives the largest return on investment.

Some economists (Sachs, 2005) argue that it is possible to end poverty within 20 years. The world's wealth and technology can surely make large improvements, if resources are wisely directed. But there are statistical issues in how this can be achieved. One strategy is for many small philanthropies to cooperate in small scale designed experiments to determine the effect of different kinds of interventions in different kinds of situations. For example, one might compare education versus infrastructure in countries that have high and low rates of corruption; this simple 2^2 experiment could give guidance on appropriate investment.

Similarly, economists use input-output models (mostly systems of linear equations) to model how resources flow through economic sectors. These models are poorly validated, but could provide important insight into the kinds of interventions that are most beneficial to developing nations. Statistical approaches to model validation are common in engineering; with some thought, data from simple experiments could help pin down input-output models for the economies of developing nations.

In this same spirit of smart philanthropic investment, we analyze the long-standing question of how much money an individual should contribute to charity. Many religions encourage their members to tithe, but this is an approximate approach to a moral problem that deserves serious attention.

Consider the following approach to determining appropriate altruism. First, identify the total need in the world. If there are I nations, suppose the distribution of income in nation i is $f_i(x)$, where x is annual income. Let t_i be the minimum annual income needed to live with dignity in that country. Then the amount needed in that nation is

$$V_i = \int_0^{t_i} (t_i - x) f_i(x)\, dx$$

and so the total need in the world is $V = \sum_{i=1}^{I} V_i$.

Next, we figure out how much those with enough should give to those in want. To do this, first consider the utility function of money for the ith nation. Denote this by $U_i(x, y)$ where x is the amount of money one has, y is an additional amount of money, and $U_i(x, y)$ is how much that additional money is valued. A pauper would value one additional dollar at a dollar, but as wealth increases the additional dollar means less. There is a large body of experimental economics and psychology (cf. Rabin, 2000; Kahneman, 2003) that charts how this utility function works, and it works differently for different people. For our purposes, we assume that all people in a given nation have the same utility function, and we note that this assumption is reasonable but invites discussion.

Using Rawl's perspective, we assert the principle that the burden of altruism is distributed fairly when everyone who contributes should feel the same pain, as measured by their nation's utility function. A citizen of nation i with x dollars of annual income who gives $g(x)$ to help the poor feels the pinch as $U_i(x - g(x), g(x))$. So set this to c, solve for $g(x)$, and one has the basis for determining what it is right to give.

To illustrate this, suppose that everyone values each dollar the same, so that $U_i(x-y,y)=y$. In this case everyone should contribute the same dollar amount, and the amount that is given for a fixed utility value c is

$$W_i(c) = \int_{t_i+c}^{\infty} cf_i(x)\,dx.$$

A different utility function model says that people who donate the same proportion of their income feel the same regret. This is a flat tax rate scenario, and the utility function is $U_i(x-cx,cx)=c$. In this case the amount that is given for a fixed utility c is:

$$W_i(c) = \int_{(1+c)t_i}^{\infty} cxf_i(x)\,dx.$$

But experiments suggest a more complicated form for the utility function. Very wealthy people seem willing to tolerate giving a larger percentage of their wealth, and this is partially reflected in the graduated tax scales used in many countries.

From our standpoint, all that is needed is a utility function that is broadly accepted as sane by the standards of the nation. For that function, the amount of capital that is raised is $W_i(c)$. To end world poverty, one just solves

$$V = \sum_{i=1}^{I} W_i(c)$$

for c. This determines how much each person should give, in order that all feel the same inconvenience. The equation is easy to solve numerically, since each $W_i(c)$ is monotone in c.

5 Discounting Grudges

What is an appropriate model for discounting the past? During the disturbances in Khosovo, both sides pointed to massacres that the other had committed in the fifteenth century. And Sein Fein points to English abuses in Ireland stretching back Henry VII. To an outsider, it seems psychotic if politicians reach back 500 years to inflame their people.

But jurists may see this as a continuum. The heritage of African-American slavery still lingers, and one can make a reasonable argument that compensation is due. There is still living memory of the genocide of Armenians by the Young Turks in 1915-16, but it is slipping rapidly off the table of active history. Many Palestinians living today hold valid legal deeds to property in downtown Jerusalem—how angry do they deserve to be, and what recompense is right? More topically, how can an international court calculate what a country owes when it is ultimately judged to have waged an unjust war? The Treaty of Versailles levied a penalty of 6,600 million pounds against Germany at the end of World War I, but the statistician John Maynard Keynes held that the penalty was too high and made World War II inevitable.

There are some relevant data points. The German government pays compensation to survivors of the Holocaust. The United States is paying the survivors of the Japanese internment camps. And recently the xx court refused to pay reparations to African-Americans for slavery, on the grounds that no one directly affected is still alive. Each of these decisions represents a some kind of negotiated agreement that, fairly or unfairly, balances the rights of victims against the assets of government authority. These data could be used to fit curves for the degree to which claims get discounted with time. In practice, we expect that the estimated function will be extremely noisy, in part because there is little data and in part because different countries and courts have different attitudes. Of course, the calculation should be made in terms of standard dollars, to correct for changes in value over time.

A different approach to the problem of discounting past injuries lies in agent-based models (again). Axelrod (1984) created a competition he called the "Iterated Prisoner's Dilemma" in which large numbers of agents competed against each other, over and over again. The simulated competition led to surprising insights about cooperation and may enable estimation of the cost of holding a grudge.

The Prisoner's Dilemma is a game theory problem that took its name from the following hypothetical situation: Suppose Abelard and Heloise are suspected of having jointly committed a crime, but there is insufficient evidence to convict. The police isolate them both, and separately offer each a deal. If

- Abelard confesses but Heloise does not, then he testifies against her and plea bargains down to no time while she gets twenty years;
- if Abelard confesses and Heloise does too, then both get ten-year sentences;
- if Abelard does not confess and neither does Heloise, then both get sentences of five years on trumped-up charges;
- If Abelard does not confess, but Heloise does, he is sentenced to twenty years.

Abelard's dilemma is whether to rat on Heloise. In game-theoretic terms, his payoff matrix is:

	Heloise Confesses	Heloise Does Not Confess
Abelard Confesses	-10	0
Abelard Does Not	-20	-5

The minimax solution to this game is for both parties to confess. But this does not account for the possibility of repeated play.

Axelrod's competition allowed agents to have different rules on whether to confess or stand silent (defect or cooperate). The agents played each other repeatedly, choosing their partner at random from the pool of all partners. Each agent got a payoff according to the table above (actually, Axelrod added 20 points to each cell, since it was more convenient to talk about gains than losses, but this does not change the structure of the problem). The rules for the agents were generated by 63 experts in decision theory, psychology, and economics.

The winning rule set was submitted by Anatol Rapaport, a psychologist at the University of Toronto. His rule has been called "Tit-for-Tat" and specifies that the agent should start by cooperating, and thereafter do whatever the partner agent did in their previous round of play. Thus if Rapaport's agent plays a copy of itself, it always cooperates (does not rat) and gains 15 points. If it plays an agent that defected (ratted) the last time it played, then Rapaport's agent will defect and gain just 10 points. Rapaport's rule set was the shortest of all entries, just five Fortran lines, and it won decisively.

Axelrod published the results of the competition, and invited all 63 experts to resubmit. Rapaport did not change his code, but everyone else did—as before, Rapaport's agent was the clear winner. Axelrod was intrigued, and altered the simulation so that agents which were more successful made duplicate copies of themselves in the pool of agents. He found that if he created an pool of cooperators and seeded it with defectors, the defectors gradually failed and disappeared. And if he created a pool of defectors and seeded it with cooperators, then again the defectors gradually failed and disappeared. All of this is quite surprising, from the perspective of conventional minimax theory.

Axelrod's study prompts the provocative possibility that free and fair trade, the rule of law, and human rights may create a stable equilibrium. If so, that would be wonderful news for human rights workers. Perhaps more study can elucidate this issue. It also suggests that populations of defectors are at an unstable equilibrium, and that a few cooperators will lead to total change.

From this perspective, Rapaport's agents need to learn to forgive. If they are placed in a society of defectors, they will all defect too, and no improvement is possible. But if, over a span of time, they are able to set aside an old grudge, then the population of defectors will eventually get replaced by a population of cooperators, and the net payoff will steadily increase. Agent-based simulations of this kind can explore the speed with which Tit-for-Tat agents need to forgive in order to maximize the net gains in payoff. And this curve will show how much it costs to hold a grudge as a function of time.

6 Conclusions

Statistical thinking can address new problems in the evolving theory of social and economic rights. This paper has explored some of the nonstandard things we can contribute, suggesting:

- ways in which human rights behavior can be simulated, in the context of a theory of justice based upon exchangeable agents;
- a procedure for determining how to fairly distribute the burdens of charity among the well-to-do and wealthy;
- a methodology for studying the costs of compensation as a function of time, and thus appropriately discounting past injuries.

Each of these topics is to be taken lightly—neither of the authors is an expert in these areas. But we think the problems are important, and that statisticians have a unique perspective on some of these issues.

But we do not mean to minimize the other important, and more traditional, work that statisticians can do in this arena. It is natural for us to develop procedures for tracking trends over time in quality of life indices, poverty measures, and so forth. Similarly, we can analyze comparative data from developing nations to identify factors that accelerate or retard growth. For example, there is much current discussion over the effect that official corruption has on the efficacy of public and private charities, and this impact could probably be quantified. And there are longstanding debates on what kind of interventions do the most to help the poor—most agree that ultimately one wants to invest in education, but the immediate need may be for food or medicine or security. This kind of tradeoff can be formulated as an optimization problem under large uncertainties, so statisticians have a role to play.

And that is the main message. Statisticians have a role to play. We have done it before and are doing it now, and there is good reason to hope that continued effort will bring fresh successes.

References

[1] R. Axelrod, *The Evolution of Cooperation*, Basic Books, New York, N.Y. (1984).

[2] K. Baynes, *The Normative Grounds of Social Criticism: Kant, Rawls, Habermas*, SUNY Press, Albany, NY (1992).

[3] J. M. Epstein, and R. L. Axtell, *Growing Artificial Societies*, MIT Press, Cambridge, MA (1996).

[4] P. Hoff, A. Raftery, and M. Handcock, *Latent space approaches to social network analysis*, Journal of the American Statistical Association, 97 (2002), pp. 1090-1098.

[5] M. O. Jackson and A. Watts, *The evolution of social and economic networks*, Journal of Economic Theory, 106 (2002), pp. 265-295.

[6] D. Kahneman, *A psychological perspective on economics*, American Economic Review Papers and Proceedings, 93, pp. 162-168.

[7] M. Rabin, *Diminishing marginal utility of wealth cannot explain risk aversion*, in *Choices, Values, and Frames*, Daniel Kahneman and Amos Tversky, eds., Cambridge University Press, New York, N.Y. (2000), pp. 202-208.

[8] J. B. Rawls, *A Theory of Justice*, Harvard University Press, Cambridge, MA (1971).

[9] J. Sachs, *The End of Poverty: Economic Possibilities for Our Time*, Penquin Press, New York, N.Y. (2005).

[10] B. Sallans, A. Pfister, A. Karatzoglou and G. Dorffner, *Simulation and validation of an integrated markets model*, Journal of Artificial Societies and Social Simulation, 6 (2003), pp. 1-33.

[11] J. H. Wellbank *John Rawls and His Critics: An Annotated Bibliography*, Garland Publishers, New York, NY (1982).