

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

AN EXPERIMENTAL ANALYSIS OF NASH REFINEMENTS IN SIGNALING GAMES

Jeffrey Banks
University of Rochester

Colin Camerer
University of Pennsylvania

David Porter
California Institute of Technology



SOCIAL SCIENCE WORKING PAPER 740

June 1990

AN EXPERIMENTAL ANALYSIS OF NASH REFINEMENTS
IN SIGNALING GAMES

Jeffrey Banks
University of Rochester

Colin Camerer
University of Pennsylvania

David Porter
California Institute of Technology

ABSTRACT

This paper investigates the refinements of Nash equilibrium in two person signaling game experiments. The experimental games cover the watershed of the nested refinements: Bayes-Nash, Sequential, Intuitive, Divine, Universally Divine, NWBR, and Stabel. In each game an equilibrium selection problem is defined in which adjacent refinements are considered.

The pattern of outcomes suggest that individuals select the more refined equilibria up to the divinity concept. However, an anomaly occurs in the game in which the stable equilibrium is a clear preference among the subjects. Since the concepts are nested this suggests that the outcomes are game specific. Sender behavior does not seem to follow any specific decision rule (e.g., Nash, minmax, PIR, etc.) while receiver actions tend to correspond to the Nash equilibrium outcomes.

Key Words: Game Theory, Nash Refinements, Experimental Economics, Signaling Games

AN EXPERIMENTAL ANALYSIS OF NASH REFINEMENTS IN SIGNALING GAMES*

Jeffrey Banks
University of Rochester

Colin Camerer
University of Pennsylvania

David Porter
California Institute of Technology

1. INTRODUCTION

In a noncooperative game a Nash equilibrium point is a set of strategies, one for each player, which are best responses to each other. The Nash concept has many virtues. It always exists in finite games and is intuitively appealing, precise and simple to teach and apply. Indeed, the Nash concept seems too simple in many games, especially those with complicated dynamics and information structure because it permits equilibrium points which seem implausible or illogical. Many game theorists have tried to sharpen the Nash concept by proposing logical rules stronger than the mutual-best-response requirement which strategies must satisfy. These rules are called "refinements" of Nash equilibrium.

There are remarkably few data guiding the theoretical process of refinement. This paper attempts to fill that gap by testing various refinements experimentally in two-player signaling games with one-sided incomplete information. Equilibrium refinements have been applied to signaling games in many areas including finance (Harris and Raviv, 1985), choice of product quality (Grossman, 1981; Milgrom and Roberts, 1986), education (Spence, 1974), bargaining (Rubinstein, 1985) and predatory pricing (Selten, 1978; Kreps and Wilson, 1982b; Milgrom and Roberts, 1982a,b). Fudenberg and Tirole (1987) review some of the many applications.

Our experiments test six refinements: *sequentiality* (Kreps and Wilson, 1982a); *the intuitive criterion* (Cho and Kreps, 1987); *divinity and universal divinity* (Banks and Sobel, 1987); *the never-*

* We are indebted to Mark Olson for help in statistical analysis and programming. Audiences at Washington University, University of Arizona, University of Houston, and University of Pennsylvania provided helpful comments. This research was partially funded by the Wharton Risk and Decision Processes Center. Support for the second author was provided by the National Science Foundation (SES 87-08566).

a-weak-best-response (NWBR) criterion (Kohlberg and Mertens, 1986); and *stability* (Kohlberg and Mertens, 1986).

We constructed two games to test for Nash behavior and six games to test the refinements against each other. In each of the six games there are two equilibria, one more refined and one less refined; e.g., a Nash equilibrium which is not sequential and a Nash equilibrium which is sequential.

An alternative, more conservative approach would be to pick one refinement and study it in a variety of games with different parameters. But if subjects then consistently choose the more refined of two equilibria we will not know whether they would choose even more refined equilibria if they were available. Our experiments avoid this shortcoming.

Furthermore our experiment *does* test each refinement in a variety of games with different parameters. Because the sets of equilibria are nested, (all sequential equilibria are Nash, all intuitive equilibria are sequential, etc.), our experiments thus test each refinement in several games.

The paper is organized as follows. In the next section we define several refinements and illustrate them with the games used in the experiments, in section 3 we describe the experimental design, in section 4 we report results, and section 5 is a conclusion with ideas for further research.

2. A PRIMER ON REFINEMENTS

In the generic form of a signaling game there are two players, a sender S and a receiver R . The sender S has private information summarized by his type $t \in T$. Knowing his type S selects a message $m \in M$ which R observes before choosing an action $a \in A$. R does not know t before making his choice, but his beliefs about t are characterized by a prior probability distribution $P(t)$ over the set T which is *common knowledge*. Preferences for S and R are represented by von Neumann-Morgenstern utility functions $u(t, m, a)$ and $v(t, m, a)$, respectively. The sets T, M, A are assumed to be finite.

For any finite set K let Δ_K denote the set of all probability distributions over K . A signaling strategy for S is a function from types into messages

$$q : T \rightarrow \Delta_M,$$

where $q(m | t)$ denotes the probability that S sends the message m given type t . A response strategy for R is a function from messages into actions

$$r : M \rightarrow \Delta_A$$

where $r(a | m)$ denotes the probability that R takes action a in response to the message m . Since $u(\cdot)$ and $v(\cdot)$ are assumed to be von Neumann-Morgenstern utility functions, we can extend the utility functions u and v to the strategy spaces associated with Δ_A by taking expected values.

For any $\lambda \in \Delta_T$ and $m \in M$, define the best response correspondence for R by

$$BR(\lambda, m) = \operatorname{argmax}_{r \in \Delta_A} \sum_{t \in T} v(t, m, r(m)) \cdot \lambda(t)$$

and for any $\Lambda \subset \Delta_T$ let

$$BR(\Lambda, m) = \bigcup_{\lambda \in \Lambda} BR(\lambda, m).$$

The concept of Nash equilibrium provides the starting place for analysis in games of complete information. The generalization of the Nash equilibrium concept of games of incomplete information in which signaling games are a subset is Bayesian-Nash equilibrium (Harsanyi, 1967-68).

Definition.¹ A Bayesian Nash equilibrium consists of strategies q, r , and beliefs $\mu(\cdot | m) \in \Delta_T$ such that

i) $\forall t \in T, q(m' | t) > 0$ only if

$$u(t, m', r(m')) = \max_{m \in M} u(t, m, r(m)) :$$

ii) $\forall m \in M$ s. t. $q(m | t) > 0$ for some $t \in T, r(a' | m) > 0$ only if

$$\sum v(t, m, a') \cdot \mu(t | m) = \max_{a \in A} \sum_{t \in T} v(t, m, a) \cdot \mu(t | m) :$$

iii) $\forall m \in M$ s. t. $q(m | t) > 0$ for some $t \in T,$

$$\mu(t' | m) = \frac{q(m | t') \cdot P(t')}{\sum_{t \in T} q(m | t) \cdot P(t)}$$

Part i) of the definition says that the signaling strategy of S is optimal for each type $t \in T$ given the response strategy of R . Part ii) says that R is selecting actions optimally for messages which are sent with positive probability ("along the equilibrium path"). Part iii) says that R uses Bayes' rule to update the prior belief $P \in \Delta_T$ after observing the message m knowing the sender's signaling strategy q .

Consider the signaling game depicted in Table 1. The payoffs in the matrices are $u(t, m, a)$ and $v(t, m, a)$. Each part of the table represents payoffs for a different message (m_1, m_2 , or m_3) for the possible types (t_1 and t_2) and actions (a_1, a_2 , and a_3). We assume throughout that $P(t_1) = P(t_2) = .5$. In game 1 there is a unique Bayesian-Nash equilibrium path² generated by the following strategies: $q(m_1 | t_i) = 1, i = 1, 2$, and $r(a_1 | m_1) = r(a_3 | m_2) = r(a_1 | m_3) = 1$. This is a pooling equilibrium because both types send the same message (m_1). No information about S 's type is revealed by the message, so $\mu(t_i | m_1) = P(t_1) = .5, i = 1, 2$.

1. Although Bayesian-Nash equilibrium is usually not defined explicitly with respect to R 's beliefs, defining it this way makes it easy to compare with the more refined equilibrium concepts discussed below.

2. That is, there are other Bayesian-Nash equilibria, but they are the same along the equilibrium path.

Table 1: Game 1, Unique Nash Pooling

m_1	a_1^n	a_2	a_3	m_2	a_1	a_2	a_3	m_3	a_1	a_2	a_3
t_1	2,1	2,0	0,2	t_1	3,1	1,0	0,0	t_1	1,2	1,1	3,0
t_2	1,3	2,0	2,1	t_2	2,1	0,0	0,6	t_2	0,2	3,1	1,1

ⁿ - Nash equilibrium path.

The game in Table 2 has a unique Bayesian-Nash equilibrium which is separating (each type sends a different message). Hence, upon observing either equilibrium message R can infer S 's type. The equilibrium strategies are $q(m_1|t_1) = q(m_2|t_2) = 1$ and $r(a_1|m_1) = r(a_2|m_2) = r(a_3|m_3) = 1$.

Table 2: Game 2, Unique Nash Separating

m_1	a_1	a_2	a_3	m_2	a_1	a_2	a_3	m_3	a_1	a_2	a_3
t_1	2,3 ⁿ	1,1	2,2	t_1	0,3	0,2	3,1	t_1	1,2	0,3	3,1
t_2	1,1	3,1	0,3	t_2	1,0	2,2 ⁿ	1,1	t_2	1,3	1,2	3,2

ⁿ - Nash equilibrium path.

In games 1 and 2 the Bayesian-Nash equilibrium prediction is unique. However, in many games there is more than one equilibrium, and hence the theoretical prediction is ambiguous. Several game theorists have proposed refinements of the Bayesian-Nash equilibrium concept criteria which select equilibria that satisfy certain rationality restrictions, thus giving more precise behavioral predictions.

The most common criticism of the Bayesian Nash equilibrium concept is that it does not restrict R 's choice of actions "off the equilibrium path," that is, for $m \in M$ such that $q(m|t) = 0 \forall t \in T$. Since the choice of signaling strategy endogeneously determines which messages are off the equilibrium path, "perverse" out-of-equilibrium behavior can generate implausible equilibrium predictions. One selection criterion rules out equilibria in which behavior off the equilibrium path is not optimal according to *some* belief about the type of S . For signaling games this criterion is equivalent to the sequential equilibrium concept of Kreps and Wilson (1982a).

Definition. A *sequential equilibrium* consists of strategies q, r and beliefs $\mu(\cdot|m) \in \Delta_T$ such that conditions i) and iii) of the definition of Bayesian-Nash equilibrium hold and condition ii) is replaced by:

ii) $\forall m \in M, r(a'|m) > 0$ only if

$$\sum v(t, m, a') \cdot \mu(t|m) = \max_{a \in A} \sum_{t \in T} v(t, m, a) \cdot \mu(t|m).$$

Table 3 shows a game in which there are two Bayesian-Nash equilibria, only one of which is sequential. The Bayesian-Nash equilibria which is not sequential is $q(m_1|t_i) = 1, i = 1, 2$ and $r(a_2|m_1) = 1, r(a_2|m_2) = 1, r(a_2|m_3) = 1$. There is no belief over T which makes the action a_2 optimal following m_3 . However, if R plays a response involving only a_1 and a_3 rather than a_2 , S will send the message m_3 rather than m_1 , thus upsetting the equilibrium. The unique sequential equilibrium for this game is $q(m_3|t_i) = 1, i = 1, 2$, and $r(a_2|m_2) = 1, r(a_2|m_3) = 1, r(a_1|m_3) = 1$.

Table 3: Game 3, Sequential vs. Nash

m_1	a_1	a_2^s	a_3	m_2	a_1	a_2	a_3	m_3	a_1^s	a_2	a_3
t_1	1,2	2,2	0,3	t_1	1,2	1,1	2,1	t_1	3,1	0,0	2,1
t_2	2,2	1,4	3,2	t_2	2,2	0,4	3,1	t_2	2,2	0,0	2,1

^s - sequential equilibrium path, ⁿ - Nash equilibrium path.

The sequential equilibrium concept refines Bayesian-Nash equilibria by requiring that out-of-equilibrium responses be supported by *some* belief. However, these beliefs might themselves be unreasonable. The equilibrium refinements proposed by Cho and Kreps (1987) and Banks and Sobel (1987) require beliefs to reflect some thought about which types are likely to benefit from a particular defection.

Consider the game in Table 4. There is a sequential equilibria with $q(m_2|t_i) = 1, i = 1, 2$, and $r(a_1|m_1) = r(a_3|m_2) = r(a_2|m_3) = 1$. To support this equilibrium $\mu(t_1|m_1)$ must be greater than $\frac{2}{3}$; R must believe that the out-of-equilibrium message m_1 is more likely to have been sent by t_1 than by t_2 . (Recall that we assume priors $P(t_1) = P(t_2) = .5$.) However, the equilibrium payoff of a t_1 type from choosing m_2 is greater than any possible payoff from defecting by choosing m_1 whereas an $a't_2$ type could conceivably do better by defecting to m_1 than by sending the equilibrium message m_2 . Hence, it seems unreasonable to *increase* the probability placed on t_1 after observing a defection m_1 . Since t_1 could not possibly benefit and t_2 might, it seems reasonable to believe a defection was from t_2 (i.e., $\mu(t_2|m_1) = 1$). Cho and Kreps (1987) call this "the intuitive criterion". This belief implies a best response of a_2 which provides an incentive for t_2 to switch and send a message m_1 thus upsetting the equilibrium.

Table 4: Game 4, Intuitive vs. Sequential

m_1	a_1	a_2^i	a_3	m_2	a_1	a_2	a_3^s	m_3	a_1	a_2	a_3
t_1	0,3	2,2	2,1	t_1	1,2	2,1	3,0	t_1	1,6	4,1	2,0
t_2	1,0	3,2	2,1	t_2	0,1	3,1	2,6	t_2	0,0	4,1	0,6

i - intuitive equilibrium path, s - sequential equilibrium path.

The intuitive criterion can be stated more formally. Fix a sequential equilibrium with the associated equilibrium payoffs $u^*(t)$ for S , and for each out-of-equilibrium message define

$$T_I(m) = \{t \in T : u^*(t) > \max_{r \in BR(\Delta_r, m)} u(t, m, r)\},$$

(where the subscript I denotes $T(m)$ defined by the intuitive criterion) and let

$$\Delta_I(m) = \{\lambda \in \Delta_I : \lambda(t) > 0 \text{ only if } t \notin T_I(m)\}.$$

Definition. A sequential equilibrium satisfies the *intuitive criterion* if for all out-of-equilibrium messages m , $\mu(\cdot | m) \in \Delta_I(m)$.

For the signaling game in Table 4 the unique sequential equilibrium satisfying the intuitive criterion is: $q(m_1 | t_i) = 1$, $i = 1, 2$, and $r(a_2 | m_1) = r(a_1 | m_2) = r(a_1 | m_3) = 1$.

The logic of the intuitive criterion can be extended in various ways. Consider the game in Table 5 which has the following sequential equilibrium: $q(m_2 | t_i) = 1$, $i = 1, 2$, and $r(a_3 | m_1) = r(a_3 | m_2) = r(a_2 | m_1) = 1$. This equilibrium satisfies the intuitive criterion since $T_I(m_1) = \emptyset$, implying that $\mu(t_1 | m_1) = 1$ is contained in $\Delta_I(m_1)$, and a_3 is a best response to $\mu(t_1 | m_1) = 1$, while $T_I(m_3) = \{t_1\}$, and a_2 is a best response to $\mu(t_2 | m_3) = 1$. However, given the equilibrium payoffs t_2 would like to defect from the equilibrium and send the message m_1 for any possible response that would make t_1 want to defect, but there are some responses (e.g., a_2) for which t_1 does not want to defect when t_2 does. That is, the set of responses to m_1 which induce t_1 to defect are strictly contained in the set which induces t_2 to defect. It seems reasonable to require that the belief $\mu(\cdot | m_1)$ should assign greater weight (relative to the prior) to t_2 when m_1 is observed. This restriction is implied by the concept of divine equilibria (Banks and Sobel, 1987).

Table 5: Game 5, Divine vs. Intuitive

m_1	a_1	a_2	a_3	m_2	a_1	a_2	a_3^i	m_3	a_1^d	a_2	a_3
t_1	4,0	0,3	0,4	t_1	2,0	0,3	3,2	t_1	2,3	1,0	1,2
t_2	3,4	3,3	1,0	t_2	0,3	0,0	3,2	t_2	4,3	0,4	3,0

i - intuitive equilibrium path, d - divine equilibrium path.

Fix an equilibrium with S payoffs $u^*(t)$ and let m be an out-of-equilibrium message. For all $r \in \Delta_A$ define

$$\bar{\mu}(t, r : m) = \begin{cases} 1 & \text{if } u(t, m, r) > u^*(t) \\ [0, 1] & \text{if } u(t, m, r) = u^*(t) \\ 0 & \text{if } u(t, m, r) < u^*(t) \end{cases}$$

as the frequency that $t \in T$ would send m if m induces a response of r . Let

$$\Gamma(r, m) = \{\gamma \in \Delta_T : \exists \mu(t) \in \bar{\mu}(t, r) \text{ and } C > 0 \text{ s.t. } \gamma(t) = C \cdot \mu(t) \cdot P(t) \forall t \in T\},$$

where C is a constant normalizing the expression into the form of a probability. $\Gamma(r, m)$ is the set of beliefs consistent with R responding to m with r where each t has the option of obtaining $u^*(t)$ or $u(t, m, r)$. Finally for any set $\Delta \subseteq \Delta_A$ let

$$\bar{\Gamma}(\Delta, m) = \text{convex hull} \left[\bigcup_{r \in \Delta} \Gamma(r, m) \right].$$

Note that $\bar{\Gamma}(\Delta_A, m) \subseteq \Delta_I(m)$ so that restricting beliefs to be in $\bar{\Gamma}(\Delta_A, m)$ implies that any resulting equilibrium will satisfy the intuitive criterion. Further, if there exists $t, t' \in T$ such that $\bar{\mu}(t, r, m) = 1$ implies $\bar{\mu}(t', r, m) = 1$, then for all $\mu \in \bar{\Gamma}(\Delta_A, m)$, $\frac{\mu(t' | m)}{\mu(t | m)} \geq \frac{P(t')}{P(t)}$, thus satisfying the logic described above.

If it is common knowledge that R holds beliefs in $\bar{\Gamma}(\Delta_A, m)$, then S should expect the message m to induce an action in $BR(\bar{\Gamma}(\Delta_A, m), m)$. This suggests the following iterative procedure: let

$$\Gamma_0 = \Delta_T, A_0 = \Delta_A, \text{ and } \forall n > 0,$$

$$\Gamma_n(m) = \begin{cases} \bar{\Gamma}(A_{n-1}, m) & \text{if } \bar{\Gamma}(A_{n-1}, m) \neq \emptyset \\ \Gamma_{n-1}(m) & \text{else} \end{cases}$$

where $A_n = BR(\Gamma_n(m), m)$ and $\Delta_d(m) = \bigcap_n \Gamma_n(m)$.

Definition. A sequential equilibrium is *divine* if for all out-of-equilibrium messages $m, \mu(A | m) \in \Delta_d(m)$.

For the game in Table 5 the divine equilibrium is: $q(m_3 | t_i) = 1, i = 1, 2$, and $r(a_3 | m_1) = r(a_2 | m_2) = r(a_1 | m_3) = 1$.

The set of beliefs in $\Delta_d(m)$ clearly can depend on the prior belief $P(\cdot)$, since divinity may only require that the beliefs assign greater (or equal) weight than the prior on types more likely to defect. The requirement that beliefs assign *all* positive probability to those types which are most likely to defect is called "universal divinity".³

3. This is equivalent to requiring the out-of-equilibrium beliefs to be divine for all possible priors.

Consider the signaling game in Table 6. There exists a divine equilibrium of the following form: $a(m_3|t_i) = 1, i = 1, 2$ and $r(a_2|m_1) = r(a_1|m_2) = r(a_1|m_3) = 1$. Although t_2 is more likely to defect to m_1 than t_1 since both types could potentially gain from a defection the prior belief $P(t_i) = \frac{1}{2}$ is in the set $\Delta_d(m_1)$. Since this belief supports the equilibrium action a_2 the divinity criterion is satisfied. However, universal divinity requires us to place positive probability only on t_2 because t_2 wants to defect whenever t_1 does, but not vice versa (e.g., t_2 would defect, but t_1 would not if $r(a_1|m_1) = 4, r(a_2|m_1) = .6$). That belief implies a best response by R of a_1 breaking the divine equilibrium. The unique universally divine equilibrium for the game in Table 6 is: $q(m_1|t_i) = 1, i = 1, 2$, and $r(a_2|m_1) = r(a_1|m_2) = r(a_3|m_3) = 1$.

Table 6: Game 6, Universally Divine vs. Divine

m_1	a_1	a_2^u	a_3	m_2	a_1	a_2	a_3	m_3	a_1^d	a_2	a_3
t_1	4,1	2,4	1,5	t_1	1,3	3,1	4,2	t_1	3,3	0,0	0,4
t_2	5,6	2,5	2,2	t_2	1,3	1,4	3,3	t_2	3,4	1,5	0,1

^u - universally equilibrium path, ^d - divine equilibrium path.

As above, fix an equilibrium and let $m \in M$ be an out-of-equilibrium message. Let $\bar{\Delta}_T = \{P \in \Delta_T : P(t) > 0 \forall t \in T\}$ be the set of all nondegenerate priors over T . Define

$$\Delta_u(m) = \bigcap_{P \in \bar{\Delta}_T} \Delta_d(m).$$

Definition. A sequential equilibrium is *universally divine* if for all out-of-equilibrium messages $m, \mu(\cdot|m) \in \Delta_u(m)$.

Consider the signaling game in Table 7. There exists a universally divine equilibrium where $q(m_3|t_i) = 1, i = 1, 2$, and $r(a_1|m_1) = r(a_3|m_2) = r(a_2|m_3) = 1$. At message m_2 there are responses by R which make t_1 want to defect while t_2 does not (e.g., $r(a_2|m_2) = \frac{2}{3}, r(a_3|m_2) = \frac{1}{3}$ and there are responses such that t_2 wants to defect while t_1 would not (e.g., $r(a_1|m_2) = 1$). Hence, universal divinity does not place any restrictions on $\mu(\cdot|m_2)$. The set of best responses following m_2 which sustain the equilibrium are $r(a_1|m_2) = 0, r(a_2|m_2) \leq \frac{1}{2}, r(a_3|m_2) \geq \frac{1}{2}$ (no belief supports mixing between a_1 and a_3). The interior boundary of this set is $r(a_2|m_2) = \frac{1}{2}, r(a_3|m_2) = \frac{1}{2}$; this strategy makes t_1 indifferent between staying along the equilibrium path and defecting while t_2 would prefer the equilibrium path. Hence, we can think of t_1 as the most likely type to defect if we define "most likely" as those types for which the set of equilibrium strategies includes an out-of-equilibrium strategy for which the type has a weak-best response of staying along the equilibrium path.

Table 7: Game 7, Universally Divine vs. NWBR

m_1	a_1	a_2	a_3	m_2	a_1^N	a_2	a_3	m_3	a_1	a_2^u	a_3
t_1	2,2	0,3	5,2	t_1	1,6	5,3	1,0	t_1	2,1	3,3	0,4
t_2	0,2	2,0	5,1	t_2	4,0	4,1	0,2	t_2	1,4	3,3	2,1

N - NWBR equilibrium path, u - universally divine equilibrium path.

Fix an equilibrium where m is an out-of-equilibrium message and define

$$T_n(m) = \{t \in T : \forall r \in BR(\Delta_T, m) \text{ s.t. } \bar{\mu}(t, r) = [0, 1] \exists t' \in T \text{ s.t. } \bar{\mu}(t', r) = 1\}.$$

If $t \in T_n(m)$ then if t is ever indifferent between staying along the equilibrium path and defecting there always exists some other type t' which strictly prefers to defect. In the above example $t_2 = T_n(m_2)$, since t_2 is only indifferent when $r(a_2 | m_2) = \frac{3}{4}$, $r(a_3 | m_2) = \frac{1}{4}$, and this strategy gives t_1 the incentive to defect. Define

$$\Delta_n(m) = \{\lambda \in \Delta T; \lambda(t) > 0 \text{ only if } t \notin T_n(m)\}.$$

Definition. A sequential equilibrium satisfies the *never-a-weak-best response* (NWBR) criterion if for all out-of-equilibrium messages m $\mu(\cdot | m) \in \Delta_n(m)$.

The unique sequential equilibrium satisfying NWBR for the signaling game in Table 7 is: $a(m_2 | t_i) = 1$, $i = 1, 2$, and $r(a_2 | m_1) = r(a_1 | m_2) = r(a_3 | m_3) = 1$.

Universal divinity and NWBR were initially attempts to characterize the restrictiveness of the concept of stable equilibrium (Kohlberg and Martens, 1986) in signaling games. The stable equilibrium concept requires that every possible "tremble" of strategies have an equilibrium "close to" the candidate equilibrium. Since for signaling games there is a one-to-one relationship between trembles of signaling strategies and beliefs we can discuss stability in terms of every possible out-of-equilibrium belief generating a nearby equilibrium. Table 8 consists of a game which demonstrates that NWBR does not completely capture stability. Consider the following sequential equilibrium: $q(m_1 | t_i) = 1$, $i = 1, 2$, and $r(a_2 | m_1) = r(a_1 | m_2) = r(a_2 | m_3) = 1$. This equilibrium satisfies NWBR. At m_2 there exists a response strategy making t_1 indifferent while t_2 prefers the equilibrium, namely $(r(a_1 | m_2) = \frac{1}{2}, r(a_2 | m_2) = \frac{1}{2})$. Thus, NWBR places no restrictions on $\mu(\cdot | m_2)$. (At m_3 it is clear that NWBR implies $\mu(t_2 | m_3) = 1$ supporting the equilibrium path.)

Table 8: Game 8, Stable vs. NWBR

m_1	a_1	a_2^N	a_3	m_2	a_1	a_2^S	a_3	m_3	a_1	a_2	a_3
t_1	1,6	2,5	2,0	t_1	0,5	3,4	1,2	t_1	4,2	1,1	0,3
t_2	2,0	2,5	0,6	t_2	1,2	3,4	0,5	t_2	1,2	0,4	3,3

^s - stable equilibrium path, ^N - NWBR equilibrium path.

To see whether this equilibrium is stable we focus on trembles at m_2 . If a tremble generates a belief $\mu(t_1 | m_2) \geq \frac{2}{3}$, then a_1 is a best response and the equilibrium path is supported. If $\mu(t_1 | m_2) \leq \frac{1}{3}$ then a_3 is a best response, and again the equilibrium path is supported. Suppose a tremble induces a belief $\mu(t_1 | m_2) \in (\frac{2}{3}, \frac{1}{3})$; if neither type sends m_2 with positive probability R responds with a_2 thus upsetting the equilibrium. What is needed is the ability for one or the other type to send m_2 with sufficient probability to make R indifferent between two actions and R to mix between the actions in such a way as to make the type(s) indifferent between m_1 and m_2 thus rationalizing the original mixing. Hence t_1 or t_2 must send m_2 in such a way as to induce a belief of either $\mu(t_1 | m_2) = \frac{1}{3}$ or $\mu(t_1 | m_2) = \frac{2}{3}$. At the former R could mix between a_2 and a_3 and such a mix could leave t_1 indifferent while t_2 prefers the equilibrium path. At the latter R could mix between a_1 and a_2 and such a mix could leave t_2 indifferent while t_1 prefers the equilibrium path. For a tremble inducing a belief in $(\frac{1}{3}, \frac{2}{3})$ if t_1 sends m_2 this would push the belief to $\mu(t_1) = \frac{2}{3}$; however, this is the belief generating a response leaving t_2 indifferent. A similar conclusion holds if t_2 would send m_2 . Thus, a tremble inducing a belief in $(\frac{1}{3}, \frac{2}{3})$ cannot be *stabilized* by a judicious choice of signaling strategy. Note that a necessary condition for such a stabilization in this game is that both types be on the boundary of the set of sequential equilibria; that is, stability is a refinement of the NWBR equilibrium concept.

Fix an equilibrium with m being an out-of-equilibrium message and for all $J \subset T$ define

$$I(J, m) = \{r \in \Delta_A : u^*(t) \geq u(t, m, r) \forall t \in T \text{ and } u^*(t) = u(t, m, r) \text{ if } t \in J\};$$

the set $I(J, m)$ contains those responses making types in J indifferent between the equilibrium path and m while those in $T \setminus J$ prefer the equilibrium path. If a tremble induces a belief μ at m and only NWBR types in J send m with positive probability then R 's posterior belief at m will be a convex combination of μ and $e(t)$. Define

$$\hat{\Lambda}(J, r, m) = \{\lambda \in \hat{\Delta}_T : \exists \lambda' \in \Delta_T \text{ where } r \in BR(\lambda', m) \text{ s. t. } \lambda = \sum_{t \in J} \alpha(t) \cdot e(t) + \beta \cdot \lambda' \text{ for } \alpha(t) \geq 0,$$

$$1 - \sum_{t \in J} \alpha(t) = \beta > 0\};$$

$\hat{\Lambda}(J, r, m)$ thus consists of the beliefs that cannot be stabilized by types in J sending m if R responds with r . Define

$$\Lambda(J, m) = \begin{cases} \bigcap_{r \in I(J)} \hat{\Lambda}(J, r, m) & \text{if } I(J) \neq \emptyset \\ \Delta_T & \text{else} \end{cases}$$

and set $\Lambda^*(m) = \bigcap_{J \subset C} \Lambda(J, m)$. Thus $\Lambda^*(m)$ consists of those trembles in beliefs that cannot be stabilized at the out-of-equilibrium message m . Banks and Sobel (1987) and Cho and Kreps (1987) provide the following characterization result.

Theorem. A sequential equilibrium is *stable* if and only if for all out-of-equilibrium messages m , $\Lambda^*(m) = \emptyset$.

In the above example the NWBR equilibrium is not stable since $\Lambda^*(m_2) = (\frac{1}{3}, \frac{2}{3}) \neq \emptyset$. The unique-stable equilibrium in this game is $q(m_2 | t_i) = 1$, $i = 1, 2$ and $r(a_2 | m_1) = r(a_2 | m_2) = r(a_2 | m_3) = 1$.

In summary we have the following nesting of equilibrium concepts: Bayesian Nash \supseteq Sequential \supseteq Intuitive \supseteq Divine \supseteq Universally Divine \supseteq NWBR \supseteq Stable.

Since Kohlberg and Mertens (1986) prove that every game has at least one stable equilibrium and the stable equilibrium is also NWBR, universally divine, etc., existence of each kind of less-refined equilibrium is also guaranteed.

3. EXPERIMENTAL DESIGN

Our experimental design uses the games in Tables 1-8 to determine which refinements subjects play most often. Games 1 and 2 test whether subjects play unique Nash equilibria, one pooling (game 1) and one separating (game 2). Games 3 to 8 all have two pooling Nash equilibria, one more refined and one less refined. In addition to the two messages which are equilibrium choices in games 3 to 8 we include a third message which is not an equilibrium choice. Each game then has Nash and non-Nash messages so we can test the robustness of Nash play in several different games.

Our experiment has three goals:

1. Test for robustness of Nash equilibrium play in several different games.
2. Test whether subjects play more refined equilibria (for several different nested refinements).
3. Test whether decision criteria (other than Nash equilibrium) can explain individual choices.

a. *The Experimental Session*

The games were presented to subjects in two 3 x 3 payoff tables (see Appendix A for an example). Each of their choices were made and communicated on a computer network.

An experimental session consisted of many periods. In each period six subjects were divided into three pairs. One subject in each pair was the sender and the other subject was the receiver. To make each play as much like a one-shot game as possible, pairings were random and anonymous each period.

Subjects assigned to be senders were told their type (1 or 2) which was randomly determined. Receivers did not know the sender's type, but the two types were commonly known to be equally likely. After receivers were told their sender's message (m_1, m_2 or m_3) they chose an action (a_1, a_2 or a_3). A period ended when the receiver picked an action and all results (type, message, action and payoff) were transmitted to both players.

In each experimental session several games were played for 10 consecutive periods each. Subjects knew the payoff matrix on their screen was the same for everyone and would be used for the ten periods. Subjects were given history of their own paired plays, but *not* the entire cohort history of plays. An experimental session took two hours and lasted 30-50 periods; i.e., ten periods each of three to five different games. Subjects had unlimited time to make decisions. Typically the first period of a new game took about five minutes to finish; later periods took one to two minutes each. Subjects earned \$.25 for each payoff point minus \$5.00. Their earnings averaged about \$20.00. Payments were made privately at the end of each session

Note that the equilibrium predictions which assume payoffs in Tables 1 through 8 are units of utility. To apply the predictions to our experiment requires a method for inducing risk neutrality (e.g., Roth and Malouf, 1979) or the assumption that subjects are risk neutral for gambles involving about a dollar. Since we share the reservations of many experimenters about the risk-induction procedure (e.g., Cox, Smith, and Walker, 1985) we simply assume risk neutrality.⁴

b. *Treatment Variables*

We conducted 13 experimental sessions. Each session has two important learning variables. First, within an experimental session the sequence in which different games were played might affect learning. (Subjects were usually most confused and made fewer equilibrium choices in the first game in a session.) Second, we expect some learning across the ten plays of a given game. To check for sequence effects we varied the order in which the games were played in each session. Each game was played first in a session at least once.

To distinguish behavior of subjects with different potential mathematical sophistication we used four subject pools. We ran experiments using students at the California Institute of Technology and the Universities of Arizona and Pennsylvania, and members of the technical staff at the Jet Propulsion Laboratory. Table 9 lists the experiments we conducted.

4. In most of the games the equilibria still hold under risk aversion unless subjects are sufficiently risk-averse that they prefer a payoff of $$.25n$ to an even gamble between 0 and $$.25(2n + 1)$ (for $0 \leq n \leq 3$). Studies of choice indicate subjects are roughly risk neutral for bets of this size; e.g., Camerer (1988).

Table 9: Experimental Design

Experiment Number	Subject Pool	Sequence of Games
1	Caltech	8,2,4,6
2	Penn, Ph.D.	3,2,6,4
3	Caltech*	1,3,4
4	JPL*	2,8
5	JPL	3,6,1,4,8
6	Penn, undergrad.*	1,8
7	Penn, undergrad.*	6,8,2,3
8	Caltech	4,8,6,1
9	Arizona	1,4,1
10	Penn, undergrad.	5,7,9,1**
11	Penn. undergrad.	9,7,5,3
12	Caltech	7,9,5,4**
13	Caltech	9,5,7,6**

* These sessions included an unreported game. ** These games were conducted for 20 periods.

NOTE: Game 9 is the same as Game 8 except that we rearranged the payoff matrix so that the optimal action for receivers was not the same for both equilibria. We made this change after observing the strength of the stable equilibrium in Game 8. However, the results were indistinguishable (see the logit analysis in Appendix A and Figures 1 and 2). Therefore, we will pool all data from Games 8 and 9.

4. EXPERIMENTAL RESULTS

a. *Nash Equilibrium Behavior*

First we look at how often outcomes constitute a Nash-equilibrium message-action pair. Table 10 shows the proportion of Nash outcomes aggregated across all games. About two-thirds of the outcomes are Nash. We can easily reject the hypothesis that choices are random and in most games there is some convergence toward Nash outcomes between periods 1-5 and 6-10.⁵

5. There is also a slight sequence effect: The overall proportions of Nash outcomes in periods one through five and six through ten for games played first in an experimental session were .64 and .71.

Table 10 - Frequency of Nash Outcomes

Periods	Percent of Outcomes	If first in Sequence	Number of Obs
1-5	66		645
		64	65
6-10	70		645
		71	65

However, the time series and summary statistics suggest the amount of Nash behavior is not consistent across games. Table 11 shows Nash and non-Nash behavior by game and refinement. The Bonferonni X^2 statistics test the hypothesis that the proportions of Nash outcomes are equal in all games. We cannot reject the equality hypothesis in periods 1-5 ($X^2 = 10.95$, $P = .19$), but we can reject it ($X^2 = 32.97$, $P = .01$) in periods 6-10.

The variation in amount of Nash play is hard to explain parsimoniously. Logit analyses of the subject pool and learning variables (see Appendix B) suggest few systematic effects which can explain the variation in Nash play across games. Whether equilibria are unique appears helpful. Table 11 shows nearly 80% of the outcomes are Nash in later periods of games 1 and 2, but game 3 has two equilibria and even more (95%) of the plays are Nash. Or one might conjecture there is less Nash play in higher-numbered games with deeper refinements, but games 6 through 8 yield more Nash play than game 5.

b. *Refinement Results*

Table 11 shows the fraction of responses consistent with each refinement. There is a lot of non-Nash play, but there is also some tendency to choose the more refined equilibrium. That tendency also grows stronger between periods one to five and six to ten. The time series graphs of outcomes for each game are in Appendix C.

Table 11
Contingency Table of Outcomes by Refinement

Game	Periods	Proportion of Observed Outcomes			Sample Size
		More Refined	Less Refined	Non-Nash	
1	1 - 5	<i>Nash</i> .56		<i>Non-Nash</i> .44	75
	6 - 10	.76		.24	75
2	1 - 5	<i>Nash</i> .68		<i>Non-Nash</i> .32	75
	6 - 10	.81		.19	75
3	1 - 5	<i>Sequential</i> .63	<i>Nash</i> .12	<i>Non-Nash</i> .25	60
	6 - 10	.73	.22	.05	60
4	1 - 5	<i>Intuitive</i> .53	<i>Sequential</i> .13	<i>Non-Nash</i> .34	90
	6 - 10	.68	.03	.29	90
5	1 - 5	<i>Divine</i> .37	<i>Intuitive</i> .22	<i>Non-Nash</i> .41	60
	6 - 10	.43	.10	.47	60
6	1 - 5	<i>U-Divine</i> .28	<i>Divine</i> .32	<i>Non-Nash</i> .40	75
	6 - 10	.37	.32	.31	75
7	1-5	<i>NWBR</i> .30	<i>U-Divine</i> .35	<i>Non-Nash</i> .35	60
	6 - 10	.20	.43	.37	60
8,9	1 - 5	<i>Stable</i> .59	<i>NWBR</i> .13	<i>Non-Nash</i> .28	150
	6 - 10	.56	.06	.38	150

Figure 1:
Confidence Regions for Outcome Proportions
periods 1-5

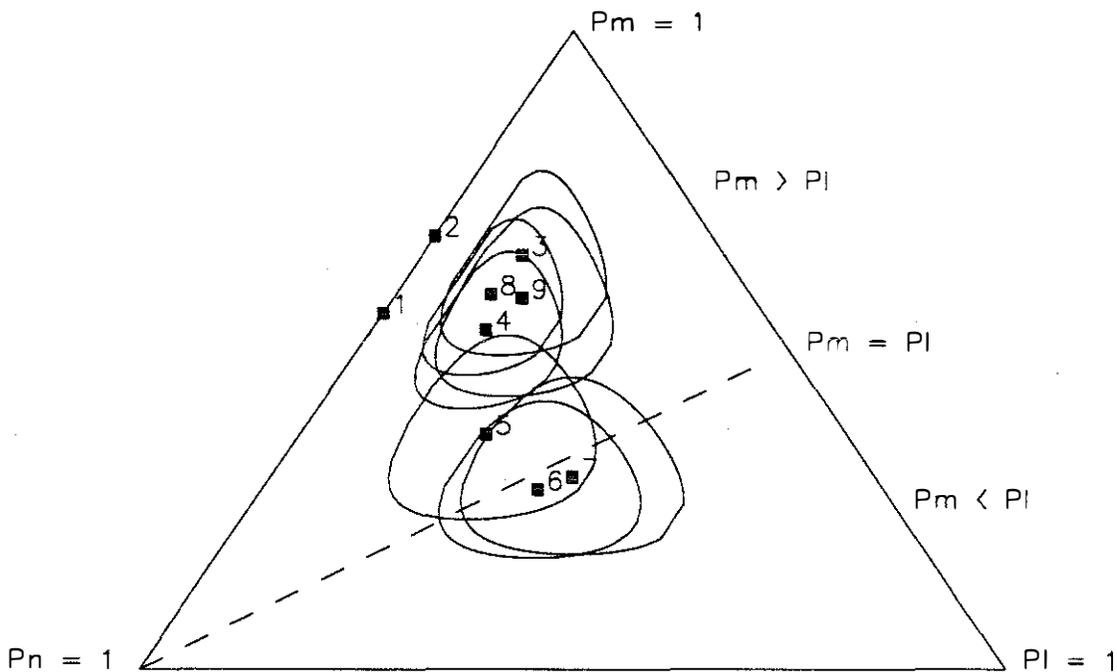
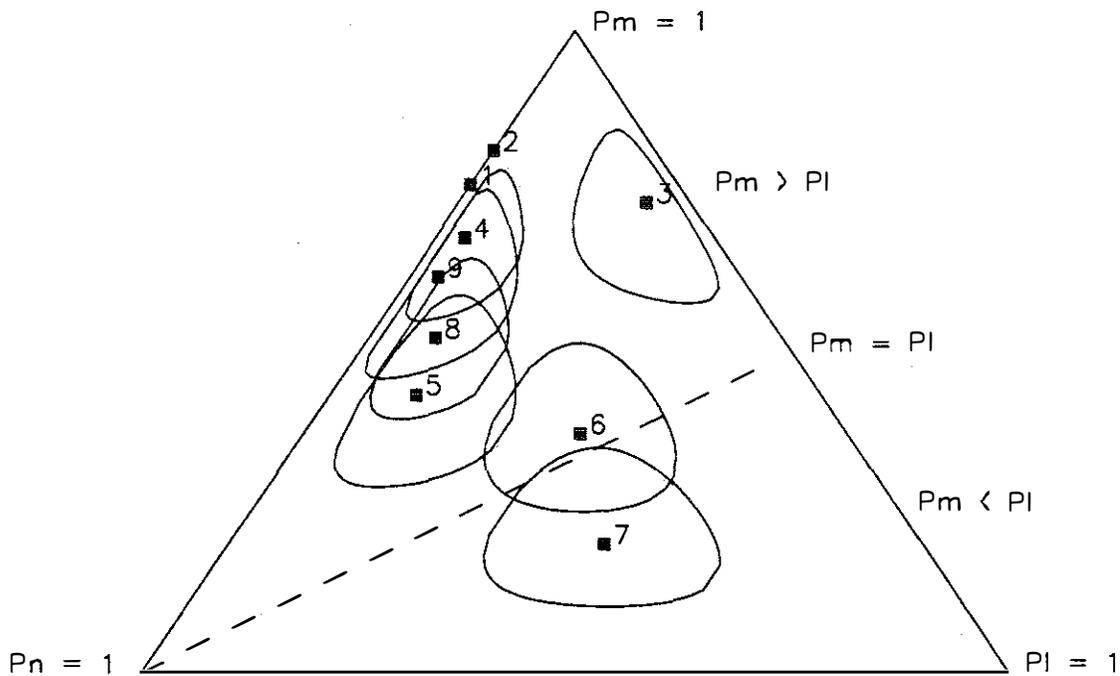


Figure 2:
Confidence Regions for Outcome Proportions
periods 6-10



at the 5% level.

In earlier periods (Figure 1) the confidence regions are centered near the middle of the simplex where more-refined and less-refined play are equally likely. But in later periods (Figure 2) the confidence regions for some games move toward the left edge of the triangle where $P_l = 0$ while games 6 and 7 stay near $P_m = P_l$ and game 3 moves toward $P_n = 0$.

One is tempted to conclude that subjects simply converge toward more refined equilibria up to a point--games 6 and 7 and game 8 is an outlier. This interpretation is wrong because of the nesting of refinements. In game 4, for instance, 68% of the outcomes in later periods correspond to the intuitive criterion. But both equilibria in game 5 are intuitive (and hence sequential and Nash) and they are only played a total of 53% of the time. Thus the refinement which predicts well in game 4 works less well in game 5. (This is simply our point about Nash outcomes--their frequency varies across games--extended to the refinements.)

Similarly, the stable equilibrium is played 56% of the time in later periods of game 8. The nesting implies that in all other games the more refined equilibrium is stable and the less refined equilibrium is not, so every game is a test of stability. The stable equilibrium is played two-thirds of the time in games 3--4, but only a third of the time in games 5--7.

Thus the frequency of play of various Nash refinements varies across games rather mysteriously just as Nash play does. One way to explain the variation is to examine sender messages and receiver actions separately. It might be that senders always choose the more-refined equilibrium message, but receivers sometimes choose the best-response action from the less-refined equilibrium. (In our analysis so far such a message-action pair would be classified as a non-Nash outcome.)

c. *Sender Behavior*

A sender in our signaling games must make a decision knowing his type, but not knowing (with certainty) the reaction of his receiver counterpart. Except for game 2 all our equilibria require senders to pool -- ignore their type and choose the same message.

Table 13 shows the proportion of sender messages by game and type in later rounds. Only 10-20% of the messages are not Nash. However, there is a clear dependence of message choice on sender type in games 4 to 7, contrary to the pooling equilibrium prediction. For example, in game 4 50% of the t_1 senders chose the intuitive message, but 90% of the t_2 senders chose it.⁷ Brandts and Holt (1988) also observed type dependence in their experimental tests of the Cho-Kreps (1987) "beer-quiche" game which pits an intuitive equilibrium against an unintuitive one.

7. The logit model for sender messages in Appendix B suggests that: i) Type dependence relative to the non-Nash (less refined) message is game specific and significant in games 1, 4, 5 and 7. ii) Subject pool and sequence have no systematic affect on the outcomes. Learning is game specific; when it is significant it increases the probability of Nash message relative to non-Nash messages.

Table 13
Proportions of Sender Messages by Type (t_1, t_2) for periods 6 - 10.

Game	Proportion of Observations			Sample Size
1	<i>Nash</i> (.75, .95)		<i>Non-Nash</i> (.25, .05)	(36, 39)
2	<i>Nash</i> (.94, .90)		<i>Non-Nash</i> (.06, .10)	(33, 42)
3	<i>Sequential</i> (.81 ^b , .76)	<i>Nash</i> (.19, .24)	<i>Non-Nash</i> (.00, .00)	(31, 29)
4	<i>Intuitive</i> (.50, .90 ^b)*	<i>Sequential</i> (.18 ^b , .08)	<i>Non-Nash</i> (.32, .02)*	(40, 50)
5	<i>Divine</i> (.56, .54 ^b)*	<i>Intuitive</i> (.34 ^b , .00)*	<i>Non-Nash</i> (.10, .46)*	(32, 28)
6	<i>U-Divine</i> (.37, .71)*	<i>Divine</i> (.50 ^b , .29 ^b)	<i>Non-Nash</i> (.13, .00)*	(38, 37)
7	<i>NWBR</i> (.28, .52 ^b)*	<i>U-Divine</i> (.62 ^b , .48)	<i>Non-Nash</i> (.10, .00)	(29, 31)
8	<i>Stable</i> (.78 ^b , .78 ^b)	<i>NWBR</i> (.09, .04)	<i>Non-Nash</i> (.12, .18)	(72, 78)

* Rejection at the .05 level of the hypothesis that conditional on the message being sent, the probability is equal for each type, against the alternative of unequal probabilities.

^b Denotes message which is *best* for the sender of each type, in equilibrium.

The equilibrium message which is better for (i.e., yields a higher payoff for) a particular sender type is marked with a superscript "b" in Table 13. Sender types generally choose the better-for message, though not always. Type-dependence occurs if one equilibrium is better for type one, the other equilibrium is better for type two, and senders choose which is better for their type. But even the better-for rule is sometimes violated, as in type two's in game 6. Type-dependence due to the better-for rule helps explain why sender play is game specific. Other decision rules might help explain deviations too. We considered several decision criteria:⁸

8. We considered several other decision criteria: Pareto-optimality, maximize-total-payoff (tacit cooperative strategy), and lexicographic preference for higher payoffs subject to minimizing payoff difference to reflect equity concerns. (The latter

- n. Nash message
- m. maximum message (maximizes sender's minimum payoff); see Luce and Raiffa (1957, p. 278)
- p. principle of insufficient reason (sender regards each receiver action as equally likely)
- e. empirical-best response (selects message which is a best response to the message-dependent actions that were actually played in periods six to ten).

Table 14 shows the proportion of message choices by sender type along with letters marking which message is selected by each criterion. No criterion or pair of criteria accounts for deviations from Nash messages especially well. However, when all the criteria (including better-for, not shown here in Table 14) select a message then senders choose it more than 80% of the time.

criterion often picks weakly-dominated strategies.) None of these criteria helped explain deviations from Nash.

Table 14
 Comparison of Sender Decision Rules Periods 6-10
 (Raw Data Messages)

Game	Type					
	m_1	1 m_2	m_3	m_1	2 m_2	m_3
1	ne 27	p 8	m 1	nmpe 37	2 2	0
2	nmpe 31	1	1	2	nmpe 38	m 2
3	n * 5	0	nmpe 26	np 7	0	nme 22
4	n 20	nme 7	mp 14	nmpe 45	n 4	1
5	3	np 12	nme 17	mp 13	n 0	ne 15
6	nmpe 15	m 4	ne 19	nmpe 26	* 0	n * 11
7	e 3	nmp 8	n 18	* 0	np 16	nme 15
8	nmpe 7	n 56	* 9	nm 5	nm 58	mpe 15

* These strategies are *weakly* dominated if the sender eliminates the *weakly* dominated strategies of the receiver.

d. *Receiver Behavior*

In our games once receivers get a message they must choose an action without knowing the sender's type with certainty. Table 15 shows the number of receiver actions for each possible message in periods six to ten. Decision-criteria-selection actions are marked with letters. We denote the more-refined-message-action pairs by a box and the less-refined with a circle (deleting "n")

markings), adding "d" to denote a weakly-dominated action.⁹

Table 15
Comparison of Receiver Decision Rule Periods 6-10

<i>Game</i>	<i>m</i> ₁			<i>m</i> ₂			<i>m</i> ₃		
	<i>a</i> ₁	<i>a</i> ₂	<i>a</i> ₃	<i>a</i> ₁	<i>a</i> ₂	<i>a</i> ₃	<i>a</i> ₁	<i>a</i> ₂	<i>a</i> ₃
1	mpe 57	d 0	m 6	m 5	d 0	pe 6	mpe 1	d 0	d 0
2	e 28	d 0	pm 4		mpe 38	d 0	m 2	mpe 1	d 0
3	d 0	mpe 13	m 0	m 0	pe 0	d 0	mpe 45	d 0	dm 2
4	2	mpe 63	d 0	mpe 14	dm 0	4	pe 7	m 0	0
5	me 4	p 11	1	0	e 5	mp 6	mpe 29	4	d 0
6	9	mpe 31	0	mpe 5	0	d 0	mpe 24	3	3
7	m 2	pe 0	d 1	pe 12	m 2	10	3	mpe 26	4
8	0	mpe 16	0	18	mpe 40	5	d 0	0	mpe 11

Receivers rarely choose dominated actions. Their infrequency is a bit surprising because many of the actions are only dominated by mixed-strategy combinations of other actions. Also, receivers rarely choose actions if they are not empirical-best responses to messages.¹⁰

9. The logit model for receiver actions in Appendix B suggests that; i) subject pool and sequence have no systematic effect, ii) learning is game specific.

10. On the other hand, criteria "p" and "e" often coincide because the belief interval which makes an action a best response

Given that a sender transmits an equilibrium message the receivers do chose equilibrium actions most of the time. It seems that the blame for non-Nash outcomes must rest mostly on the senders, and their tendency to separate in search of equilibrium payoffs that are better for each type.

5. CONCLUSIONS AND EXTENSIONS

We conducted a series of experiments to test whether subjects chose refined subsets of Nash equilibria in signaling games. In the experiments a sender was privately informed of a randomly-drawn type then chose a message. A receiver knew the message, but not the sender's type, and chose an action. Each game was repeated ten times with subjects randomly reassigned in pairs in each repetition. Two of the games had unique Nash equilibria. The six other games each had two equilibria one of which obeyed a more stringent refinement criterion than the other. Our experimental design and treatments were ambitious and exploratory.

We conducted 13 experiments with four subject pools. Our conclusions are as follows:

- a. About 70% of the message-action pairs gave Nash equilibrium outcomes. The frequency of Nash play differed across games.
- b. There was some tendency for subjects to choose the more refined equilibrium, but it depended on the specific game they played. No refinement predicted well in every game.
- c. Even though all games except one predicted pooling equilibria senders often chose different messages depending on their types (they separated). Senders tended to choose the equilibrium message which gave the better payoff for their type which often caused them to separate rather than pool.
- d. We tested whether several simple decision criteria such as minimax and principle of insufficient reason could explain why senders chose non-Nash and unrefined messages. No criterion worked very well, but when several criteria select to a particular message it was picked about 90% of the time.

There are two ways of thinking about equilibrium refinements. An extreme "logical" position is that refining is a search for an ideal set of formal rules which will refine the set of Nash equilibria to a single point in every game. "Intuitionists" take the opposite position, contending that refinements are useless because the plausibility of equilibria is self evident, varies across settings, and cannot be captured in a set of formal rules. The fact that refinements differ in accuracy across games supports the intuitionist position. Our results might be disheartening for logical-school theorists, but they suggest one positive lesson: Always check whether equilibria are consistent with decision criteria like minimax. If they are consistent with several appealing criteria they are more likely to be played.

The game-specificity of refinement choices suggests further research should proceed in three directions:

usually encompasses both $P(t_1) = P(t_2) = .5$ (as the principle of insufficient reason assumes) and the empirically-observed frequency of t_1 .

- (i) Simple refinements like sequentiality and the intuitive criterion predict fairly well in these and other experiments¹¹ games could be constructed in which those refinements make implausible unappealing choices because they conflict with other decision criteria to put those refinements to a tougher test.
- (ii) Further work might suggest decision criteria which explain anomalous choices better than the several criteria we tried. It might help to gather more detailed data; e.g., one could measure subjects' beliefs and elicit contingent strategies (message choices for each possible type and actions for each possible message).
- (iii) Game-specificity of results suggest refinements should be studied in specific institutional settings of economic interest. For instance, the Spence (1974) signaling model would be interesting to experiment with because it permits a wide variety of equilibria and has been widely applied. The Milgrom and Roberts (1982b) limit-pricing model is interesting because an incumbent firm's type can be known to entrants and reputation building still occurs provided the incumbent's type is not commonly known. Other examples are easy to find.

11. See Miller and Plott (1985); Brandts and Holt (1987); Cramerer and Weigelt (1988); Pitchik and Schotter (1988); Cadsby, et al (1989).

Appendix A

Subject's Screen of Game Payoffs
and Instructions

You are about to participate in an experiment designed to provide insight into certain features of decision processes. If you follow the instructions carefully and make good decisions, you might earn a considerable amount of money. You will be paid in cash.

The type of currency used in this market is francs. All transactions will be in terms of francs. Each franc is worth .25 dollars to you. At the end of the experiment 30 francs will be subtracted from your total franc earnings. Your francs will then be converted into dollars at your specified rate, and you will be paid in dollars.

All communication during the experiment will be done through your computer terminal. The experiment will proceed as a series of periods during which you will make decisions and obtain earnings.

At the beginning of a period each participant will be randomly paired with another participant in this experiment. Within each pair one individual will be assigned randomly as X and the other as Y. Each participant assigned as X in a pair will roll a die to determine their *index*. If the numbers 1 through 3 appear on the roll, the index will be I; if the numbers 4 through 6 appear, the index will be II. The actual value of the role will be X's own private information (the outcome of the roll will not be known to Y).

The participant assigned as X will then make a selection (A, B or C). Participants assigned as X will indicate their selections by using the cursor and selecting one of the rows A, B or C.

The participants assigned as Y will then be informed of the selection made by X, and then will choose one of the columns labeled D, E or F. Participants assigned as Y will indicate their selection by using the cursor key and selecting one of the columns D, E or F.

Individual earnings from the selections will be determined from a *Payoff Chart* given to you each period. Each Payoff Chart specifies the amounts (in francs) X and Y will receive depending on the choices made as well as the index. Suppose, for example, that the Payoff Chart had the following entries:

Payoff Chart

	I			II			
	D	E	F	D	E	F	
A	5,6	2,12	3,3	A	6,0	22,1	32,2
B	8,1	4,13	1,12	B	0,12	2,32	0,0
C	21,6	0,0	15,1	C	4,2	21,3	2,23

The first entry in each cell in the Payoff Chart is the amount to be received by X and the second entry is the amount to be received by Y. For example, under index I the cell B, F shows a potential earning of 1 francs for X and 12 francs for Y.

Individual earnings for each period will be determined as follows. Each participant assigned as X will be informed of their index and will be asked to make selection of A, B or C. The paired participant assigned as Y will then be informed of X's choice but *not* the index. Y will then make a choice of D, E or F. The index will then be announced to Y and Y choice will be announced to X.

Given the choices made by the paired participants, individual earnings will be calculated and highlighted on your screen. For example, suppose the Payoff Chart were as above and the paired participant assigned as X selected the letter C and the paired participant assigned as Y selected the letter D. If the index was I then the earnings for X would be 21 and the earnings for Y would be 6. If the index was II the earnings for X would be 4 and the earnings for Y would be 2.

In review, the process is as follows:

1. At the beginning of a period, you will be paired randomly with another participant.
2. You will be assigned randomly as either X or Y.
3. You will be given a Payoff Chart.
4. Each individual assigned as X will roll a die to determine their index.
5. Each X participant will make a choice of A, B or C.
6. Each individual assigned as Y will be informed of X's selection and will make a selection of D, E or F (without knowing the index).
7. The index and selection of participants will be announced.
8. Earnings are calculated and placed in your Record of Earnings Sheet.

Feel free to earn as much as you can. Are there any questions?

Period No. 2 Participant ID #3

Choose - Arrows
History-H Select-Enter

Last Period Payoff: 3

-----I-----
[D] [E] [F] -----II-----
[D] [E] [F]

- [A] 1.6 2.5 2.0 2.0 2.5 0.6
- [B] 0.5 3.4 1.2 1.2 3.4 0.5
- [C] 4.2 1.1 0.3 1.2 0.4 3.3

Your Index is I, please make a selection

Appendix B

Logit Estimates for Games, Sender,
and Receiver Responses

Recall we have used various subject pools, have played the game for several periods, and have played the games in various sequences. To check for these effects across the games we estimate the following unordered logit model for each game.

$$\text{Prob}(Y_i = \text{Non-Nash [Less]})/\text{Prob}(Y_i = \text{More}) = a_0 + b_1 * a_i + b_2 * j_i \\ + b_3 * p_i + b_4 * t_i + b_5 * s_i + b_6 * g_i$$

where: a = dummy for Arizona subject pool
 j = dummy for JPL subject pool
 p = dummy for Penn subject pool
 t = period play of game
 s = sequence play of the game
 g = game 9 dummy for row change

Thus, subject pool estimates are relative to the Caltech subject pool. Table 4 supplies the estimates of the model for each game and joint hypothesis test. The standard error of the estimates is in parentheses.

Logit Estimates for Treatments

Game	Coefficients	P(NonNash)/P(More)					
		a_0	b_1	b_2	b_3	b_4	b_5
1	-0.35 (0.64)	0.35** (.577)	0.39 (.492)	1.04** (.558)	0.177** (.0065)	0.24 (0.18)	na
2	2.12 (1.32)	-1.05 (0.86)	-0.44 (0.90)	1.9 (0.62)	0.156** (0.007)	-0.87 (0.61)	na
3	0.14 (0.94)	na	-0.05 (0.86)	0.37 (0.74)	-0.26** (0.11)	-.34 (0.26)	na
4	-0.81 (0.77)	-0.15 (0.48)	-0.61 (0.63)	-0.67 (0.63)	-0.05 (0.06)	0.22 (0.26)	na
5	1.41 (0.83)	na	na	-0.54 (0.42)	-.0001 (.071)	-0.46 (0.27)	na
6	-2.10 (1.07)	na	1.26 (0.80)	1.16** (0.60)	-0.10 (.07)	0.75** (.276)	na
7	-1.72 (0.92)	na	na	-0.195 (0.50)	0.107 (.087)	0.79** (0.356)	na
8	-0.62 (0.46)	na	0.70 (0.71)	-0.632 (0.52)	0.03 (.046)	-.07 (0.19)	-.07 (.3)

* This is the estimate for the column change in game 8 (game 9)

** Significant at .10 level (Bonferroni Joint Hypothesis test)

Game	P(Less)/P(More)						
	a_0	b_1	b_2	b_3	b_4	b_5	b_6
3	-0.21 (0.91)	na	-2.42** (0.89)	-1.07 (0.56)	0.11 (0.10)	-0.47 (0.29)	na
4	1.46 (1.03)	1.15 (0.48)	3.34** (1.63)	2.98 (1.66)	-0.24** (0.11)	-1.27** (0.56)	na
5	1.07 (1.10)	na	na	0.77 (0.64)	-0.22** (.10)	-0.63 (0.56)	na
6	-0.57 (1.26)	na	1.49** (0.77)	-0.72 (0.75)	0.05 (.07)	-0.06 (0.33)	na
7	-0.51 (0.83)	na	na	0.18 (0.48)	0.124 (.083)	0.11 (0.35)	na
8	-1.541** (0.62)	na	0.56 (0.89)	-1.95** (0.72)	-0.10 (.065)	0.16 (0.24)	0.08 (.5)

For the senders we estimate the following Logit model:

$$\text{Prob}(Y_i = \text{Refinement}) / \text{Prob}(Y_i = \text{Non-Nash}) = b_0 * c_i + b_1 * a_i + b_2 * j_i + b_3 * p_i + b_4 * t_i + b_5 * s_i + b_6 * g_i + b_7 * z_i$$

where: c = dummy for Caltech subject pool
a = dummy for Arizona subject pool
j = dummy for JPL subject pool
p = dummy for Penn subject pool
t = period play of game
s = sequence play of the game
g = dummy for game 8 and 9
z = dummy for type = 1, 2

Logit Estimates for Treatments

Game	P(Less)/P(Non-Nash)							
	b_0	b_1	b_2	b_3	b_4	b_5	b_6	b_7
1	-1.88 (0.77)	-1.06 (0.94)	-1.85 (0.97)	-1.20 (0.72)	0.234 (.0765)	0.328 (0.21)	na	1.59 (0.45)
2	-1.57 (2.42)	-0.15 (1.39)	0.28 (1.45)	-2.83 (2.54)	0.128 (0.11)	1.82 (1.12)	na	-.25 (0.60)
3	-11.2 (486.0)	na	6.011 (627.0)	5.633 (482.6)	-9.0 (61.5)	-4.45 (139.1)	na	-11.61 (258.0)
4	1.52 (1.10)	1.039 (1.93)	4.87 (1.75)	2.91 (1.09)	.090 (.097)	-0.73 (0.40)	na	-.036 (0.66)
5	-0.913 (1.24)	na	na	2.31 (1.09)	-0.16 (0.10)	0.17 (0.40)	na	-4.23 (1.11)
6	-0.64 (1.68)	na	-1.34 (1.26)	-0.13 (1.09)	0.28 (0.11)	-0.11 (0.44)	na	12.39 (74.6)
7	14.06 (73.6)	na	na	10.65 (49.1)	0.139 (0.14)	-4.71 (24.54)	na	1.34 (1.13)
8	-0.94 (0.75)	na	-1.14 (0.85)	-.016 (1.13)	-.086 (.077)	0.72 (0.31)	-4 (.5)	.006 (0.44)

Game\Coefficients	P(More)/P(Non-Nash)							
	b_0	b_1	b_2	b_3	b_4	b_5	b_6	b_7
3	0.19 (0.85)	na	1.89 (0.73)	1.212 (0.55)	-0.07 (.084)	0.47 (0.28)	na	-0.53 (0.48)
4	-0.07 (1.12)	-0.04 (1.03)	1.81 (1.90)	0.56 (1.69)	0.196 (1.03)	-0.24 (.090)	na	2.27 (0.56)
5	-0.63 (1.03)	na	na	-0.058 (0.92)	-0.019 (.084)	0.71 (0.32)	na	-0.45 (0.49)
6	1.01 (1.07)	na	1.17 (0.80)	-0.32 (0.60)	0.29 (0.07)	-0.54 (.276)	na	11.08 (74.6)
7	14.57 (0.92)	na	na	10.70 (0.50)	-0.012 (0.087)	-4.63 (0.356)	na	2.13 (1.12)
8	0.36 (0.58)	na	0.40 (0.65)	-0.57 (0.92)	0.057 (.060)	0.70 (0.27)	-0.58 (0.3)	.047 (0.37)

For the receivers we estimate the following Logit model:

$$\text{Prob}(Y_i = \text{Refinement}) / \text{Prob}(Y_i = \text{Non-Nash}) = b_0 * c_i + b_1 * a_i + b_2 * j_i + b_3 * p_i + b_4 * t_i + b_5 * s_i + b_6 * g_i + b_7 * z1_i + b_8 * z2_i$$

where:

- c = dummy for Caltech subject pool
- a = dummy for Arizona subject pool
- j = dummy for JPL subject pool
- p = dummy for Penn subject pool
- t = period play of game
- s = sequence play of the game
- z1 = dummy for less refined message sent
- z2 = dummy for more refined message sent
- g = dummy for game 8 and 9

Game	P(More)/P(Non-Nash)							
	b_0	b_1	b_2	b_3	b_4	b_5	b_7	b_8
3	0.19 (0.854)	na	1.88 (0.727)	1.21 (0.548)	-0.071 (0.084)	0.47 (0.28)	-0.537 (0.45)	1.15 (0.80)
4	-8.86 (91.5)	-9.50 (80.3)	-6.24 (62.5)	-5.51 (37.1)	-0.25 (0.21)	-1.26 (0.56)	11.96 (61.7)	12.05 (57.5)
5	0.40 (1.79)	na	na	0.94 (1.61)	-0.18 (0.15)	-0.95 (0.56)	5.28 (1.59)	-10.19 (72.37)
6	-2.08 (1.11)	na	1.71 (2.51)	3.33 (2.76)	0.32 (0.48)	0.45 (0.66)	1.63 (1.21)	2.62 (1.34)
7	-4.97 (1.69)	na	na	-4.42 (1.72)	0.22 (0.11)	0.35 (0.39)	2.13 (1.35)	1.62 (1.19)
8	4.46 (1.14)	na	4.88 (1.23)	4.11 (1.38)	-0.19 (0.087)	-0.54 (0.34)	0.006 (1.15)	-2.86 (0.577)

Game	Coefficients	P(Less)/P(Non-Nash)						
1	-3.93 (956.46)	-3.76 (691.31)	-2.49 (859.11)	9.70 (250.209)	-9.08 (67.53)	.08 (240.908)	-11.005 (181.13)	na
2	-3.12 (1.30)	-1.19 (0.79)	-0.49 (0.93)	-0.32 (0.71)	0.067 (0.08)	0.29 (0.23)	-2.23 (0.53)	na
3	-11.20 (486.0)	na	6.01 (62.7)	5.63 (48.7)	-8.99 (61.5)	4.45 (139.7)	-11.61 (257.7)	2.27 (63.3)
4	-9.631 (92.7)	-8.40 (91.1)	-5.62 (82.1)	-4.94 (78.1)	-0.10 (0.14)	-1.528 (0.88)	9.528 (90.88)	16.91 (29.7)
5	2.07 (1.04)	na	na	1.84 (0.95)	-0.003 (0.009)	-0.4 (0.323)	1.26 (1.16)	-2.22 (0.54)
6	1.07 (1.86)	na	1.01 (2.33)	-3.06 (4.56)	0.005 (0.90)	0.501 (1.39)	2.72 (0.96)	0.676 (1.07)
7	-2.07 (1.56)	na	na	-2.01 (1.62)	-0.009 (0.10)	0.201 (0.39)	4.72 (1.266)	0.98 (1.71)
8	4.02 (3.78)	na	2.95 (4.15)	3.73 (5.13)	-0.18 (0.91)	-0.20 (0.45)	1.51 (2.10)	-0.68 (1.11)

Appendix C

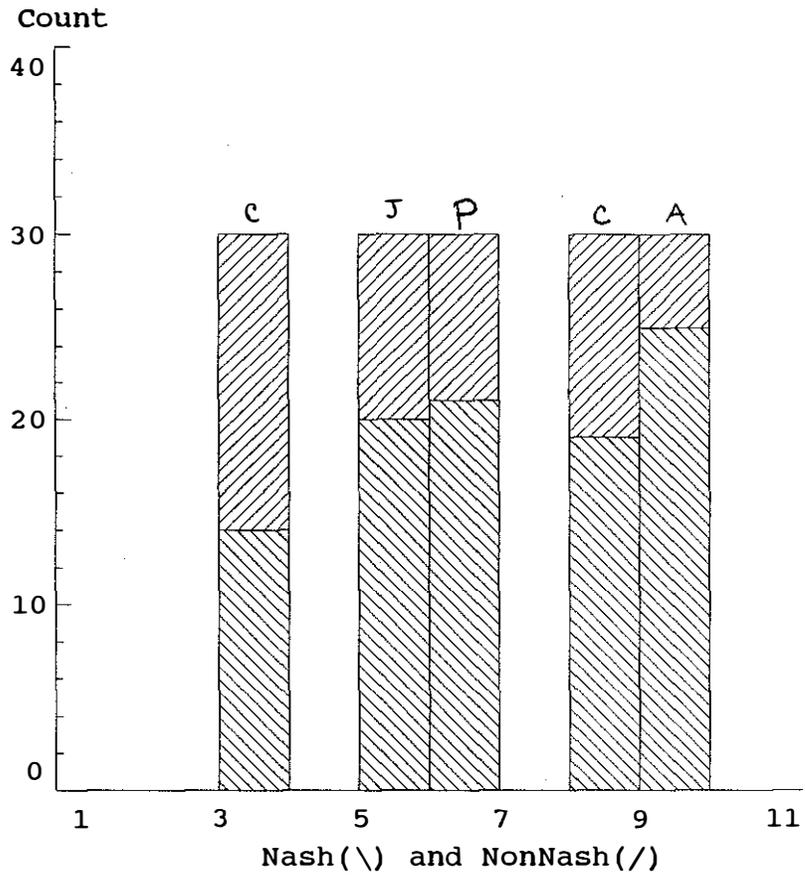
Cohort and Time Series of

Outcomes by Game

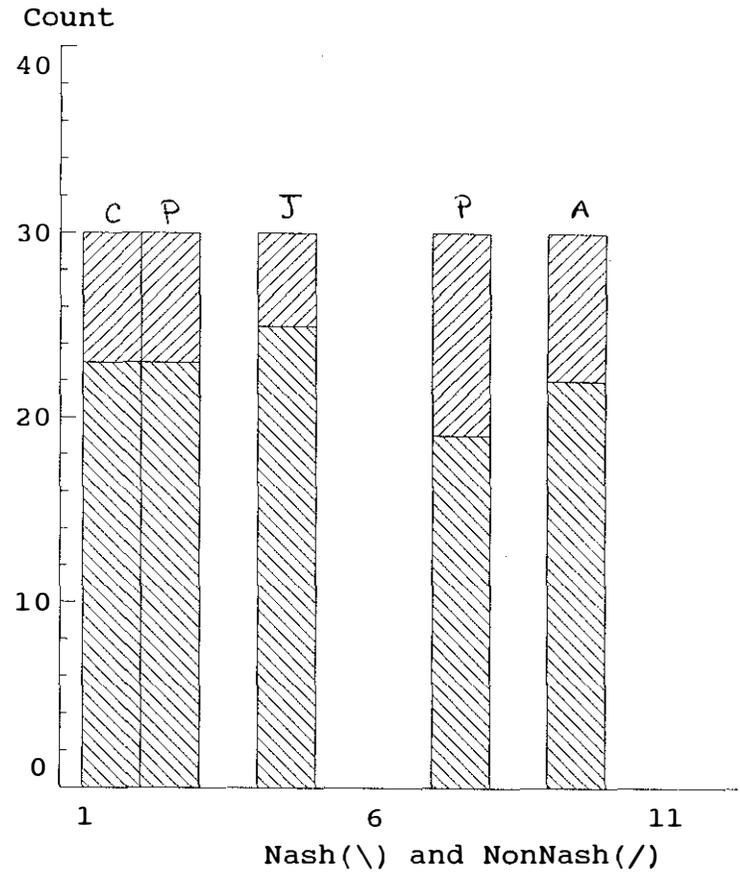
Note: The cohort number at the bottom of each histogram are the experiment numbers used in Table 9 of the text. The letters above each stacked bar is the related subject pool where:

- A ≡ University of Arizona
 - C ≡ Caltech
 - J ≡ JPL
 - P ≡ University of Pennsylvania
-

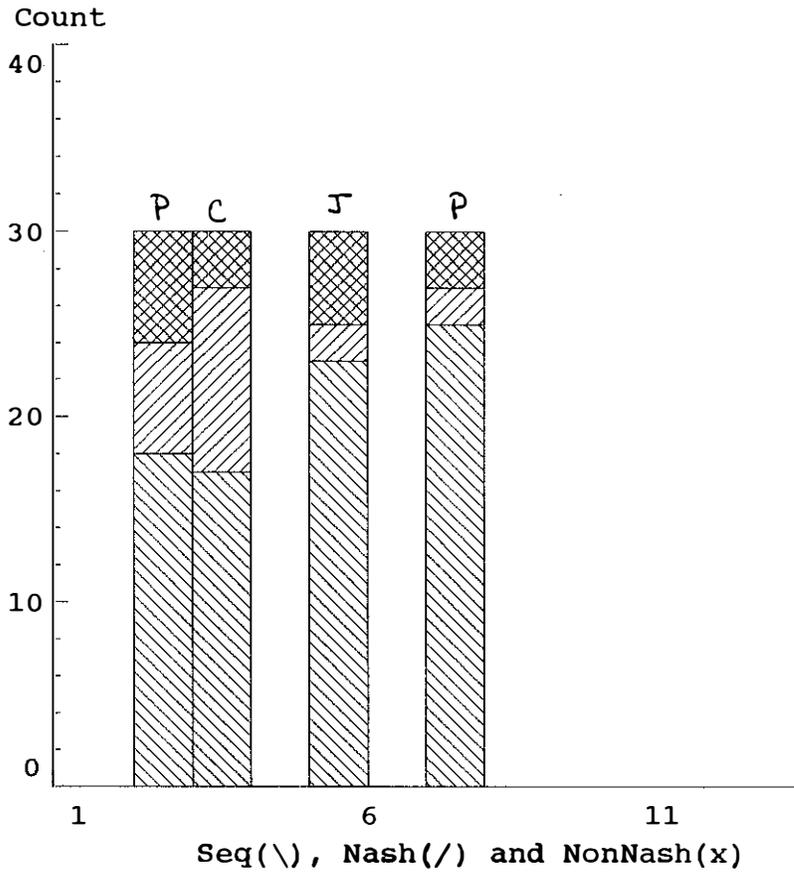
Game1



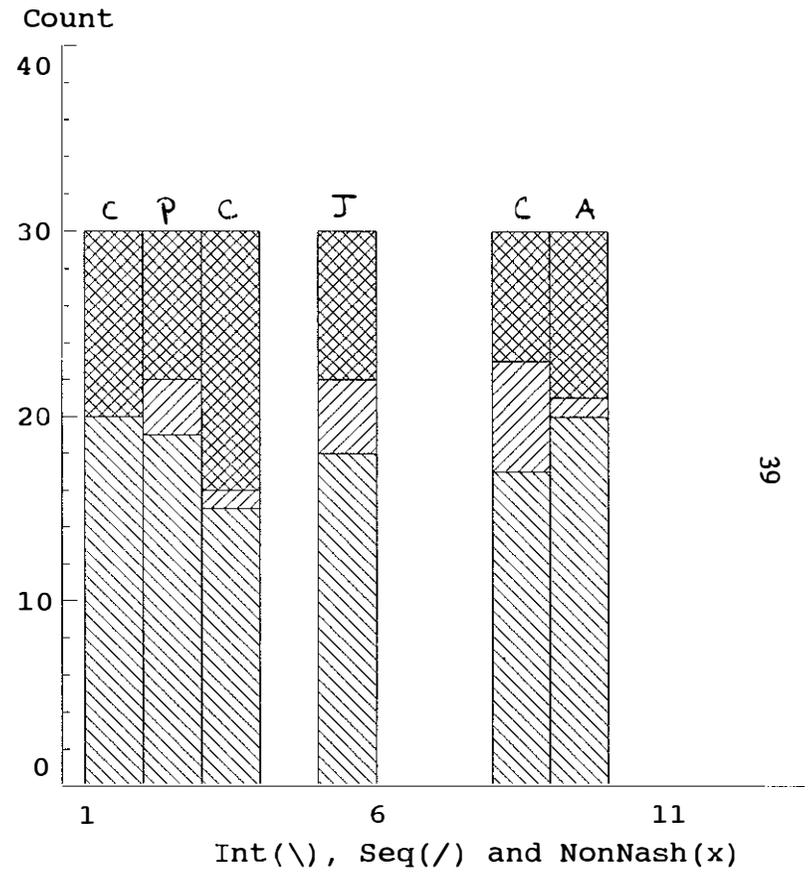
Game2



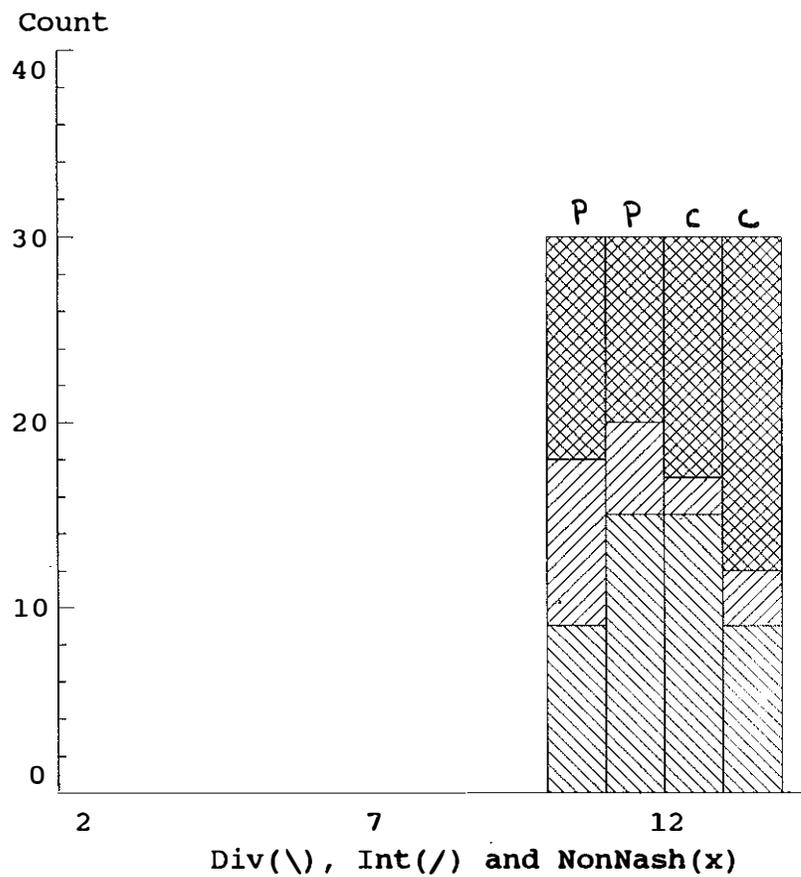
Game3



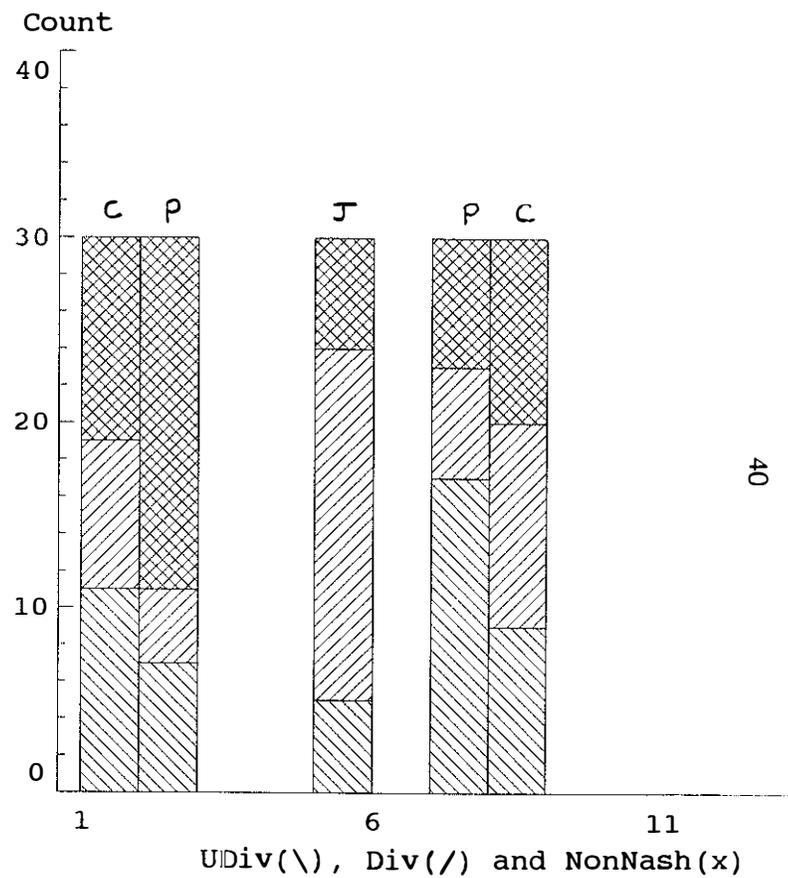
Game4



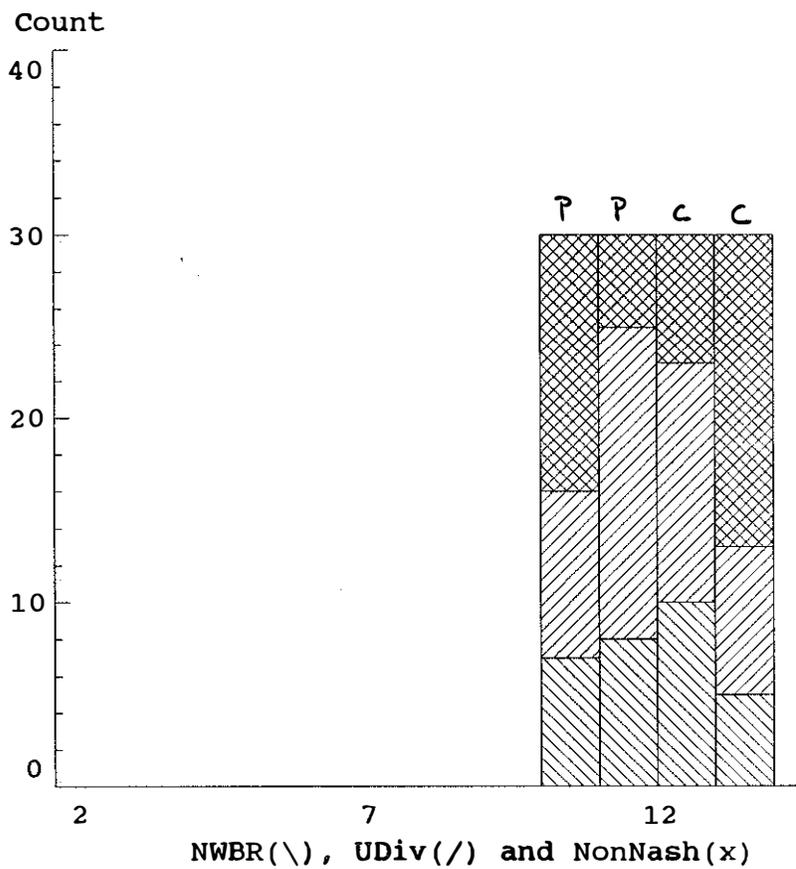
Game5



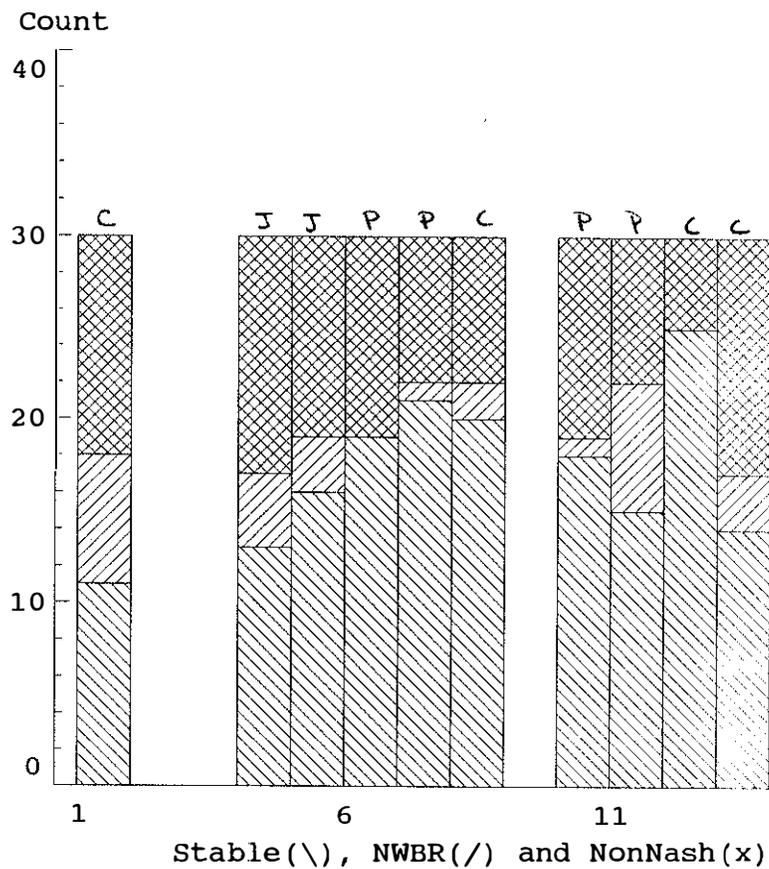
Game6



Game7



Game8&9



REFERENCES

- Banks, Jeffrey S. (1988): "Pricing and Regulation with Verifiable Cost Information," University of Rochester, Department of Economics, working paper.
- Banks, Jeffrey S. and Sobel, Joel (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55, 647-661.
- Berg, Joyce E.; Daley, Lane A.; Dickhaut, John W.; and O'Brien, John R. (1986): "Controlling Preference for Lotteries on Units of Experimental Exchange," *Quarterly Journal of Economics*, 101, 281-306.
- Brandts, Jordi and Holt, Charles A. (1987): "An Experimental Test of Equilibrium Dominance in Signaling Games," University of Virginia, Department of Economics, working paper.
- Cadsby, Charles; Frank, Murrig; and Maksimovic, Vojislav (1989): "Pooling, Separating and Semi-Separating Equilibria in Financial Markets: Some Experimental Evidence," mimeo, University of British Columbia.
- Camerer, Colin F. and Weigelt, Keith (1988): "An Experimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica*, 56, 1-36.
- Cho, In-Koo and Kreps, David M. (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179-221.
- Cooper, Russell; Dejong, Douglas V.; Forsythe, Robert; and Ross, Thomas W. (1990): "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 80, 218-233.
- Cox, James; Smith, Vernon L.; and Walker, James (1985): "Experimental Development of Sealed-Bid Auction Theory: Calibrating Controls for Risk Aversion," *American Economic Review*, 75, 160-165.
- Fudenberg, Drew and Tirole, Jean (1987): *Dynamic Models of Oligopoly*. Chur, Switzerland: Harwood Academic Publisher.
- Grossman, Sanford (1981): "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 461-483.
- Harris, Milton and Raviv, Arthur (1985): "A Sequential Signaling Model of Convertible Debt Call Policy," *Journal of Finance*, 40, 1263-1281.
- Harrison, Glenn and McCabe, Kevin (1988): "Testing Bargaining Theory in Experiments," University of Western Ontario, Department of Economics, working paper, February.
- Harsanyi, John (1967-68): "Games with Incomplete Information Played by 'Bayesian' Players, I, II, III," *Management Science*, 14, 159-182, 320-324, 486-502.
- Harsanyi, John (1975): "The Tracing Procedure: A Bayesian Approach to Defining a Solution for n -Person Noncooperative Games," *International Journal of Game Theory*, 4, 61-94.
- Kahneman, Daniel and Tversky, Amos (1979): "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263-291.
- Kohlberg, Elon and Merten, Jean-Francois (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003-1038.
- Kreps, David and Wilson, Robert (1982a): "Sequential Equilibria," *Econometrica*, 50, 863-894.

- Kreps, David and Wilson, Robert (1982b): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253-279.
- Luce, Duncan and Raiffa, Howard (1957): *Games and Decisions*. New York: John Wiley & Sons.
- Miller, Ross M. and Plott, Charles R. (1985): "Product Quality Signaling in Experimental Markets," *Econometrica*, 53, 837-872.
- Milgrom, Paul and Roberts, John (1982a): "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis," *Econometrica*, 50, 443-459.
- Milgrom, Paul and Roberts, John (1982b): "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, 27, 280-312.
- Milgrom, Paul and Roberts, John (1986): "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 94, 796-821.
- Ochs, Jack and Roth, Alvin E. (1989): "An Experimental Study of Sequential Bargaining," *American Economic Review*, 79, 355-384.
- Pitchik, Carolyn and Andrew Schotter, (1988): "Perfect Equilibria in a Budget-Constrained Sequential Auctions: Experimental Study," *RAND Journal of Economics*, 19, 363-382.
- Quesenberry, C. P. and Hurst, D. C. (1964): "Large Sample Simultaneous Confidence Intervals for Multinomial Proportions," *Technometrics*, 6, 191-195.
- Reinganum, Jennifer and Wilde, Louis (1986): "Settlement, Litigation, and the Allocation of Court Costs," *Rand Journal of Economics*, 17, 557-566.
- Riley, John (1979): "Informational Equilibrium," *Econometrica*, 47, 331-359.
- Roth, Alvin and Malouf, Michael K. (1979): "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, 86, 574-594.
- Rubinstein, Ariel (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97-109.
- Selten, Reinhard (1978): "The Chain-Store Paradox," *Theory and Decision*, 9, 127-159.
- Snee, Ronald d. (1974): "Graphical Display of Two-way Contingency Tables," *American Statistician*, 28, 9-12.
- Spence, A. Michael (1974): *Market Signaling*, Cambridge, Mass.: Harvard University Press.
- Sutton, John (1986): "Non-Cooperative Bargaining Theory: An Introduction," *Review of Economic Studies*, 53, 709-724.
- Wilson, Charles (1977): "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 16, 167-207.