Complex Adaptive Systems: Energy, Information and Intelligence: Papers from the 2011 AAAI Fall Symposium (FS-11-03)

Values and Evaluation in Game Theoretical Models

Christopher M. Caldwell and Majid Amini
Virginia State University
Dept. of History and Philosophy, Box 9070
Virginia State University, VA 23806
ccaldwell@vsu.edu
mamini@vsu.edu

Abstract

The paper provides an investigation and identification of the values and assumptions that influence the applications and results derived from the applications of game theoretical models to solve logistical problems. The ultimate goal is to identify and begin to eliminate prejudices and blind spots within the formulation of the game theory constraints and the application itself. We explore the investigation and identification of blind spots through the application of game theoretical models to solve problems concerning resource allocation. In particular, the paper investigates the application of game theoretical models to allocate resources in crisis situations.

Introduction

The paper provides an investigation and identification of the values and assumptions that influence the applications and results derived from the applications of game theoretical models to solve logistical problems. The ultimate goal is to identify and begin to eliminate prejudices and blind spots within the formulation of the game theory constraints and the application itself. We explore the investigation and identification of blind spots through the application of game theoretical models to solve problems concerning resource allocation. In particular, the paper investigates the application of game theoretical models to allocate resources in crisis situations.

Fairness in Resource Allocation

Concepts such as max-min fairness, utility max-min fairness, and proportional fairness are used in attempts to answer resource allocation questions. In particular, such concepts are used to attempt to quantify the quality of fairness in terms of utility, more broadly construed. Such projects and applications encounter various problems and limitations which must each be addressed in the formation of a more complete decision-making process through the application of game theory to problems such as resource allocation in crisis situations. The following discussion highlights only a few such problems.

The initial problem encountered is the ambiguity of the notion of fairness itself. As is well known, there are multiple notions of fairness. The concept generally employed in discussions of resource allocation is some form of distributive fairness as opposed to a form of retributive fairness, for example. However, any model for allocating resources must provide justification for the preference of a form of distributive fairness over some form of retributive fairness. Indeed, such a justification can be easily provided by appealing to the goal of determining how to distribute the resources available in a crisis situation over determining what each individual deserves in some retributive sense. However, why this is the goal instead of using the crisis situation as a way to teach individuals a lesson in some retributive sense must be justified. Hypothetically, if a natural disaster hits an area, one must decide whether to first help those who are in the most dire situations due to their own lack of preparedness, or whether one should use this opportunity to teach those individuals a lesson and help others first, thus pursuing the goal of retributive justice and not the goal of distributive justice. It might be difficult to ascertain the deserving victims, i.e., those who took precautions, as opposed to the undeserving victims, i.e., those who failed to prepare. Game theoretical models are used to tackle difficult tasks and assist with the decision-making process in crisis situations, so it cannot be argued that just because a task is difficult that the game theoretical model ignore it. Further, detailed justifications and arguments cannot be provided in every paper which presents a game-

theoretical model. However, the application of game theoretic models without the acknowledgment that such value judgments are in fact playing a role in the application of the model leaves the discussion of the game theoretical model in question incomplete and lacking adequate conceptual foundations. An acknowledgement of the theoretical assumptions concerning the concepts being employed at the level of the kind of fairness being pursued is necessary. Importantly, it appears that most applications of such game theoretical models are done not only without the recognition of such assumptions, but without realizing the significance of the assumptions which have been made in describing the problem to be addressed.

The next problem encountered is in the lack of clarification of the notion of utility itself that is the very foundation for attempts to quantify the quality of fairness. The notion of utility has many forms and can be measured in many different ways. Classically, the notion of utility has been connected to happiness simpliciter, or pleasure. However, contemporary consequentialists have developed the notion in more nuanced ways to attempt to deal with various problems. For example, the notion of utility has been defined in terms of satisfying an individual's preferences. Such a move requires a further determination of whether all preferences are equal in such a calculation or whether some preferences themselves should be preferred over others due to other considerations, e.g. the rationality of the preferences under consideration. Additionally, as one attempts to satisfy preferences, one must determine whether or not one is going to attempt to satisfy the greatest number of preferences or whether there are some preferences that are worth more than others. For example, should preferences concerning long-term well-being be privileged over preferences for short-term comfort in a crisis situation? This question is exacerbated by the uncertainties surrounding survival itself in certain crisis situations. The problem is that these difficulties are generally ignored in applications of game theoretical models in a way which we find to make the applications incomplete in significant and serious ways.

As an example of such problems, consider "Multievent Crisis Management Using Noncooperative Multistep Games" by Gupta and Ranganathan. The paper is motivated to address the meaningful, important, and significant problem of allocating and scheduling response units in a crisis situation concerning multiple events. In the paper, they apply both utility max-min fairness and proportional fairness very quickly. The paper argues that the application of both of these concepts is often inadequately developed and employed. However, Gupta and Ranganathan also employ evaluations outside of those of fairness. They explicitly make evaluative judgments concerning feasibility, criticality of the event, and optimal allocation without the necessary care or exposition in order to provide an adequate justification for those evaluations. The ultimate problem is apparent in the following paragraph, which occurs early in the paper:

> The uniqueness of the multiemergency scenario lies in the fact that the modeling of the problem needs to pivot around the objective of socially optimum and "fair" allocation because, even if some event is of low criticality, it has the potential to spawn more crisis in case of delayed service. As an example, a road accident with no major injuries is a low criticality event with a requirement of three or four police cars for traffic redirection. However, if the resources are not available on time, then the traffic congestion may give rise to a disastrously critical crisis situation like delay of service to other crises. (578)

Multiple difficulties are immediately apparent. First, the separation of social optimality and fairness reveal the assumption that the two concepts are somehow different. However, a utilitarian conception of fairness would not necessarily accept this split. A result that is socially optimal would be fair and a result that is fair would be socially optimal for at least some utilitarians. Secondly, note the evaluative judgments being made throughout the passage without the acknowledgement of those judgments. The following judgments are all evaluative in nature and impact the analysis in important ways: that an event is a crisis, the level of criticality of the crisis, the timeliness of the availability of resources, the significance of the injuries sustained, the level of traffic congestion requiring the additional police cars, and the disastrous situation of delay of service mentioned at the end of the paragraph. The point here is not that all such evaluative judgments must be fully developed in every analysis, which is clearly impossible. Instead the focus here is upon the effects of such evaluations upon analysis and game theoretical applications as well as exploring the significance upon them within such applications.

The Effects of Evaluative Judgments and Concepts

Unacknowledged and unjustified evaluations, and the effects those evaluations have on game theoretical model applications, illustrate a clear weakness of such applications. Although the attempt to quantify fairness is a common type of evaluation method, as mentioned previously, other evaluations are clearly evident in these attempts to determine how to allocate resources. Importantly, some of these evaluations are used to change game theoretical models in significant ways by removing constraints related to moral evaluations if those morally significant evaluations turn out to cause inconvenient results. The goal in this section is to discuss the effects of evaluative judgments and concepts within such concepts to identify some of the key shortcomings of such applications. The goal,

ultimately, should be to better account for these difficulties in future models to minimize these problems as much as possible. However, overcoming the difficulties is not possible until the extent of the difficulties is recognized by those applying the models. Thus, if a model results in a distribution of resources that is determined to be fair but not to be feasible, then the limits employed in the model are changed to aid in gaining results that are feasible in the relevant way, regardless of whether the results continue to be fair.

To continue with the primary example, Gupta and Ranganathan also employ evaluations outside of those of fairness. They explicitly make evaluative judgments concerning feasibility, criticality of the event, and optimal allocation without the necessary care or exposition in order to provide an adequate justification for those evaluations. The consequences for the success of the model are significant.

Concerning feasibility, Gupta and Ranganathan state the following, "In a multicrisis scenario the overall optimization is feasible only if each crisis has been satisfied with respect to all other crises." (579) In this passage there is an attempt to explicitly identify what constitutes feasibility. Although there is an attempt to define feasibility, the definition simply relies upon other evaluative judgments lacking the necessary and sufficient conditions for fulfillment. In particular, the suggestion that feasibility relies upon the satisfaction with respect to all other crises depends upon the evaluative judgment of satisfaction and the ability to determine this in relation to all other crises. If one is attempting to come to a quantifiable answer, such immense and underdetermined qualitative judgments are clearly problematic.

Even more significantly, consider the statement that any allocation vector is fair if, for any other "feasible allocation…the aggregate of proportional changes is nonpositive in terms of utility and is nonnegative in terms of the total cost of allocation." (587) The very notion of feasibility is thus central to the judgments of fairness being made. The evaluative claim of feasibility, which is problematic as previously stated, is thus embedded in the notion of fairness in a manner which undermines the clarity and even the determinacy of the notion of fairness being employed in the application of the game theoretic model in this instance.

The judgment of the criticality of the event is of central importance when determining how to allocate resources in a crisis situation. Gupta and Ranganathan are right to acknowledge this and attempt to say something about criticality itself. The problem is the lack of specificity and development in their notion of criticality. They state that there is to be a criticality level "associated with each crisis which decides the sequence in which the games are to be played." (583) The details of how this evaluation is to be made are not clearly articulated. The judgment itself is embedded in an algorithm, which highlights our concern. Evaluative concepts are working at a level of importance not recognized by those applying the theoretical models to solve problems. The mathematics in the models may be perfect, but this underlying problem undermines their ability to provide useful and applicable results. This is a problem to which we will return later in the paper.

The notion of the optimality of allocation is central to the very issue of resource allocation in a crisis. What one wants is a quantifiable answer from a method for determining optimality. The idea is that the model in question will provide such an answer. The central question, then, concerns whether a qualitative evaluative judgment such as optimality can even be quantified. However, for present purposes, we will put that question aside and simply assume that such quantification is possible. Significant problems remain. The kind of optimization will clearly change the results. For example, if one is using a max-min model of measurement or if one is using a proportional method of measurement, one will come to different conclusions concerning optimality. Thus, depending upon the judgment concerning the optimal type of measurement to be used, the results will clearly be affected. Even if one uses different types to run the model in different ways, then the underlying judgment remains to be made at the moment of implementation if the results of the models run. If the goal is to actually use these models to guide the responses in crisis situations, which we take it to be, then the particular evaluative judgment concerning optimality itself is playing a central role and needs to be much more carefully examined and explored.

Ultimately, in some cases, the definition of morally significant evaluations, such as fairness, are altered in order to meet constraints such as feasibility, but the revised definition is problematic given that the evaluative nature of the constraints themselves are overlooked. The goal is to come to a course of action that is fair, in some sense, but when the model comes to a result which seems fair in the relevant sense then some other consideration overrides the result. The problem occurs when the evaluative nature of the overriding consideration is either overlooked, ignored, or buried in some fashion. The difficulty is that their attempt at quantification is undermined due to the incomplete recognition of the evaluative judgments which are being made throughout the process.

The point here is not that these problems cannot be overcome nor that such models are not of significant use and importance. The point is simply that there are some unstated assumptions which are more complicated and need to be more significantly addressed if the move to quantify the seemingly qualitative judgments necessary in coming to result in how to allocate resources in crisis situations is to be successful.

Theory Application Problems More Generally Considered

Hitherto, game theoretical models have been assayed and assessed from an *ethical* perspective with the primary purpose of appropriately pruning the models to overcome pitfalls and shortcomings hindering their applicability. Yet, there are some game theory practitioners, let alone detractors, that go even further in their ethical criticism of models by arguing that the "confusion between 'is' and 'ought' in game theory is widespread and is a serious obstacle to developing the theory along more productive lines" to the point that it is astounding that so many years of elegant mathematical "theorizing and experimentation could produce so little of value in instructing people on how they *should* behave in conflict situations and in predicting how they *do* behave in conflict situations." (Kadane & Larkey, 1370)

However, in addition to ethical concerns, there are also *epistemological* problems that seem to have contributed to the difficulties faced by game theoretical models in real-life contexts. In this respect, there seems to be an uncanny correspondence between game theory and its mathematically inspired elder sibling, *viz.* econometrics, where similar complaints and criticisms have been raised. Despite the presence of sophisticated and intricate formalizations, as Milton Friedman once complained, the extensive application of mathematics in econometrics "is often used to impress rather than inform." (36) On the incongruence between modeling and application, Coddington, amongst others, points to the difficulty of investing the formal framework with an application by noting that the framework is syntactical and fails to offer any rules for transmitting the internal precision of the syntax to the semantic problem of application. (550-3) Thus, the rest of this brief section is concerned with some of these epistemological issues that game theory modeling has to contend with in the pursuit of practical applicability.

There are a number of standard epistemological problems raised against game theoretical models such as: (i) they are too hard to use, (ii) they form an incoherent collage suggesting no general principles, (iii) they are hard to test, and (iv) they can explain anything. (Camerer, 1991) However, the focus here is narrowed onto the notion of rationality *vis-à-vis* equilibrium.

Equilibrium game theory has borrowed its rationality precept from neoclassical economics, and by deploying the method of backward induction and subgame perfection it has attempted to construct models of social interaction for various contexts. According to the theory, in any such social game there is a *unique* outcome provided that each player is entirely rational, and then the behavior of each agent is tracked through backward induction whereby the model attempts to generate a *unique* equilibrium set of beliefs that allows each player to arrive at a conclusion about the best course of action. But, as Sugden (1991) shows, backward induction is fraught with difficulties. Basically, the problem here is that the model ignores the real possibility of an agent deciding not to follow the rules of the game and at the same time not incur any accusation of irrationality as axiomatically assumed by the theory. In other words, the theory imposes a very rigid sense of rationality and fails to allow *counterfactual* reasoning on the part of one or more players in the game. As Varoufakis notes, thinking "about the possibility of defying the theory that is supposed to govern one's behaviour, is a uniquely human capacity. It is also a capacity that makes the life of the social scientist inordinately demanding. To disallow counterfactuals within a theory … is to ask for serious trouble since human rationality has the bad habit of instructing agents to ask, 'what if I do not obey the theory's rules?'" (400)

Generally, it seems as if counterfactual reasoning is both rational and irreconcilable with equilibrium game theory. By relaxing the stringent conception of rationality as imposed by game theory, one can easily see that some of the game's players may rationally come to different conjectures which may after all turn out to be the actual choice of strategies, obviously not predicted by the other party. (Bernheim, 1984) Players may engage in attempts such as outmaneuvering and embark upon thoughts that will effectively engender the impossibility of securing symmetry and instead lead in the long run towards what Skyrms (1990) calls deliberational disequilibrium.

The fundamental problem here may be traced back to an endemic confusion over the concept of equilibrium. On the one hand, one may identify an *ontological* reading of equilibrium whereby the term is intended to denote an aspect of the actual reality that the theorist is attempting to understand, describe, explain, or represent, whereas a second *theoretical* rendering of it attempts to portray a property of a mathematical-deductive model formulated apparently with the intention of explaining, representing, or somehow increasing our understanding of social reality. But as Lawson points out "not only are these two conceptions distinct … but also, in practice, they have little bearing on each other." (455) For, the theoretical notion does not bear a one-to-one correspondence to the agents of the ontological conception where each player holds his or her own beliefs, desires, attitudes and other intentional states – states that through being intentional manage to introduce *non-uniqueness* into the framework. As in the case of econometrics, Hoover notes, "Economic values and motivations may be powerful, but other – higher – values and motivations may dominate in particular cases and mask the economic tendencies." (48) Moreover, the problems are further compounded when theoretical models of equilibrium are constructed by incorporating concepts and notions like optimization that do not seem to have an equivalent counterpart in the ontological variety. (Chick, 1998) Thus, until and unless, econometric or game theoretic models manage to capture the non-uniqueness element of human

decision-making and refrain from belaboring human behavior with conditions such as of optimization, the models are no more than self-referential constructs, even though mathematically very elegant and sophisticated.

Acknowledgment

The authors would like to thank the Office of the Provost at Virginia State University and the Provost/Vice President for Academic Affairs Dr. W. Weldon Hill for the funding which supported this research.

References

Bernheim, D. 1984. Rationalisable Strategic Behaviour. *Econometrica*. 52: 1007-28.

Camerer, C.F. 1991. Does Strategy Research Need Game Theory? *Strategic Management Journal*. 12: 137-152.

Chick, V. 1998. On Knowing One's Place: The Role of Formalism in Economics. *The Economic Journal*. 108 (451): 1859-1869.

Coddington, A. 1975. The Rationale of General Equilibrium Theory. *Economic Inquiry*. 13: 539-58.

Gupta, U., and Ranganathan, N. 2007a. Multievent Crisis Management Using Noncooperative Multistep Games. *IEEE Transactions on Computers*. 56 (5): 577-589. 1991.

Friedman, M. Old Wine in New Bottles. *The Economic Journal*. 101 (404): 33-40.

Hoover, K.D. 2002. Sutton's Critique of Econometrics. *Economics and Philosophy*. 18: 45-54.

Kadane, J.B., and Larkey, P.D. 1983. The Confusion of Is and Ought in Game Theoretic Contexts. *Management Science*. 29 (12): 1365-1379.

Lawson, T. 2005. Reorienting History (of Economics). *Journal of Post Keynesian Economics*. 77 (3): 455-470.

Skyrms, B. 1990. *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.

Sugden. R. 1991. Rational Choice: A Survey of Contributions from Economics and Philosophy. *The Economic Journal*. 101: 751-85.

Varoufakis, Y. 1993. Modern and Postmodern Challenges to Game theory. *Erkenntnis*. 38 (3): 371-404.