GAMES, GOALS, AND BOUNDED RATIONALITY

by

Leigh Tesfatsion

Discussion Paper No. 76-63, January 1976

Center for Economic Research
Department of Economics
University of Minnesota
Minneapolis, Minnesota 55455

ABSTRACT

A generalization of the standard n-person game is presented, with flexible information requirements suitable for players constrained by bounded rationality. Strategies (complete contingency plans) are replaced by "policies," i.e., end-mean pairs of candidate goals and "controls" (partial contingency plans). The existence of individual objective functions over the joint policy choice set is axiomatized in terms of primitive preference and probability orders. Conditions are given for the existence of pure policy Nash equilibrium points in n-person games, and pure policy Nash bargaining and equilibrium threat solutions in 2-person policy games. Connectedness of the policy and payoff sets is not required.

GAMES, GOALS, AND BOUNDED RATIONALITY*

by

Leigh Tesfatsion

1. INTRODUCTION

In standard games the choice set for player i is assumed to be a collection of "strategies," i.e., functions which assign to each of player i's information sets S_i^j one of the arcs which follows a representative vertex of S_i^j (see Owen [7]). Strategies are thus complete contingency plans for playing the game at hand. Once a strategy has been chosen by every player (possibly including a strategy "chosen by chance"), a unique outcome for the game is determined.

The usefulness of the standard game theory framework for real-world problems is somewhat limited. In actual problem contexts the specification of available actions in the form of complete contingency plans is often not feasible. Information may necessarily be incomplete; alternatively, the required calculations may be too costly. Players in real-world games generally plan in advance for only a limited number of moves. Secondly, the implicit requirement that the chance strategies be defined independently of the other players' strategy choices often imposes an awkward formulation on real-world problems.

In this paper a "policy game" (p-game) is presented, with flexible information requirements suitable for group decision problems constrained by bounded rationality. The players are allowed to specify their available actions in the form of partial contingency plans ("controls"). Their choice sets are assumed to be collections of end-mean, candidate goal-control pairs ("policies"). The candidate goals

---

*Research underlying this paper was supported by National Science Foundation Grant GS-31276X.

are operationally interpreted as potential objectives (e.g., market share aspiration levels) whose realization the players can attempt to achieve by appropriate choice of control. Chance strategies and consequences (e.g., monetary payoffs) are subsumed into "state flows" over which the players' policy-conditioned preference and probability orders are both defined. Thus, in a manner to be made precise below, chance strategies need not be defined independently of the players' policy choices. The existence of objective (expected utility) functions representing the players' preferences among joint policies is axiomatized.

For games with only one player, the p-game reduces to the policy model (goal-control model) formulated in Tesfatsion [9]. As shown there, the Savage expected utility model, the Marschak-Radner team model, the Bayesian statistical decision model, and the standard optimal control model can be viewed as special cases of the policy model. Similarly, it is shown in this paper that the standard n-person game in normal form can be viewed as a special type of p-game.

The importance of explicit, nontautological[1] goal specification in problems of individual choice is discussed and illustrated in Tesfatsion [9]. Briefly, specified goals and controls play distinct strategic roles in many decision problems. Moreover, goals can be important for the feedback evaluation of chosen policies; i.e., the utility or cost of a chosen policy may be a function of the "distance" between the realized outcome and the target[2] outcome (goal) specified in the chosen policy.

In game theory the importance of explicit, nontautological goal specification is potentially even greater. As Aumann [1] notes, for a standard n-person game the implications of a particular coalition structure (partition of the player set) are often quite clear; but what is often not clear is why a particular coalition structure should initially form. For p-games, with nontautological goals explicitly introduced, it seems reasonable to predict that players with complementary candidate goals will form coalitions. Thus certain initial restrictions on the players' candidate goal sets might in some cases be used to restrict the number of coalition structures that can "rationally" form, prior to any consideration of imputed coalition values. Indeed, real-world players free to cooperate but constrained by bounded rationality would probably not bother to impute values for coalitions containing players with highly conflicting candidate goals.

The organization of this paper is as follows. In section 2 the p-game is presented and discussed in normal form; i.e., with individual objective (expected utility) functions for the players simply assumed given. In section 3, drawing on results from Tesfatsion [10], it is shown that necessary and sufficient conditions exist for "primitive" p-games with finite state flow sets to have this normal form. It is also shown that a standard n-person game in normal form can be viewed as a special type of axiomatized primitive p-game in normal form.

In section 4, drawing on results from Tesfatsion [11], conditions are given for an n-person p-game to have at least one pure policy Nash equilibrium point. Also, drawing on results from Tesfatsion [12], a broad class of 2-person threat p-games for which a unique pure policy Nash bargaining solution exists is characterized in terms of three, simple, empirically meaningful properties of the joint objective function: compact domain, continuity, and "corner concavity." Conditions are also given for the existence of at least one pure policy Nash equilibrium threat solution. As will be discussed further below, pure policy sets corresponding to p-games must realistically be allowed to be disconnected if the candidate goal sets are nontrivial. The existence results in section 4 are nontrivial extensions of standard game theory in that connectedness of the policy and payoff sets is not required.

In section 5 an example is given of a primitive 2-person p-game with expected utility representation as in section 2.

2. THE POLICY GAME

Let $G = \{g,\ldots\}$ be a set of (*joint*) *candidate goals*, and for each $g \in G$ let $\Lambda(g) = \{\lambda(g),\ldots\}$ be a set of *(joint) controls*. Then the components for an *n-person policy game (p-game) in normal form* are given by

$$(1) \qquad (n^*; \Theta; \Pi_{i\varepsilon n^*} U_i : \Theta \to R^n)$$

where: $n^* \equiv \{1,\ldots,n\}$ is the *player set*; $\Theta = \{\theta,\ldots\} \equiv \cup_{g\varepsilon G}(\{g\}\times\Lambda(g))$ is the *joint policy choice set* consisting of candidate goal-control pairs *(joint policies)*; and for each $i\varepsilon n^*$ $U_i:\Theta\to R$ is the *objective function for player i*. The function $\Pi_{n^*}U_i:\Theta\to R^n$ will be referred to as the *joint objective function*.

In the special case where G is a cross-product $\Pi_{i\varepsilon n^*}G_i$ of *individual candidate goal sets* $G_i=\{g_i,g_i',\ldots\}$, and, for each $(g_1,\ldots,g_n) \varepsilon G$, $\Lambda(g)$ is a cross product $\Pi_{i\varepsilon n^*}\Lambda_i(g_i)$ of *individual control sets* $\Lambda_i(g_i)$, the p-game will be called *free*. For free p-games there is a natural identification between the policy choice set $\Theta \equiv \cup_{g\varepsilon G}(\{g\}\times\Lambda(g))$ and the cross-product $\Pi_{i\varepsilon n^*}\Theta_i$, where $\Theta_i \equiv \cup_{g_i\varepsilon G_i}(\{g_i\}\times\Lambda_i(g_i))$. The sets Θ_i will be referred to as *individual policy choice sets*.

The candidate goals $g \varepsilon G$ can be operationally interpreted as potential objectives (e.g., production quotas) whose realization the players can attempt to achieve by appropriate choice of control. The controls can be operationally interpreted as possibly conditioned sequences of actions (i.e., partial contingency plans) under the joint control of the players at the time of their choice. The grouping of the controls into sets $\{\Lambda(g) \mid g \in G\}$ reflects the possibility that different controls may be relevant for different goals.

Each player is assumed to have full knowledge of the joint objective function $\Pi_{n*}U_i : \Theta \to R^n$. The objective of the p-game is assumed to be the selection by the n players of a joint policy (candidate goal - control pair) in the policy choice set Θ which is "optimal" in some sense with respect to this joint objective function.

The game is fully specified only when "optimal" is defined and rules are given concerning allowable cooperation. These rules must in turn be combined with specific assumptions concerning the structure of G and the control sets $\Lambda(g)$, $g \in G$. For example, by itself the formalization of the p-game in (1) implies that total cooperation is needed. On the other hand, a free p-game may be played without any form of cooperation (each player i may independently choose an individual policy in Θ_i). Hybrid games are also possible, e.g., cooperative choice of a goal and noncooperative choice of controls. A certain group of environmental protectionists may jointly agree on a platform (goal) for their organization; at the same time they may heartily disagree (act noncooperatively) in voting for a particular presidential-vice presidential ticket as an instrument for achieving this platform.

3. AXIOMATIC APPROACH TO THE POLICY GAME

The normal form for standard games is generally interpreted as a reduced form representation for an "extensive" game described in more detailed terms. Similarly, as will be shown below, a p-game in normal form can be interpreted as a representation for a more primitive p-game. However, unlike extensive games in which functional representations for probability and preference are already assumed to exist, this "primitive p-game" is defined entirely in terms of sets and relations (3.1). Thus the normal form representation for the primitive p-game must be axiomatized; it is not a reduced form. As will be seen below (3.2), necessary and sufficient conditions exist for primitive p-games with finite state flow sets to have a normal form representation. In 3.3 it will be shown that the standard n-person game in normal form can be interpreted as an axiomatized primitive p-game in normal form.

3.1 THE PRIMITIVE P-GAME: Let G be a set of *(joint) candidate goals,* and for each $g \in G$ let $\Lambda(g) = \{\lambda(g), \ldots\}$ be a set of *(joint) controls*. Then the components for a *primitive n-person p-game* are given by

$$(\{(\Theta, \overset{i}{\succsim})\,|\,i \varepsilon n^*\};\{(\Omega_i(\theta), \succsim^i_\theta, \geq^i_\theta)\,|\,\theta \varepsilon \Theta, i \varepsilon n^*\})$$

where

$n^* \equiv \{1, \ldots, n\}$ is the *player set;*

$\Theta \equiv \bigcup_{g \varepsilon G}(\{g\} \times \Lambda(g))$ is the *joint policy choice set* consisting of candidate-goal control pairs *(joint policies);*

$\overset{i}{\succsim}$ *(policy preference order* for player i) is a weak order[3] on Θ, for all $i \varepsilon n^*$;

and for each joint policy $\theta \varepsilon \Theta$ and $i \varepsilon n^*$,

$\Omega_i(\theta) = \{\omega_i(\theta),\ldots\}$ is a nonempty set of *state flows* associated with θ by player i;

$\succcurlyeq^i_\theta$ (θ-*conditioned preference order* for player i) is a weak order[3] on $\Omega_i(\theta)$;

$\geq^i_\theta$ (θ-*conditioned probability order* for player i) is a weak order on $2^{\Omega_i(\theta)}$, the algebra[4] comprising all subsets *(event flows)* $E_i(\theta) \subseteq \Omega_i(\theta)$.

The candidate goals and controls can be interpreted as in section 2. The weak orders $\succcurlyeq^i$, $i \in n^*$, can be interpreted as preference orders as follows. For all joint policies $\theta', \theta'' \in \Theta$:

$\theta' \succcurlyeq^i \theta''$ <=> the joint selection of θ' is at least as desirable to player i as the joint selection of θ''.

For each $\theta \in \Theta$ and $i \in n^*$, the set $\Omega_i(\theta)$ of state flows $\omega_i(\theta)$ can be interpreted as player i's answer to the following question: "If joint policy θ were to be chosen by the n players, what distinct situations (i.e., state flows $\omega_i(\theta)$) might obtain?" The state flows may include references to past, present, and future happenings. In order for subsequent probability assessments to be realistically feasible, each state flow should include player i's background information concerning the decision problem at hand.

The θ-conditioned preference order $\succcurlyeq^i_\theta$ can be interpreted as follows. For all $\omega', \omega'' \in \Omega_i(\theta)$,

$\omega' \succcurlyeq^i_\theta \omega''$ <=> the realization of ω' is at least as desirable to player i as the realization of ω'', given the event "the n players choose θ."

Similarly, the θ-conditioned probability order $\geq^i_\theta$ can be interpreted as follows. For all event flows E', E'' ε $2^{\Omega_i(\theta)}$,

$$E' \geq^i_\theta E'' \iff$$ in the judgment of player i, the realization of E' is at least as likely as the realization of E'', given the event "the n players choose θ."

A state flow may be relevant for player i's decision problem under distinct potential policy choices; i.e., $\Omega_i(\theta) \cap \Omega_i(\theta') \neq \emptyset$ for some $\theta, \theta' \varepsilon \Theta$. Similarly for event flows. Given state flows $\omega, \omega' \varepsilon \Omega_i(\theta) \cap \Omega_i(\theta')$, it may hold that $\omega \succcurlyeq^i_\theta \omega'$ whereas $\omega' \succcurlyeq^i_{\theta'} \omega$. Verbally, the relative desirability of the state flows ω and ω' for player i may depend on which conditioning event he is considering, "the n players choose θ" or "the n players choose θ'." Similarly for the relative likelihood of event flows.

An example illustrating these interpretations is given in section 5. Other examples are given in Tesfatsion [9].

3.2 EXPECTED UTILITY REPRESENTATION: In Tesfatsion [10,6.1] necessary and sufficient conditions A* are given which ensure that a one-person primitive p-game (equivalently, a primitive p-model)

$$((\Theta, \succcurlyeq^1); \{(\Omega_1(\theta), \succcurlyeq^1_\theta, \geq^1_\theta) \mid \theta \varepsilon \Theta\})$$

with finite state flow sets $\Omega_1(\theta)$ has an expected utility representation in the following sense. For each policy $\theta \varepsilon \Theta$ there exists a finitely additive probability measure $\sigma_1(\cdot|\theta):2^{\Omega_1(\theta)} \to [0,1]$ satisfying

(2) $$\sigma_1(E|\theta) \geqq \sigma_1(E'|\theta) \iff E \geq^1_\theta E',$$

for all E, E' ε $2^{\Omega_1(\theta)}$, and a utility function $u_1(\cdot|\theta):\Omega_1(\theta) \to R$

satisfying

$$(3) \qquad u_1(\omega|\theta) \geqq u_1(\omega'|\theta) \iff \omega \succsim^1_\theta \omega',$$

for all $\omega, \omega' \in \Omega_1(\theta)$, such that

$$(4) \qquad \int_{\Omega_1(\theta)} u_1(\omega|\theta)\sigma_1(d\omega|\theta) \geqq \int_{\Omega_1(\theta')} u_1(\omega|\theta')\sigma_1(d\omega|\theta') \iff \theta \succsim^1 \theta',$$

for all $\theta, \theta' \in \Theta$.

The conditions A* are restrictions on player 1's preference and probability orders

$$\{ \succsim^1, \succsim^1_\theta, \geq^1_\theta \mid \theta \in \Theta \}.$$

For primitive n-person p-games with finite state flow sets, necessary and sufficient conditions for the existence of an individual objective (expected utility) representation as in (2), (3), (4) for each player's policy preference order can be obtained by requiring conditions A* to hold for each player's preference and probability orders

$$\{ \succsim^i, \succsim^i_\theta, \geq^i_\theta \mid \theta \in \Theta \}, \ i \in n^*.$$

The expected utility representation (2), (3) and (4) for the policy preference orders $\succsim^i$, $i \in n^*$, can be interpreted as follows. To each state flow $\omega \in \Omega_i(\theta)$, $\theta \in \Theta$, player i assigns a utility number $u_i(\omega|\theta)$ representing the desirability of $\{\omega\}$ obtaining, conditioned on the event "the n players choose θ," and a probability number $\sigma_i(\omega|\theta)$ representing the likelihood of $\{\omega\}$ obtaining, conditioned on the event "the n players choose θ." He then calculates his expected utility

$$U_i(\theta) \equiv \int_{\Omega_i(\theta)} u_i(\omega|\theta)\sigma_i(d\omega|\theta)$$

corresponding to each choice of joint policy $\theta \in \Theta$. The function $U_i: \Theta \to R$ is then the *objective function for player i*. The components for an axiomatized primitive n-person p-game can thus be presented in normal form

$$(n^*;\ \Theta;\ \Pi_{i\in n^*}\ U_i:\Theta \to R^n),$$

as in section 2.

3.3 STANDARD GAMES AS P-GAMES: For standard n-person games Γ in normal form (see Owen [7]), the choice of strategies $\delta_1,\ldots,\delta_n$ by players 1,...,n and a strategy δ_o "by chance" determine that a unique outcome $t(\delta_o,\ldots,\delta_n)$ will obtain. Player i then receives a certain real-valued (utility) payoff $f_i(t(\delta_o,\ldots,\delta_n))$. Since the players do not know in advance which chance strategy will obtain, it is assumed that each player i calculates his expected payoff $\psi_i(\delta_1,\ldots,\delta_n)$ corresponding to each joint strategy $(\delta_1,\ldots,\delta_n)$. For example, denoting the strategy set for player i by Δ_i and the set of possible chance strategies by Δ_o,

$$\psi_i(\delta_1,\ldots,\delta_n) = \int_{\Delta_o} f_i(t(\delta_o,\delta_1,\ldots,\delta_n)\ P_i(d\delta_o),$$

where P_i represents player i's assessments concerning the likelihood of the chance strategies in Δ_o. The goal of player i is assumed to be g_i: (player i maximizes his expected payoff).

The game Γ can be viewed as a free p-game as follows. Define

$n^* \equiv \{1,\ldots n\}$, the player set;

$G_i' \equiv \{g_i\}$, a one-element candidate goal set for player i,

$i \varepsilon n^*$;

$\Lambda'_i(g_i) \equiv \Delta_i$, the control set for player i corresponding to

g_i, $i \varepsilon n^*$;

$\Theta'_i \equiv \{g_i\} \times \Lambda'_i(g_i) = \{(g_i,\delta_i),(g_i\ ,\delta_i'),\ldots\}$, the individual

policy choice set for player i, $i \varepsilon n^*$;

and for each joint policy $\theta = ((g_1,\delta_1),\ldots,(g_n,\delta_n)) \varepsilon \Pi_{i\varepsilon n^*} \Theta_i'$ and $i \varepsilon n^*$,

$\Omega_i'(\theta) \equiv \Delta_o$, the state flow set associated

with θ by player i;

$u_i'(\cdot|\theta):\Omega_i'(\theta) \to R$, the utility function associated with θ

by player i, given by $u_i'(\delta_o|\theta) = f_i(t(\delta_o,\ldots,\delta_n))$,

$\delta_o \varepsilon \Omega_i'(\theta)$;

$\sigma_i'(\cdot|\theta)$: $2^{\Omega_i'(\theta)} \to [0,1]$, the probability measure associated

with θ by player i, where $\sigma_i'(\delta_o|\theta) = P_i(\delta_o), \delta_o \varepsilon \Omega_i'(\theta)$.

Then for each joint policy $\theta = ((g_1,\delta_1),\ldots,(g_n,\delta_n)) \varepsilon \Pi_{i=1}^{n} \Theta_i'$, and each $i \varepsilon n^*$,

$$\int_{\Omega_i'(\theta)} u_i'(\omega|\theta)\ \sigma_i'(d\omega|\theta) \equiv \int_{\Delta_o} f_i(t(\delta_o,\ldots,\delta_n))P_i(d\delta_o)$$

$$\equiv \psi_i(\delta_o,\ldots,\delta_n).$$

Hence Γ can be identified with the free p-game

$$(n^*;\ \Pi_{n^*} \Theta_i';\ \Pi_{n^*} U_i':\ \Pi_{n^*} \Theta_i' \to R^n),$$

where, for each $i \varepsilon n^*$, the objective function U_i': $\Pi_{n^*} \Theta_i' \to R$ for player i is given by

$$U_i'(\theta)=\int_{\Omega_i'(\theta)} u_i'(\omega|\theta)\sigma_i'(d\omega|\theta),\ \theta \in \Pi_{n*} \Theta_i'.$$

4. SOLUTIONS FOR POLICY GAMES

Researchers investigating the existence of Nash equilibrium points for standard games in normal form have generally assumed that the players can "mix" their strategy choices by resorting to various random devices; e.g., player i may decide to implement pure strategy b if a flipped coin lands heads and pure strategy d if it lands tails. Similarly, researchers investigating the existence of bargaining solutions in 2-person standard games have generally assumed that the two players can "correlate" their choice of strategies; e.g., players 1 and 2 may agree to implement joint strategy (b,c) if a flipped coin lands heads and joint strategy (d,e) if it lands tails. Individual strategy sets for standard games with mixed strategies are convex. Moreover, since (utility) payoff functions in standard games are assumed to be linear with respect to lotteries, the payoff regions for standard games with correlated strategies are convex. For such games a distinction is not made between pure and mixed or correlated strategy solutions (see Owen [7]).

On the other hand, in real-world group decision contexts such as boardroom meetings, union versus management, and presidential elections, the flipping of coins is seldom observed. Conditions ensuring the existence of pure solutions are thus of particular interest. As is well known, all finite standard n-person games in extensive form with "complete information" have pure strategy Nash equilibrium points (Owen, [7, I.4.5]). Moreover, the existence theorems established for Nash equilibrium points in mixed strategy games in normal form can be extended to pure strategy games in normal form having convex strategy sets. In addition, Debreu [2] has established the existence of pure strategy Nash equilibrium points for games in normal form whose individual strategy sets are contractible (hence connected) polyhedra. On the other hand, no Nash equilibrium existence results appear to have been established for games in normal form having disconnected strategy sets. Similarly, Nash's bargaining solution has been investigated only for games with convex (hence connected) payoff regions.

The joint policy choice set Θ for p-games is a disjoint union $\bigcup_{g\varepsilon G} (\{g\} \times \Lambda(g))$ of policy subsets $\{g\} \times \Lambda(g)$ corresponding to the distinct candidate goals $g \varepsilon G$. Unless G is a trivial one-element set, it cannot be assumed that Θ is connected; similarly for the payoff region $\Pi_{n*}U_i(\Theta)$. In theorems 4.2 and 4.5 below, conditions are given which are sufficient to guarantee the existence of Nash equilibrium and bargaining solutions for p-games in normal form without requiring either Θ or $\Pi_{n*}U_i(\Theta)$ to be connected. Since standard games in normal form have been shown (3.3 above) to be a special type of p-game, these theorems represent an extension of existing game theory.

4.1 DEFINITION: Let a free n-person p-game $(n^*; \Pi_{n*}\Theta_i; \Pi_{n*}U_i: \Pi_{n*}\Theta_i \to R^n)$ be given. A joint policy $\theta' \equiv (\theta_1',\ldots,\theta_n') \varepsilon \Pi_{n*}\Theta_i$ will be called a *pure policy Nash equilibrium point* if for each $i \varepsilon n^*$,

$$(5) \qquad U_i(\theta') \geqq U_i(\theta_1',\ldots,\theta_{i-1}', \theta_i, \theta_{i+1}',\ldots,\theta_n')$$

for all $\theta_i \varepsilon \Theta_i$. (See Luce and Raiffa [3] for a critique of this solution concept.)

4.2 THEOREM (Tesfatsion [11,2.8]): *Let a free n-person p-game*

$$\Gamma \equiv (n^*;\ \Pi_{n^*}\Theta_i;\ \Pi_{n^*} U_i : \Pi_{n^*}\Theta_i \to R^n)$$

be given such that:

1) $\Pi_{n^*}\Theta_i$ *is a compact metrizable absolute neighborhood retract;*

2) $\Pi_{n^*} U_i$ *is continuous;*

3) $T(\theta)$ *is* C_F*-acyclic (i.e., acyclic with respect to Čech homology over a field F) for each* $\theta \in \Pi_{n^*}\Theta_i$*, where* $T \equiv \Pi_{n^*}T_i : \Pi_{n^*}\Theta_i \to \Pi_{n^*}\Theta_i$ *is a multivalued map defined by*

$$T_i(\theta') \equiv \{\theta_i^\circ \in \Theta_i \mid U_i(\theta_1',\ldots,\theta_i^\circ,\ldots,\theta_n') = \max U_i(\theta_1',\ldots,\theta_i,\ldots,\theta_n'),\ \theta_i \in \Theta_i\}$$

for $\theta' \equiv (\theta_1',\ldots,\theta_n') \in \Pi_{n^*}\Theta_i$ *and* $i \in n^*$;

4) *The Lefschetz number of T (with respect to Čech homology over F) is not zero.*

Then Γ *has at least one pure policy Nash equilibrium point.*

REMARK: For a detailed discussion of 4.2, please refer to Tesfatsion [11]. As discussed there, many of the spaces commonly used in economic and game theory are compact metrizable absolute neighborhood retracts: for example, compact convex subsets of Banach spaces; finite dimensional locally contractible compact metrizable spaces (e.g., finite discrete spaces); and locally euclidean compact metrizable spaces (e.g., compact n-manifolds). Contractible (e.g. convex) subsets of compact Hausdorff spaces are C_F-acyclic. If $\Pi_{n*}\Theta_i$ is a C_F-acyclic compact Hausdorff space, then the Lefschetz number of T is equal to 1. In general, however, the hypotheses of 4.2 do not require any form of global connectedness for $\Pi_{n*}\Theta_i$.

4.3 DEFINITIONS: A 2-person p-game $\Gamma \equiv (\{1,2\};\ \Theta\ ;\ U_1 \times U_2 : \Theta \to R^2)$ will be called a *threat p-game* if the two players bargain with each other in three stages as follows: Stage 1) A status quo *threat* $\theta^* \varepsilon \Theta$ is announced; Stage 2) Players 1 and 2 attempt to agree on a joint policy $\theta \varepsilon \Theta$; Stage 3) If an agreement on a joint policy θ is reached, it is implemented and player i receives $U_i(\theta)$. If an agreement on a joint policy is not reached, then the threat θ^* is enforced and player i receives $U_i(\theta^*)$.

By refusing to come to an agreement, player i can ensure himself of the payoff $U_i(\theta^*)$. Hence the effective range of joint payoffs for players 1 and 2 arising from pure policy choices $\theta \varepsilon \Theta$ is given by the *barter set*

$$B(u^*,v^*) \equiv B \cap \{(u,v) \ \varepsilon\ R^2 \mid u \geqq u^*,\ v \geqq v^*\},$$

where $u^* \equiv U_1(\theta^*)$, $v^* \equiv U_2(\theta^*)$, and $B \equiv \{U_1 \times U_2(\theta) \mid \theta \ \varepsilon\ \Theta\}$.

A function $J:D \to R^2$, D an arbitrary set, will be said to be *corner concave* if for every pair d, d' ε D and every r ε [0,1] there exists d* ε D such that

(6) $\qquad rJ(d) + [1-r]\ J(d') \leqq J(d^*)$.

If D is interpreted as a collection of pure joint policies and J is interpreted as a joint objective function, then (6) has an obvious interpretation: Each available correlated joint policy (lottery among pure joint policies) involving two pure joint policies is "dominated" by at least one available pure joint policy in the sense that, for each player i, the expected utility of the correlated joint policy is no greater than the utility of the pure joint policy.

A function $J:D \to R^2$ will be said to be *corner concave with respect to* $(u^*,v^*) \in R^2$ if the restriction of J to $J^{-1}(\{(u,v) \mid u \geqq u^*, v \geqq v^*\})$ is corner concave.

In the nineteen fifties Nash ([4],[5]) proposed six axioms as an empirically meaningful set of conditions which should be satisfied by any "barter rule" devised to settle 2-person games by assigning to each barter set $B \subseteq R^2$ a "solution" $(u,v) \in B$. He then established the existence of a unique barter rule satisfying these six axioms with respect to the class of all standard 2-person threat games with compact convex barter sets. In 4.5 below this result will be extended to the class of all 2-person threat p-games with compact policy sets and continuous, corner concave joint objective functions. Barter sets corresponding to such games need not be connected.

The six Nash axioms will first be presented in general form.

4.4 THE NASH AXIOMS: Let $\mathcal{D}^*$ denote any collection of subsets of the form

$$D(u^*,v^*) \equiv D \cap \{(u,v) \in R^2 \mid (u,v) \geqq (u^*,v^*)\}$$

with $(u^*,v^*) \in D \subseteq R^2$. A function $\phi: \mathcal{D}^* \to R^2$ will be said to satisfy the Nash axioms with respect to $\mathcal{D}^*$ if for every $D(u^*,v^*) \in \mathcal{D}^*$:

Axiom 1 (Feasibility): $\phi(D(u^*,v^*)) \in D$;

Axiom 2 (Individual Rationality): $\phi(D(u^*,v^*)) \geqq (u^*,v^*)$;

Axiom 3 (Pareto Optimality): If $(u,v) \in D$ and

$(u,v) \geqq \phi(D(u^*,v^*))$, then $(u,v) = \phi(D(u^*,v^*))$;

Axiom 4 (Independence of Irrelevant Alternatives): If

$A(u^*,v^*) \in \mathcal{D}^*$ with $(u',v') \in A(u^*,v^*) \subseteq$

$D(u^*,v^*)$ and $(u',v') = \phi(D(u^*,v^*))$, then

$(u',v') = \phi(A(u^*,v^*))$;

Axiom 5 (Independence of Linear Transformations): Suppose

for some $r_1,r_2,s_1,s_2 \in R$,

$$E \equiv \{(r_1u+s_1,r_2v+s_2) \mid (u,v) \in D(u^*,v^*)\} \in \mathcal{D}^*.$$

Then if $\phi(D(u^*,v^*)) = (u',v')$, it must hold

that $\phi(E) = (r_1u' + s_1,r_2v' + s_2)$;

Axiom 6 (Symmetry): Suppose $(u,v) \in D(u^*,v^*)$ iff $(v,u) \in$

$D(u^*,v^*)$, and suppose $u^* = v^*$ and $\phi(D(u^*,v^*)) = (u',v')$.

Then $u' = v'$.

REMARK: See Smorodinsky and Kalai [8] and Luce and Raiffa [3] for a critical appraisal of the Nash axioms; also Nydegger and Owen [6] for an experimental test of these axioms.

Let $\mathcal{C}^*$ denote the collection of all barter sets corresponding to 2-person threat p-games with compact joint policy choice set Θ and continuous joint objective function $U_1 \times U_2 : \Theta \to R^2$, corner concave with respect to the threat payoff $U_1 \times U_2(\theta^*)$.

4.5 THEOREM (Tesfatsion [12, 4.4 and 4.10]): *There exists a unique barter rule* ϕ: $C^* \to R^2$ *which satisfies the Nash axioms with respect to* C^*.

REMARK: For a detailed discussion of 4.5, please refer to Tesfatsion [12]. As proved there, corner concavity of a continuous function $J:D \to R^2$, D compact, is equivalent to "corner concavity" of the set J(D), defined as follows. For any set $A \subseteq R^2$, let A^- denote its closed convex hull and let A^p denote the set of all pareto optimal (efficient) points in A. Then A is said to be "corner concave" if A^p is a nonempty compact set which coincides with $(A^-)^p$. The set and function definitions of corner concavity both have obvious extensions to R^n, $n > 2$, which leads to the following conjecture: for $n \geqq 2$, corner concavity for a continuous function $J:D \to R^n$, D compact, is equivalent to corner concavity for $J(D) \subseteq R^n$.

Suppose Γ is a free 2-person p-game with compact joint policy choice set $\Theta_1 \times \Theta_2$ and continuous corner concave joint objective function $U \equiv U_1 \times U_2$: $\Theta_1 \times \Theta_2 \to R^2$. The question arises whether there exists a pure policy Nash equilibrium threat for Γ; i.e., letting $W_1 \times W_2(\cdot) \equiv \phi(B(U(\cdot)))$: $\Theta_1 \times \Theta_2 \to R^2$, where ϕ is the barter rule in 4.5, a joint policy $(\theta_1^*, \theta_2^*) \in \Theta_1 \times \Theta_2$ which satisfies (5) with respect to $W_1 \times W_2$. If the derived game $\Gamma \equiv (\{1,2\};\ \Theta_1 \times \Theta_2;\ W_1 \times W_2: \Theta_1 \times \Theta_2 \to R^2)$ satisfies conditions 1), 2), 3), and 4) in 4.2, then the answer is affirmative.

5. EXAMPLE OF A P-GAME

An illegal parking problem is formalized in terms of a free 2-person p-game with expected utility representation as axiomatized in section 2. The example illustrates three points. First, certain group decisions can be given an expected utility rationalization even though the decision makers specify their available acts in the form of partial rather than complete contingency plans. Second, the choice among end-mean, goal-control pairs arises naturally in many decision problems. Third, choice of goal as well as choice of control can affect a decision maker's probability and preference judgments concerning future events.

The campus Chief of Police C must devise a new policy to meet the problem of bicycles parked illegally next to the library rather than in a bikerack some distance away. At present C's officers issue only warning tags. All collected fines are traditionally spent by C at his discretion. All confiscated bikes are traditionally given to the Sports Club.

Assume C views his problem as the following free noncooperative p-game between himself and a typical student S, in which he has the first move (i.e., C will implement his chosen policy before S).

C's Candidate Goal Set

$$G_C \equiv \{ g_1\text{: maximize collected fines, to be used for merit awards} \; ; \; g_2\text{: minimize illegal parking} \} \; ;$$

Remark: More realistically, the candidate goals $g \in G_C$ could be vectors (p,r), where p varies over a set of target ceilings for the average illegal parking rate (e.g., at most twice a week) and r varies over a set of target floors for average collected revenues (e.g., at least \$1.00 per week).

S's Candidate Goal Set

$G_S \equiv \{ g_o$: minimize costs in terms of time and money $\}$;

S's Control Sets (one control set associated with each candidate goal in G_C)

$\Lambda_C(g_1) \equiv \{$ \$1: C announces \$1 fine rate , \$5: C announces \$5 fine rate $\}$

$\times \{$ t: C announces goal to his officers , t': C does not announce goal to his officers $\}$;

$\Lambda_C(g_2) \equiv \{$\$5, h: C announces "illegally parked bikes will be confiscated" $\} \times \{t,t'\}$;

S's Control Sets (one control set associated with each candidate goal in G_S)

$$\Lambda_S(g_o) \equiv \{ \text{S always parks illegally next to library}, \ \text{S daily flips fair coin to decide where to park}, \ \text{S always parks in bikerack} \}$$

$$\equiv \{ (\text{lib}), (\text{flip}), (\text{rack}) \} ;$$

Joint Policy Choice Set

$$\equiv \bigcup_{g \varepsilon G_C} (\{g\} \times \Lambda_C(g)) \quad \times \quad \bigcup_{g \varepsilon G_S} (\{g\} \times \Lambda_S(g))$$

$$\equiv \Theta_C \times \Theta_S .$$

Let ω° denote C's background information concerning the parking problem. In particular, assume ω° contains the information that the current daily diligence index for detecting illegally parked bikes is ½; i.e., C believes that at present, on any given day, the probability is ½ that his officers making the rounds will detect a bike parked illegally next to the library. It seems reasonable to assume that C is worried about the possible effects that the various joint policy choices by C and S might have on this diligence index. For example, C might be quite sure that the index would rise if he were to announce to his officerts the goal g_1: (maximize collected fines, to be used for merit awards), whereas at best the index would stay the same if he were to exhort them with g_2:(minimize illegal parking).

To examine the various possibilities, C asks himself the following question for each joint policy $\theta \in \Theta_C \times \Theta_S$: "Given that θ is implemented by S and myself, what distinct situations (state flows) might occur, relevant for the problem at hand?" For example, the set of state flows C associates with the joint policy $\theta' \equiv (g_1,(\$1,t)) \times (g_o,\text{lib}) \in \Theta_C \times \Theta_S$ might be

$$\Omega_C(\theta') \equiv \{\omega^\circ\} \times \left\{ \begin{array}{l}\text{diligence rate}\\ \text{increases; hence}\\ \text{expected daily fine}\\ \text{revenues} > 50¢\end{array} , \begin{array}{l}\text{diligence rate}\\ \text{fails to increase}\\ \text{hence expected}\\ \text{daily fine reve-}\\ \text{nues} \leq 50¢\end{array} \right\};$$

whereas the set of state flows he associates with the joint policy $\theta'' \equiv (g_1,(\$5,t)) \times (g_o,\text{lib}) \in \Theta_C \times \Theta_S$ might be

$$\Omega_C(\theta'') \equiv \{\omega^\circ\} \times \left\{ \begin{array}{l}\text{diligence rate}\\ \text{increases; hence}\\ \text{expected daily fine}\\ \text{revenues} > \$2.50\end{array} , \begin{array}{l}\text{diligence rate}\\ \text{fails to increase}\\ \text{hence expected}\\ \text{daily fine reve-}\\ \text{nues} \leq \$2.50\end{array} \right\};$$

REMARK: Presumably C is unable or unwilling to estimate the conditional likelihoods that S will choose any particular policy in Θ_S, following a policy choice by C. If C were willing to make such estimates, then S's possible policy choices could be embodied in state flow sets conditioned just on C's policies in Θ_C rather than on joint policies in $\Theta_C \times \Theta_S$. The p-game would then essentially reduce to a one-person policy model problem, as discussed in Tesfatsion [9].

Assume to each state flow $\omega \in \Omega_C(\theta)$, $\theta \in \Theta_C \times \Theta_S$, C assigns a utility number $u(\omega|\theta)$ representing the desirability of ω obtaining, conditioned on the event "C and S choose θ," and a probability number $\sigma(\omega|\theta)$ representing the likelihood of ω obtaining, conditioned on the event "C and S choose θ." He then calculates his expected utility

$$\int_{\Omega_C(\theta)} u(\omega|\theta)\sigma(d\omega|\theta)$$

corresponding to each choice of joint policy $\theta \in \Theta_C \times \Theta_S$.

For simplicity, let it also be assumed that C finds only three of his policy choices in Θ_C to be undominated, so that he need only estimate S's expected utility levels for nine possible joint policy choices as illustrated in the following "expected utility matrix" [(p,q)]. (The components p give C's expected utility and the components q give S's expected utility.)

C \ S	θ_S':(g_o,lib)	θ_S'':(g_o,flip)	θ_S''':(g_o,rack)
θ_C':(g_1,(\$1,t))	(4,8)	(3,7)	(2,2)
θ_C'':(g_1,(\$5,4))	(8,4)	(6,6)	(1,3)
θ_C''':(g_2,(h,t'))	(1,0)	(2,1)	(5,4)

Expected Utility Matrix

According to the expected utility matrix, if C's chosen goal is g_1: (maximize collected fines, to be used for merit awards), then, from C's point of view, the more illegal parking by S the better; and if C's chosen goal is g_2: (minimize illegal parking), then the less illegal parking by S the better. Further, if C chooses g_1 as a goal and S chooses (g_o,rack), i.e., straight legal parking, then C would prefer S made such a choice in the face of a \$1 fine rate rather than a \$5 fine rate. (For example, under a \$5 fine rate, the blame for the failure to obtain revenues might be attributed to too high a fine rate.)

From S's point of view, the higher the fine rate, the lower the expected utility of illegal parking. Confiscation also makes illegal parking less attractive than legal. On the other hand, given a low fine rate of \$1, the time saved parking illegally next to the library outweighs the risk of being detected. For the higher fine rate, prudence suggests only occasional illegal parking.

The set of pareto optimal joint policies is

$$\{(\theta'_C,\theta'_S),(\theta''_C,\theta''_S),(\theta''_C,\theta'_S)\}$$

whereas the set of Nash equilibrium joint policies is

$$\{(\theta''_C,\theta''_S),(\theta'''_C,\theta'''_S)\}.$$

The highest utility level C can guarantee for himself is 6, through choice of θ''_C, assuming that S is "rational" and consequently chooses θ''_S. The resulting joint policy (θ''_C,θ''_S) is a unique pareto optimal - Nash equilibrium solution, yielding a joint payoff of (6,6).

REMARK: With $\Theta_C \times \Theta_S$ assumed to have the discrete topology, it can be shown that the above p-game satisfies the hypotheses of 4.2. However, this application of 4.2 is rather trivial. The power of the Lefschetz condition 4) in 4.2 is more clearly revealed when the policy choice set Θ has a richer topological structure.

FOOTNOTES

[1]An example of a tautological specification of a goal by a decision maker i would be "i maximizes his utility." Such specified goals convey no information to other decision makers in the problem environment, and hence can play no strategic role.

[2]Goals are generally tinged with probability considerations; they are the end result of an act of choice in the face of uncertainty. Thus goals cannot always be identified with "desirable outcomes". (Most people would like to be billionaires, but how many adopt this as a goal?)

[3]A binary relation $\geqq$ on a set D is a *weak order* if for all $a,b,c \in D$

(1) $a \geqq b$ or $b \geqq a$ (i.e., $\geqq$ is connected);

(2) $a \geqq b$ and $b \geqq c$ implies $a \geqq c$ (i.e., $\geqq$ is transitive).

[4]A collection F of subsets of a nonempty set X is said to be an *algebra* in X if F has the following three properties:

(1) $X \in F$;

(2) If $A \in F$, then $A^c \in F$, where A^c is the complement of A relative to X;

(3) If $A,B \in F$, then $A \cup B \in F$.

REFERENCES

[1] Aumann, R. J. and J. H. Dreze: "Cooperative Games With Coalition Structures," *International Journal of Game Theory*, 3 (1974), 217-237.

[2] Debreu, G.: "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences*, 38 (1952), 886-893.

[3] Luce, R. D., and H. Raiffa: *Games and Decisions*, John Wiley & Sons, N.Y., 1957.

[4] Nash, J.: "The Bargaining Problem," *Econometrica*, 18 (1950), 155-162.

[5] Nash, J.: "Two-Person Cooperative Games," *Econometrica*, 21 (1953), 128-140.

[6] Nydegger, R. V. and G. Owen: "Two Person Bargaining: An Experimental Test of the Nash Axioms," *International Journal of Game Theory*, 3 (1974), 239-249.

[7] Owen, G.: *Game Theory*, W. B. Saunders Company, Philadelphia, 1968.

[8] Smorodinsky, M. and E. Kalai: "Other Solutions to Nash's Bargaining Problem," *Econometrica*, 43 (1975), 513-518.

[9] Tesfatsion, L.: "An Expected Utility Model with Endogenously Determined Goals," Discussion Paper 58-75, Center for Economic Research, University of Minnesota, Minneapolis, September 1975.

[10] Tesfatsion, L.: "Axiomatization for an Expected Utility Model With Endogenously Determined Goals," Discussion Paper 59-75, Center for Economic Research, University of Minnesota, Minneapolis, September 1975.

[11] Tesfatsion, L.: "Pure Strategy Nash Equilibrium Points and the Lefschetz Fixed Point Theorem," Discussion Paper 60-75, Center for Economic Research, University of Minnesota, Minneapolis, September 1975.

[12] Tesfatsion, L.: "Pure Strategy Nash Bargaining Solutions," Discussion Paper 61-75, Center for Economic Research, University of Minnesota, Minneapolis, November 1975.