Anarchy in the Laboratory
(and the Role of the State) *

John Duffy
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260 USA
Email: jduffy@pitt.edu

Minseong Kim
School of Economics
Sungkyunkwan University
Seoul 110-745, Republic of Korea
Email: minseong@skku.edu

This Draft: September 2003

Abstract

A recent literature on the economics of conflict has provided conditions under which an "anarchic" outcome may come to serve as an equilibrium for an economy, as well as conditions under which a "dictator" or "government agent" is empowered to make collective action choices that enable the economy to achieve a Pareto superior equilibrium. This paper reports results from a laboratory experiment designed to test the predictions of this theory. We find that in the absence of any government, groups of subjects choose forecasts and actions that lie within a neighborhood of the predicted anarchic equilibrium, where some players choose to be producers, while others choose to be predators. The introduction of the government agent, charged with maximizing the consumption of producers, enables the subject groups to achieve nearly perfect coordination on a Pareto superior Nash equilibrium, where the fraction of time devoted to defense is high, but predation is eliminated.

JEL Classification Nos. D60, D70, D83.
Keywords: Anarchy, Predation, Property Rights, Social Contract, Experimental Economics.

Send page proofs to:

John Duffy
Department of Economics
4S01 Posvar Hall
University of Pittsburgh
Pittsburgh, PA 15260

*Funding for this project was provided by a grant from the University of Pittsburgh Central Research Development Fund. We thank Michelle Garfinkel, Herschel Grossman, Jack Hirshleifer, Suk Jae Noh, seminar participants, and several anonymous referees for helpful comments and suggestions on earlier drafts.

1 Introduction

Economists refer to spontaneous economic order, in the absence of any government, as a state of *anarchy*. In the anarchic state, individuals may engage not only in productive activities but also in appropriative activities such as predation, i.e., appropriation of goods from others by force or threat, and defense against predation. The final distribution of goods in anarchy is determined, in part, by allocations made to appropriative activities, both predatory and defensive, rather than through production and voluntary exchange alone. A recent literature focusing on the economics of conflict (see, e.g., Usher (1987, 1992), Skaperdas (1992), Hirshleifer (1995, 2001), Grossman and Kim (1995, 1996, 2000)), Anderton, Anderton and Carter (1999),and Grossman (2002) among others) has sought to characterize the conditions under which an anarchic state may come to serve as an equilibrium outcome for an economy using game theoretic models.

The anarchic outcome is not the only possibility. Another possible outcome is an economy where every individual is specialized in production, there is no predation, and therefore no need to devote resources to defense – a cooperative, "utopian" world. A third possibility is an "amorphic" outcome – an economy where no effort is made to produce goods that can be easily stolen. A final possibility is an economy where a dictator comes to power or a social contract is formed yielding power to a government; the dictator or government rids the economy of predation, but taxes individuals to pay for this service.

The theoretical literature provides conditions under which each of these possibilities can arise as a unique equilibrium in a strategic setting. However, the theory assumes that all individual actors are rational and that such rational behavior is *common knowledge*. An obvious question, which the theory does not address, is whether the theoretical predictions are robust to environments where individuals do not begin a process of social interaction with common knowledge of the rational beliefs of others, as might have been the case when economic activity first became organized. In such an environment, coordination of individual forecasts could be quite challenging.

In this paper, we report the results of a laboratory experiment designed to determine whether a stable anarchic equilibrium can arise in a laboratory environment where individuals choose to engage either in production and the defense of their resources or in predation, but not both. We also explore whether the introduction of a government (more precisely, a "dictator" whose interests are perfectly aligned with those of the producers) enables the economy to achieve an equilibrium that is superior to the anarchic outcome. The role of the government in our experiment is to require all producers to expend a certain fraction of their endowment on defense.[1] Thus, the introduction of the government agent reduces one dimension of the coordination problem that players face – how much of their resources to devote to defense, leaving only the decision of whether players should engage in production or predation.

Several researchers have conducted laboratory experiments that are related to ours. The most closely related experimental study is by Durham, Hirshleifer and Smith (1998). In their experiment, the choice made by each subject was not between being a producer or a predator. Instead, each subject was a producer–predator who simply chose the allocation of his resources between production and appropriation, without any further distinction between predatory activities and defensive activities. In other words, individual subjects played the role of *both producers and predators*. In addition, they interacted with one another in *pairs*, either in one–shot or in repeated encounters. Each member of a pair decided how much of their endowment to devote to productive effort with the remainder going to appropriative effort. In the "anarchic" Nash equilibrium, each member of a pair devotes a certain fraction of their resource endowment to production with the remainder going to predation; productive effort levels determine total income and appropriative effort levels determine how the total income is distributed between the two members of a pair. Durham et al. examined whether changes in the technology of predation that favored appropriation led to a reduction in productive activities (it did) and, under certain predation technologies, whether the disparity in achieved incomes between members of a pair became smaller over time than the initial disparity in endowments – the "so–called paradox of power" (it did).

Carter and Anderton (2001) conducted an experiment that was similar to Durham et al. (1998). In their experiment, subjects alternated between the roles of "defense" or "offense." Each defense player was randomly matched for a one–shot encounter with each offense player. The two players then engaged in a sequential move game (following Grossman and Kim (1995)), with defense moving first. The defense player chose how much of his endowment to allocate to production and defense. After observing defense's decision, the offense player decided how much of her endowment to allocate to predation or production. The experimental treatments examined whether increases in the relative effectiveness of predation against defense led to changes in the equilibrium amount of predation (it did).

The experiment we conducted is complementary to this earlier work, with several distinctions

[1] All governments, from the very primitive to the most modern, devote some resources to protecting their own goods and that of their citizens from predatory attacks. In our paper, we extend the definition of government to include a loose organization or confederation of producers that makes and enforces a collective decision to defend their own output against predation.

and innovations. First, we are interested in the decisions made by larger *groups* of individuals where cooperation is unlikely relative to the case of *pairs* of individuals. Second, we do not allow individuals to be both predators and producers; they must repeatedly choose which type they want to be, thus allowing us to examine whether individuals follow "career paths" as predators or producers or instead switch between types. Third, we elicit each individual's forecasts of the number of the other players they believe will choose to be predators, and the average fraction of time they believe producers will devote to defense. Since these variables directly determine an individual's payoff from choosing to be a producer or a predator, the forecasts allow us to explore the extent to which each individual's choice of type constitutes a best response to his or her beliefs. We are also able to study whether forecasts and actions are revised over time, in response to the history of play. Finally, and perhaps most importantly, we consider a second treatment where a government is introduced. The role of the government is to impose on producers a certain level of expenditure on defense with the aim of deterring predation and achieving a higher equilibrium payoff for the producers. To the best of our knowledge, this role of the government has not been previously examined experimentally.

While the purpose of our experimental design is to test a model of anarchy and the postulated role of government, the experimental environment we study is similar in some respects to the "market entry" game that has been extensively studied by Amnon Rapoport and associates (see, e.g., Rapoport (1995), Rapoport et al. (1998)) as well as by Camerer and Lovallo (1999). In this game, each player in a group of players simultaneously makes a simple binary decision as to whether she will enter or stay out of a market. If she chooses to enter, her payoff (and that of all other entrants) depends on the total number of subjects who choose to enter. If she stays out, she earns a known, fixed payoff. The various equilibrium predictions for this game are that a certain number of players in the group enter the market while the others stay out. As in our study, Camerer and Lovallo (1999) elicit forecasts from players regarding the number of other players they think will choose to enter the market before entry decisions are made. The main finding from all of these market entry game experiments is that the number of subjects choosing to enter is found to be very close to the predicted number of entrants.

There are two main differences between the market entry game and the game we study in this paper. First, in our environment there are *no certain payoffs*; the payoffs to being a producer or a predator depend on the action choices of all subjects in each round. Second, the decision problem that players face in our "no government" treatment is more difficult than the problem players face in the market entry game. Subjects in our no government treatment must successfully forecast *two* variables: the number of other players they think will choose to be predators *and* the average fraction of time they believe that producers will devote to defense. Thus, one could interpret our experimental findings as revealing how well the Nash equilibrium concept fares when subjects are required to engage in *multi–dimensional* decision–making as opposed to one–dimensional decision–making. When we introduce the government agent, the decision problem of the government agent as well as the decision problem of the other players becomes *one*–dimensional, making this treatment comparable to the market entry game. However, in this "government treatment," the government agent – a single subject – is the one who chooses the fraction of time that all producers must devote to defense, a parameter that affects every player's payoffs. By contrast, in the market entry game experiments referenced above, the analogous parameter choice – the capacity of the market – is made by the experimenter.

The government agent in our environment could be given an alternative interpretation as some kind of organizational authority, e.g., the manager of a firm or the leader of a team, to whom other players ("workers," "team members") cede some of their individual decision-making power. Researchers working in the field of organizational behavior, for instance, have recognized the important role that managers play in coordinating the beliefs of workers (see, e.g., Foss (2001)). However, research in organizational behavior has been slow to adopt game-theoretic models and the testing of such models using controlled laboratory experiments. Our approach and findings are therefore of relevance to this line of research as well. Nevertheless, we choose to interpret our experimental design and findings in the context of the literature on the economics of conflict, as this literature makes use of game-theoretic models and, as noted above, there is some prior experimental research that has made use of this same context.

We report two main findings. First, while the observed aggregate outcomes in the no government treatment are generally in the neighborhood of the predicted anarchic equilibrium, subjects nevertheless have difficulty coordinating on this equilibrium where some players divide their time between production and the defense of their output while others choose to be predators. Those choosing to be producers frequently choose too low a level of defense. A low level of defense encourages an excessive number of subjects to choose to be predators. Over time, the problem of excessive predation is ameliorated, though not entirely eliminated. In one session, we observe a few complete "breakdowns of anarchy" in which the amorphic outcome (where every player chooses to be a predator and nothing is produced) is achieved. We find that subjects do not have "careers" as producers or predators, in that most do not specialize in the role of predator or producer in every round of the game (the theory is silent on this issue). Moreover, the variance in players' choices of the amount of resources to devote to defense remains high, even at the end of the session. Nevertheless, we find that, at the individual level, subjects are generally playing best responses to their individual forecasts, and are revising these forecasts in response to information on actual realizations of the previous rounds.

Our second main finding concerns the impact of introducing a state government, or more precisely, a dictatorship. The dictator, or "the agent" as he is referred to in the neutral language of our experimental design, moves first and chooses the fraction of the endowment that all those choosing to be producers must devote to defense.[2] Given this choice which is made public knowledge, each of the "players," representing the public or the citizenry, decides whether to be a producer or a predator. The government agent's payoff is equivalent to that earned by producers.

In the treatment with a government, we find that subjects achieve the predicted equilibrium much more quickly and more consistently than they do in the treatment without a government. The equilibrium with a government is Pareto superior to the anarchic equilibrium where there is no government agent. Thus, the citizenry is compensated for yielding to the government the power to enforce collective resource allocation decisions – the role of the state. In the equilibrium with a government, all agents choose to be producers, but a sizeable fraction of time is devoted to defense so this equilibrium is not the most efficient outcome. Nevertheless, this finding is important as it provides a rationale for ceding power to the state to make a collective action choice with regard to the level of defense.[3]

The payoff efficient "utopian" outcome, where all agents choose to be producers and devote all of their time to production rather than to defense, is never achieved in either treatment (with or without the government agent) which is consistent with the fact that this outcome is unstable under best response dynamics. Interestingly, in the treatment with a government, we observe government agents naively attempting to steer their economies toward this utopian outcome, albeit without much success.

2 The model

Our experimental design is based on a version of the model studied by Grossman (2002) and Grossman and Kim (2000). We use this model to generate the payoff table used in the experiment. The economy consists of a constant population of $N > 1$ individuals. In every period, each individual chooses whether to be one of two types, a "producer" or a "predator". Let $p \in [0, 1]$ denote the population–wide fraction of predators and let $P = \frac{p}{1-p}$ denote the ratio of predators to producers.

[2]Alternatively, we could have introduced an agent who maximizes the welfare of a subgroup of players or his own welfare. We leave these possibilities for future research.

[3]Sened (1997) provides a much further elaboration of this rationale. His main thesis is that "Government officials commit themselves to protect private property rights in order to induce productivity," (Sened (1997), p. 6), and in doing so, these officials are acting in their own economic or political self–interest.

Producers and predators are both endowed with a single unit of time in every period. Producers must choose the fraction of this time endowment, $d \in [0,1]$ they will devote to defense against predation, with the remaining fraction of time being devoted to production of the single nonstorable consumption good. Let the ratio of defensive to productive activity be given by $D = \frac{d}{1-d}$. Those choosing to be predators are assumed to spend all of their time endowment in predatory activities.

The actual values of p and d determine the aggregate level of output:

$$Y = N(1-p)(1-d). \tag{1}$$

Notice that aggregate output is maximized when $p = d = 0$ and that output is minimized when p and/or $d = 1$. The realizations of p and d also determine the distributive shares of aggregate output that accrue to predators and producers.

The fraction of aggregate output that accrues to all producers is given by:

$$f = \begin{cases} \dfrac{1}{1+\theta P/D} & \text{for} \quad P > 0 \\ 1 & \text{for} \quad P = 0, \end{cases} \tag{2}$$

with the remaining fraction, $1-f$ going to predators. Here $\theta \geq 0$ is a parameter that determines the relative effectiveness of predators in appropriating producers' resources for given P and D. Note that for given θ and P, a higher ratio of resources devoted to defense, D, serves to increase the fraction of aggregate output, f, that producers retain for themselves. However, at the same time, higher values of D reduce the aggregate output of the economy, Y.

Producers share equally the aggregate output retained from predation. Multiplying (1) by (2), and dividing by the number of producers, $N(1-p)$, we find that each producer consumes

$$c^{\text{prod}} = f\frac{Y}{N(1-p)} = \frac{D}{(D+\theta P)(1+D)}. \tag{3}$$

Producers choose a value d, or equivalently, D, so as to maximize (3) taking P as given. The first order condition implies that:

$$D = \sqrt{\theta P}. \tag{4}$$

Predators, too, share equally the aggregate output gained through predation. Therefore, each predator's consumption is given by:

$$c^{\text{pred}} = (1-f)\frac{Y}{Np} = \frac{\theta}{(D+\theta P)(1+D)}. \tag{5}$$

In choosing to be either a producer or a predator, an individual compares c^{prod} with c^{pred} taking as given the ratio of predators to producers, P, and the ratio of defensive to productive activity, D. This yields three possibilities: (a) If $D > \theta$, then $c^{\text{prod}} > c^{\text{pred}}$. In this case, P must equal 0. (b) If $D = \theta$, then $c^{\text{prod}} = c^{\text{pred}}$. In this case, P takes on some value in $[0, 1]$. (c) If $D < \theta$, then $c^{\text{prod}} < c^{\text{pred}}$. In this case, P must equal 1.

It is easily seen that condition (4) is only consistent with case (b), where $D = \theta$; the other two possible outcomes (a and c) violate condition (4). Thus, the unique Nash equilibrium prediction is that a certain number of predators and producers coexist with one another and a positive fraction of the producers' time endowment is devoted to defense.[4] In this equilibrium where $D = \theta$, it follows from (4) that $D = P$, and therefore we also have that $p = d = \frac{\theta}{1+\theta}$.

We will also consider a variant of the above model in which there are $N + 1$ agents. Agent number $N + 1$ is the "government agent", who moves first and chooses the fraction of time, d, that all producers will devote to defense. One can think of this decision as equivalent to the choice of a tax rate that is irreversively imposed on producer's output. The remaining N agents move second and decide whether to be producers or predators, that is, p is determined in reaction to d. The consumption of the government agent is assumed to be the same as that of producers; thus, the government agent internalizes the producers' objective by seeking to maximize c^{prod} as defined above.

The equilibrium arrived at in the model with the government agent, however, is different. Since the government agent moves first and has the sole objective of maximizing c^{prod}, he chooses to set:

$$D = \theta(1 + \epsilon), \tag{6}$$

where $\epsilon > 0$ is the smallest possible fraction by which D can be increased above θ. It follows that if $D > \theta$, $c^{\text{prod}} > c^{\text{pred}}$ and therefore, the Nash equilibrium prediction in this version of the model has $P = 0$. Notice that, since $P = 0$, and $D = \theta(1+\epsilon)$, aggregate output Y should be strictly higher in this equilibrium, although less than the output level that would obtain in the first-best outcome, where $P = D = 0$.

3 Experimental Design

The experimental design consists of two treatments: one in which there is no government agent, the "NG" treatment, and one in which there is a government agent, the "G" treatment. We conducted three sessions of each treatment. For each NG session, we recruited 10 inexperienced subjects, and for each G session we recruited 11 inexperienced subjects. One of the eleven subjects in each of the G sessions was randomly chosen at the outset to serve as the government agent for the entire session.

[4]More precisely, any profile of strategies consistent with $D = \theta$ constitutes a Nash equilibrium.

Subjects were only allowed to participate in one session. Both the NG and G treatment sessions consisted of 30 rounds in which subjects repeatedly faced the same choices, as discussed below.

In the instructions for each treatment (see the appendix for a sample copy), we used neutral language to explain the choices that subjects faced. Predators were referred to as "Type A" players while producers were referred to as "Type B" players. Defensive activity was referred to as "Action 1", while production was referred to as "Action 2." The 10 subjects in the NG and G treatments were referred to as "players" and the eleventh subject who played the role of the government in the G treatment was simply referred to as "the agent." Subjects in both treatments were given the same payoff table which provided them with the expected amount they might earn in dollars (their consumption) from choosing to be either a Type A or Type B player, as a function of (a) the number of the other 9 players (excluding themselves), $\{0, 1, 2, ..., 9\}$ who chose to be Type A players (model value $p(N-1)$) and (b) the fraction of time that all Type B players devote to Action 1, (model parameter d), which was allowed to take on the discrete set of values: $\{0\%, 10\%, 20\%, ..., 100\%\}$. A copy of this payoff table is provided in the appendix – see Table A.1. The two numbers in each cell of the payoff table represent the payoffs to being a Type A or Type B player, respectively, conditional on the values of $p(N-1)$ and d. In calculating the values in the payoff table we set the model parameter $\theta = 3.4$, and calculated the values of c^{prod} and c^{pred}, over the discrete grid of values for p and d. These consumption amounts were multiplied by a scale factor of 6 to provide more reasonable monetary incentives.

Given our choice of $\theta = 3.4$, and our discretization of the action space, the Nash equilibrium in the NG treatment has $p = d = .70$. We purposely chose a Nash equilibrium that was "off–center," i.e., different from $p = d = .50$, and which involved a lot of predation (7 out of 10 players are predicted to choose to be Type A players); the latter choice provides us with a means of contrasting our findings in the NG treatment with our findings in the G treatment, where the equilibrium prediction calls for no predation. Indeed, in line with the theory presented in section 2, the Nash equilibrium in the G treatment involves $d = .80$ and $p = 0$.[5]

The payoffs in the Nash equilibrium of the NG treatment are such that each of the 7 Type A players earns \$0.60 per round and each of the 3 Type B players earns \$0.41 per round. We chose to have asymmetric rather than equal payoffs in the anarchic Nash equilibrium to make it more challenging to figure out this equilibrium. To confirm that these are the equilibrium payoffs, we use the following reasoning: Each of the 7 Type A players forecast that 6 out of the other 9 players

[5]Since the fraction of time devoted to defense was presented in increments of 10% and the anarchic equilibrium has $d = .70$, the equilibrium with a government agent has $d = .80$.

(excluding themselves) will choose to be Type A players and that $d = 70\%$. Given these forecasts, the payoff table in the appendix reveals that choosing to be a Type A player yields a payoff of \$0.60 while choosing to be a Type B player yields a payoff of \$0.57. Thus, these players prefer to be Type A rather than Type B players. Each of the 3 Type B players forecast that 7 out of the other 9 players will choose to be Type A players and that $d = 70\%$. Given these forecasts, the payoff table reveals that choosing to be a Type A player yields a payoff of \$0.38 while choosing to be a Type B player yields a payoff of \$0.41. Thus, these players prefer to be Type B rather than Type A players. Moreover, given the forecast that 7 out of the other 9 players choose to be Type A players, these players reason that all of those choosing to be Type B players will not deviate from the choice of $d = 70\%$.

In the Nash equilibrium of the G treatment, the government agent announces that d=80% and then all players choose to be producers each earning \$1.20 per round. The government agent, who internalizes the producer's objective (is always a Type B player) also earns \$1.20 per round.[6] Notice that payoffs and output are substantially higher in the G treatment as compared with the NG treatment.

At the beginning of each session, the instructions were read aloud and any questions were answered. The experiment was conducted in the Pittsburgh Experimental Economics Laboratory (PEEL). Subjects entered their choices and received feedback using networked personal computers.

The sequence of events in every round of each "NG session" was as follows. At the beginning of each round, subjects were asked to provide private, individual *forecasts* of the number of the other 9 players who would choose to be Type A players (predators) in that round. They were also asked to provide private, individual forecasts of the *average* fraction of time that all Type B players (producers) would devote to Action 1 (defense), i.e., to forecast the average value of d. They were instructed that, *in the event they choose to be Type B players*, their forecast of the average fraction of time that all Type B players devote to Action 1 will also serve as their individual contribution (as a Type B player) to the determination of this average fraction.[7] They were further instructed that if they did not choose to be Type B players, that is, if they chose to be Type A, their forecast

[6]The payoff table for the government agent is the same as the payoff table for the other players (Table A.1) with a single exception. The first column of the agent's payoff table indicates the number of the other *10* (not 9) players choosing to be Type A. Since the "agent" is always a Type B player, this relabeling does not change any of the payoff numbers found in the payoff table used by the other 10 players. However, for the government agent, an additional row must be added to the bottom of this payoff table corresponding to the case where 10 out of 10 players choose to be Type A players. The payoffs in this row consist entirely of zeros, as nothing is produced in this case.

[7]In equilibrium, all players choosing to be Type B players (producers) would choose the same amount of time to devote to Action 1 (defense), so an individual's choice would be the same as the average choice. Of course, out of equilibrium this need not be the case. We chose to have subjects' forecast of the average amount of time devoted to defense by all producers also count as their *own* contribution to defense so as to simplify the problem subjects faced, and to get subjects to understand that their contribution affected the average contribution.

of the average fraction of time that all Type B players would devote to Action 1 *would not count* in the determination of the average fraction.[8] The actual average fraction of time that all Type B players devoted to Action 1 was used to determine the payoffs that all Type A and Type B players received in the round. For example, as the instructions noted, if four out of ten players chose to be Type B players and these four players' forecasts of the average fraction of time that all Type B players devote to Action 1 were 60%, 20%, 70% and 10%, then the actual *average* fraction of time devoted to Action 1 by all Type B players is 40%, and this is the fraction that was used to determine all players' payoffs for that round.[9] Forecasts other than those made by Type B players about the average value of d had no direct payoff consequences; in particular there was no monetary reward for forecast accuracy. However, all subjects had an incentive to accurately forecast the number of the other 9 players choosing to be Type A players, and Type A players (along with Type B players) had an incentive to accurately forecast the average value of d, as more accurate forecasts provided players with a more accurate assessment of the expected value to being a Type A or Type B player.

After subjects had entered both of their forecasts they were each asked what type of player they wanted to be in that round, Type A or Type B. They were informed that they could choose to be either type of player in every round. In making their forecasts and choice of type, subjects were not allowed to communicate with one another nor could subjects see any other subjects' forecasts or action choices. After all subjects had chosen a type, the computer program calculated the actual number of Type A players and the average fraction of time that Type B players would devote to Action 1. Given these two numbers, the program calculated the payoffs to being a Type A or Type B player from the payoff table. Subjects were informed on their computer screens of their payoff from choosing to be either a Type A or a Type B. They were also reminded of their own, individual forecasts and informed of the actual number of the other 9 players choosing to be Type A players and of the actual average fraction of time that Type B players devoted to action 1. They were asked to record all of this information on private record sheets at the end of each round. If the 30th round had not been reached, another round of this same game was begun.

The sequence of events in every round of the "G sessions" was somewhat different. Each round began with the single "agent" choosing the fraction of time that all those choosing to be Type B players would devote to Action 1. The agent's choice was announced aloud and appeared on subjects' computer screens. It was public knowledge that the agent's payoff for the round was equal to the

[8] Recall that, in the theory, only producers decide between defensive and productive activity; predators devote all of their time to predation.

[9] If the average fraction of time devoted to Action 1 did not equal one of the 11 possible values, it was rounded to the nearest of these possible values, and subjects were made aware of this rule.

payoff earned by Type B players. Furthermore, all subjects understood that the agent's choice would partially determine their payoff for the round and that there was absolutely no credibility problem, as the agent's choice, once announced, was irreversible.[10] Following the agent's announcement, each of the other 10 subjects – the "players" – were asked to forecast the number of the other 9 players who would choose to be Type A players in that round. They were then asked to choose whether they wanted to be a Type A or a Type B player. No communication among the players was allowed. Once all players' choices had been made, the computer program determined the actual number of Type A players. Using the agent's choice for the fraction of time that Type B players were required to devote to Action 1, and the number of Type A players, the program then calculated the payoffs to being a Type A or Type B player using the given payoff table. Subjects were informed on their computer screens of their forecast of the number of Type As, the actual number of Type As, the agent's choice for the fraction of time that Type B players would devote to Action 1, their own choice of Type (A or B), and their payoff for the round. Subjects were asked to record this information on record sheets at the end of each round. If the 30th round had not yet been reached, another round of this same game was begun.

At the outset of each NG and G session, subjects were shown how they could use their forecasts and the payoff table to determine their expected earnings from choosing to be a Type A or Type B player. They were given two practice tests to insure that they understood how to read the table and could calculate their expected, as well as their *ex post* actual earnings in each round. In the NG treatment, they were also quizzed to insure that they understood that their forecast of the average fraction of time that all Type B players devoted to Action 1 also constituted their contribution to this amount if they chose to be Type B players. In the G treatment, subjects were quizzed regarding the expected and actual payoffs received by both the "agent" and the 10 other players. These quizzes were included as part of the instructions. No player was allowed to play the game until they had answered all quiz questions and any mistakes they made had been corrected. See the appendix for details.

Since the game was played repeatedly by the same group of subjects for 30 rounds, there is the possibility that subjects employed dynamic, repeated game strategies that might give rise to a different equilibrium configuration than the one given by the static model described in section 2. We have several reasons to believe that such behavior is a remote possibility. First, the model that was used to generate the payoff table is a static model; there is no inherent dynamic element

[10]See Van Huyck et al. (1995) for an alternative approach, where the ability of a 'dictator' to precommit to a policy is an experimental treatment variable.

that allows the strategies of one round to affect payoffs in subsequent rounds. Second, the players move simultaneously without observing the individual choices of other players, which reduces the possibility of cooperation. See Durham, Hirshleifer and Smith (1998) for a discussion of this issue. Third, cooperation among 10 players is highly unlikely. See, e.g., the work of Van Huyck et al. (1990) on group coordination games. Finally, subjects were informed that the game had a finite end – just 30 rounds, thus eliminating the possibility of dynamic repeated game strategies that depend on an indefinite horizon.

Subjects were University of Pittsburgh undergraduates. They were recruited through newspaper advertisements and classroom visits. Each subject was paid the sum of their individual earnings from all 30 rounds plus a $5 show–up payment. Each experimental session lasted about 75 minutes. The average total earnings of each subject in the 3 NG sessions was $28.60 and average total earnings in the 3 G sessions was $34.95.

4 Hypotheses

Our parameterization and implementation of the model developed in section 2 yields the following two main hypotheses.

Hypothesis 1 The distribution of types and forecasts will, over time, correspond more closely to the non–cooperative anarchic Nash equilibrium, where $p = d = .70$. This hypothesis follows from the play of "best–response" dynamics by all N players over time as information on p and d is revealed at the end of each round.

Hypothesis 2 The introduction of the government agent, whose aim is to maximize the consumption realized by producers, will move the economy to a Nash equilibrium that is characterized by an absence of predation, $p = 0$, and that is superior in payoff terms to the anarchic equilibrium, but still dominated by the utopian outcome, in that $d = 0.8$.

Note that both hypotheses are with respect to the achievement of certain equilibrium values for p and d, which are fractions. There is no prediction as to whether individuals play mixed or pure strategies. In addition, because all individuals are ex ante identical, there is no prediction as to which individuals will choose to be producers or predators.

5 Results

Our discussion of the results from the experiment are divided up between aggregate experimental findings and an analysis of individual behavior.

5.1 Aggregate findings

[Insert Table 1 here.]

The aggregate results are summarized in Table 1, which reports the means and standard deviations of the values of d and pN for all three sessions of each treatment. These means and standard deviations were calculated using the actual realizations of d and pN over all 30 rounds as well as for the last 10 rounds of each session. The results shown in Table 1 are generally supportive of our two hypotheses. In particular, the aggregate statistics on d and pN from each NG session suggest that there is some tendency for the realizations of d and p to converge to a neighborhood of the Nash equilibrium values, $d = p = .70$. With the exception of session NG2, the mean values of d and pN are closer to the Nash equilibrium values in the last 10 rounds than in all 30 rounds. Again, with the exception of session NG2, the standard deviations of the realizations of d and pN are also smaller in the last 10 rounds than in all 30 rounds. Looking at the means and standard deviations for just the last 10 rounds we see that in the 3 NG treatments, the average choice of d by Type B players had a mean of 0.5567, which is somewhat below the Nash equilibrium prediction of 0.7. However, this low mean value was mainly due to a single session, NG2, where the average choice of d by Type B players was just 0.29 in the last 10 rounds. We will examine this particular session in more detail below. As the choice of d was somewhat below the Nash equilibrium prediction, it is not surprising to find that the number of players choosing to be Type A over the last 10 rounds of the 3 NG session was slightly above the Nash equilibrium prediction of $pN = 7.00$.

By comparison with the three NG sessions, the mean values of d and pN were quite close to the Nash equilibrium predictions and more consistent across the three G sessions. The mean choice of d by the government agents in the last 10 rounds of all three G sessions was 0.790 with a standard deviation of only 0.0305. The Nash equilibrium prediction for d is 0.800. The mean value of pN in the last 10 rounds of all three G sessions was 0.57 with the standard deviation of 1.55. The Nash equilibrium prediction for pN is 0. The standard deviation for pN is somewhat high relative to the standard deviation for d because the best response dynamics call for a large deviation in pN for small deviations in d.

While informative, the standard deviations of d and pN provide imperfect measures of convergence to a particular Nash equilibrium, as they are calculated with respect to the mean realizations of these variables. It is therefore instructive to consider the mean squared deviation (MSD) of the realizations of d and pN from the predicted Nash equilibrium values. In Table 2 we report the MSD of d and pN from the predicted Nash equilibrium values over the first, second and last (third) 10 rounds of each session. Again, we use $p = d = .70$ as our benchmark for the NG sessions and $p = 0$, $d = .80$ as our benchmark for the G sessions.

[Insert Table 2 here.]

Table 2 reveals that, with the exception of session NG2, the MSD from the Nash equilibrium predictions decrease from the first 10 rounds of each session to the last 10 rounds of each session, indicating convergence toward equilibrium is occuring in these sessions.

The results in Tables 1 and 2 suggest that convergence to a neighborhood of the Nash equilibrium was obtained in most sessions of both treatments. However, given the different values for pN and d in the Nash equilibrium of each treatment, it remains difficult to gauge how 'close' subjects were to the predicted equilibrium behavior across the two treatments. To address this issue, we calculated the sum of the payoffs earned by all 10 players in the NG treatment and by all 11 players in the G treatment at every round, and divided this sum by the payoffs that 10 or 11 players would have earned had they played according to the Nash equilibrium predictions. If players are playing according to the Nash equilibrium predictions, then this 'efficiency ratio' will equal unity; values greater than 1 indicate that players are earning more than in the Nash equilibrium and values less than 1 indicate the opposite. The mean and standard deviations of these efficiency ratios for the last 10 rounds of each session are reported in Table 3.[11] The table reveals that the means for the G sessions are generally closer to unity than the means for the NG sessions. The more striking finding is that the standard deviations in the efficiency ratios over the last 10 rounds of each session are smaller in the G sessions than in the NG sessions. Indeed, a nonparametric robust rank order test confirms that these standard deviations are significantly lower in the G sessions as compared with the NG sessions (p-value =.10, lowest possible for 3 observations per treatment). We conclude that, the one–dimensional choice problem (the G treatment) leads more quickly to behavior that is consistent with the Nash equilibrium outcome relative to the multi-dimensional choice problem (our NG treatment).

[Insert Table 3 here.]

Figures 1.1-1.3 show the round-by-round realizations of d and pN for the three NG sessions and Figures 2.1-2.3 provide the same information for the three G sessions. Each figure consists of two panels. The top panel, "a," in each figure reports the actual value of d – the average fraction of

[11] Note that the efficiency ratio we have calculated is not conditional on the choice of d, which is taken to be the Nash equilibrium value. In the G sessions it might be more reasonable to allow such conditioning, in calculating the efficiency ratio as the government agent announces the choice of d before subjects choose their type. For purposes of comparison with the NG sessions, however, we chose to report only the unconditional efficiency ratio.

time devoted to Action 1 (defense) by all Type B (producer) players in each and every round in the NG sessions (Figures 1.1-1.3), or the agent's choice of the fraction of time that Type B players will devote to Action 1 in every round in the G sessions (Figures 2.1-2.3). The bottom panel, "b," in each figure reports the actual value of pN – the number of the players who chose to be Type A (predator) players in every round. The Nash equilibrium predictions are indicated by horizontal dashed lines. The figures confirm the findings reported in Tables 1–3 and provide additional insights.

[Insert Figures 1.1-1.3 here.]

As noted earlier in the discussion of Table 1, in 2 of the 3 NG sessions, there is some tendency for the actual values of d and pN to converge to a neighborhood of the Nash equilibrium values, $d = 70\%$ and $pN = 7$. The sole exception is session NG2. In session NG2, the actual value of d never rose above 60% and was equal to 0 on 4 occasions. The number of Type A players, pN, in session NG2 typically exceeded the Nash equilibrium prediction (the average value of pN in session NG2, as reported in Table 1, is 8.03), and reached 10 on two occasions, when the actual value of d was recorded as 0.

The aggregate results found in session NG2 are of particular interest because they demonstrate the possibility of a "breakdown of anarchy" in favor of what Hirshleifer (1995) has termed "amorphy." The amorphic outcome in our model obtains when every player chooses to be Type A and nothing is produced. In session NG2, we see that the amorphic outcome obtains in rounds 8 and 30.[12] While one amorphic outcome was also observed in session NG3, this outcome occurred in the first round of the session and never occurred again. We will examine the circumstances that may have given rise to the amorphic outcomes in session NG2 later in section 5.2.

[Insert Figures 2.1-2.3 here.]

Looking next at the three G sessions, Figures 2.1–2.3, we see that the subject designated as "the agent" initially experimented with the choice of d, the fraction of time that Type B players were required to devote to defense. Approximately one-third of the way through each G session, (on average, by round 10) each player designated as "the agent" figured out that predation could be largely deterred by a choice of $d = .80$, the Nash equilibrium choice. Indeed, as can be seen in the b panels of Figure 2, the choice of $d = .80$ was met with a precipitous decline in the number of players choosing to be Type A, as this value for d (as well as higher values) implies that choosing to

[12]The amorphic outcome is not a Nash equilibrium because, given that the other 9 players choose to be Type A, choosing to be Type A is payoff–dominated by choosing to be Type B (See the payoff table, Table A.1, in the appendix.); thus, the fact that the amorphic outcome obtained in round 30 cannot be due to a "last–round" effect.

be a Type B player strictly dominates choosing to be Type A player. (See the payoff table, Table A.1 in the appendix.)

Interestingly, once the government agent had chosen to set $d = .80$ for the first time, he was not content to stick with this choice for the remainder of the session. In session G1, the agent departed from the choice of $d = .80$ (which he had made in the 10 previous rounds) and experimented with a choice of $d = .70$ in round 23. Similarly, the agent in session G2 departed from his choice of $d = .80$ in favor of $d = .70$ in two rounds, 10 and 14. The agent in session G3 experimented the most, departing from a choice of $d = .80$ five times in favor of $d = .70$ in rounds 5, rounds 9, 18, 26 and 29. Interviews with these three players at the end of each G session confirmed that each was testing the stability of the Nash equilibrium. In lowering their choice from $d = .80$ to $d = .70$ each of these subjects thought that the other 10 players might all continue to choose to be Type B players. These agents reasoned that, if the other 10 players continued to be Type B players, and d was only gradually lowered, everyone, including the agent (who earned the same payoff as a Type B player), would receive a higher payoff. Eventually, the "utopian" outcome, where every player continued to choose to be Type B and $d = 0$ might be achieved. Instead, consistent with best reply dynamics, the lowering of the choice of d from .80 to .70 resulted in a sharp rise in the number of players choosing to be Type A players. When the agent lowered d from .80 to .70 the number choosing to be Type A players ranged from 4 to 8, which consists of a neighborhood around the Nash equilibrium prediction of 7.[13]

Summarizing all of these findings, we conclude that there is support for both our Hypotheses 1 and 2. However, it appears that convergence to the Nash equilibrium was somewhat easier in the G treatment than in the NG treatment. One explanation for this result is that the G treatment, (like the market entry games studied by Rapoport et al.) involves one-dimensional decision making on the part of each subject, while the NG treatment requires two-dimensional decision making: all players have to make and act upon forecasts of both p and d. As the complexity of the decision-space increases, the convergence to the Nash equilibrium becomes more difficult and, as we have found, the aggregate results become less consistent across sessions. Nevertheless, our experiment also demonstrates that individuals are frequently able to achieve something close to the Nash equilibrium even when they must engage in multi-dimensional decision making. This finding is important, in that it suggests that results obtained in simpler, one-dimensional coordination games might extend as well, albeit with some greater noise, to multi–dimensional coordination games.

[13] Indeed, with a single exception, when d was lowered from .80 to .70 the number of subjects choosing to be Type A players was in the range 7 ± 2. This is finding is consistent with behavior in market entry game experiments where market capacity is varied – see, e.g., the work of Rapoport and associates and others referenced in section 1.

5.2 Individual behavior

5.2.1 Specialization

We defined anarchy as a state in which some players choose to engage in production and the defense of their output while other players choose to engage in predation, appropriating some of the producers' output. Since all players are assumed to be identical prior to making their choices, the Nash equilibrium concept used to characterized the anarchic outcome only requires that, in the aggregate, the fractions p and d are consistent with the Nash equilibrium prediction. In particular, the Nash equilibrium concept is silent on the issue of which individual players specialize in production or predation in a one-shot game or whether players switch with some frequency between these two roles in a repeated game.

To address the issue of whether players specialized or not, we report in Figures 3.1-3.3, the frequency with which each player chose to be a Type A player (predator) together with that player's average forecast of the value of d over the last 10 rounds of each NG session. We chose to look at choices in the last 10 rounds, rather than the full 30 rounds because Figures 1.1–1.3 suggest that there is a period of adjustment in the early rounds of each NG session, and because we wanted to allow enough time for individual player actions to possibly settle down to some stable pattern of behavior. The actual frequency with which each player (numbers 1–10) chose to be Type A is represented by a black bar. Below each black bar is a grey bar revealing this same player's average forecast of the value of d – the average fraction of time that all Type B players devote to action 1. Recall that if a player chooses to be a Type B (producer) rather than a Type A player, their forecast of d serves as their contribution to the calculation of the average value of d that is used to determine every player's payoff. Hence, it is important to examine the choices made for d among those who were not always choosing to be Type A players.

[Insert Figures 3.1–3.3 here.]

We see from Figure 3, that in the last 10 rounds of the three NG sessions, there was some, but less than complete specialization in the two roles. In session NG1, (Figure 3.1) we see that five players always chose to be Type A, while one player, number 4, always chose to be Type B. The remaining four players switched with different frequencies, all less than or equal to 0.7, between the two roles. In the aggregate, the mean value of pN was 7.10 for the last 10 rounds of session NG1. Of the five subjects who sometimes or always played the role of Type B, four chose values for d that were on average less than the Nash equilibrium predicted value of 0.7. The sole exception was the one player (number 4) who always played the role of Type B; this player's choice of d averaged 0.9, and since this player was always a Type B player, his choice brought the actual value of d up enough so that it averaged 0.70 over the last 10 rounds of session NG1 (as can be seen in Table 1).

In session NG2 (Figure 3.2), we see that over the last 10 rounds, *no player* was ever specialized in the role of Type B (producer). We also see that eight players chose to be Type A players more than 70% of the time, and that three of these players always chose to be Type A players. The choices for d made by the seven players who sometimes chose to play the role of Type B, were all well below the Nash equilibrium prediction of 0.7; the highest individual average for d was just 0.41.[14]

The individual decisions over the last 10 rounds of session NG3 (Figure 3.3) are similar to those found in session NG1. In session NG3, three players always chose to be Type A, while one player, (number 9) always chose to play the role of Type B. Of the seven subjects who sometimes or always played the role of Type B, three chose values of d that averaged more than 0.7, while the other four chose values of d that averaged between 0.25 and 0.65.

These individual results from the three NG sessions suggest that the achievement of the stable "anarchic" Nash equilibrium can be quite difficult under certain circumstances. These circumstances would appear to be 1) the absence of a core of one or more individuals specialized in the role of a Type B player or choosing to be a Type B player with high frequency, and/or 2) an allocation of time to defense, d, that is too low to deter excessive predation. Indeed, the low realized value of d in session NG2 is reflected in subjects' forecasts, which are persistently low as well.

To summarize, Figures 3.1-3.3 reveal that in the last 10 rounds of the NG treatment players were not completely specialized in the role of Type A or Type B players. Instead, subjects were frequently switching between roles. Moreover, there was a lack of consistency in subjects' forecasts of the value of d that would prevail, and hence in their contributions to the determination of this value. While such behavior need not be inconsistent with the Nash equilibrium concept, it nevertheless seems at odds with the picture that is often painted of stable, anarchic societies, where some agents are specialized in production/defense while others are skilled predators. Our experimental findings suggest that the latter scenario is unlikely to obtain in the environment studied without some further augmentation of the model, for example by supposing that some players have a comparative advantage in production/defense while others have a comparative advantage in predation, e.g., through differences in the technology used to distribute output. Such an extension might also reduce the variance in the choices of p and d. Our experimental findings thus suggest one dimension

[14] We further note that, the cost of choosing too low a value of d by those choosing to be producers was much smaller than the cost of not choosing to be a predator, so payoff salience may also have been a factor in the lower-than-equilibrium-level of defense observed in session NG2, though this does not seem to have been a factor in the other two sessions.

along which the theory of anarchy could be further developed.[15] Furthermore, the results of session NG2 suggest that anarchy may be prone to "breakdowns" in the direction of what Hirshleifer (1995) characterized as "amorphy" (absence of production and the defense of resources) despite the stability of the anarchic equilibrium under strict best reply dynamics. The circumstances under which anarchy breaks down appear to be the absence of a core of individuals who frequently (if not always) choose to produce, and/or a common awareness of the importance of an adequate level of defense.

We could also perform a similar individual level analysis of the decisions made in the three G sessions, however, from Figure 2, it is clear that by the last 10 rounds of this treatment nearly all subjects chose to be Type B players in line with our Hypothesis 2. Subjects in the G treatment nearly always chose to be Type B players because the subject designated as "the agent" nearly always chose to set $d = .80$, the predicted Nash equilibrium value, with only occasional experimentation to a the lower value of $d = .70$ as noted earlier. Hence, there seems to be no need to further examine individual decisions in the three G sessions, in the same manner as we have for the NG sessions.

Nevertheless, it is of interest to examine individual actions conditional on *forecasts* in the G sessions as well as in the NG sessions to determine whether players were playing best responses to their forecasts. We now turn our attention to this task.

5.2.2 Best responses and learning behavior

[Insert Table 4 here.]

Table 4 reports the frequency with which each subject in the three NG sessions played a best response (in choosing to be either Type A or Type B) to his/her forecasts of the values of pN and d. The table also reports the frequency with which each of the 10 "players" in the three G sessions played a best response to the agent's choice of d and their own forecast of pN. These best response frequencies are reported over the entire 30 rounds and over the last 10 rounds of each NG session. We see that, on average, subjects played best responses in the majority of all rounds and that the frequency of best response play tended to be higher in the last 10 rounds relative to the entire 30 rounds of the session; the results for session NG3, however, appear to be an exception to this rule. Looking at the overall average frequencies of best response play, it appears that these frequencies were somewhat higher in the three G sessions as compared with the 3 NG sessions. This finding should not be too surprising since midway through each G session, the "agent" was nearly always choosing to set $d = .80$, and the other 10 subjects were nearly always forecasting (correctly) that

[15] Note that this finding could not have arisen in the experimental study conducted by Durham et al. (1998), where each player was required to play the role of both producer and predator.

the number of Type A players would be approximately 0.

According to Table 4, the anomalous amorphic outcomes observed in session NG2 were not due to an overall lack of best response behavior; indeed, on average, subjects in session NG2 tended to play best responses with a somewhat higher frequency on the whole than in sessions NG1 or NG3. If the actual realization of d is consistently below 0.7, and forecasts are consistent with such realizations, then it is always a best response to choose to be a Type A player, which was the choice made by most subjects in every round of session NG2. Hence the higher frequency of best response play in session NG2. However, a disaggregation of best response frequencies *by player type* in the NG sessions reveals that the frequency of best response play by subjects choosing to be Type B players was particularly low in session NG2. This low frequency of best response play by Type Bs in session NG2 relative to sessions NG1 and NG3, together with the low forecast of d would appear to be the source of the anomalous outcomes in session NG2.

Given that individuals generally played best responses to their forecasts of pN and d, it is of interest to examine how forecasts changed over time in response to the *ex post* feedback subjects received concerning the actual realizations of p and d. As a first approximation to modeling subjects' forecast revision process in the NG sessions, we made use of the following (naive) adaptive learning model, based on Selten and Stoecker's (1986) *learning direction theory*. Let $\hat{p}(t)$ denote each agent's forecast of p at the beginning of round t, and let $p(t)$ denote the ex post realization of this value for round t. Similarly, let $\hat{d}(t)$ and $d(t)$ denote the forecast of d and ex post realization of d in round t. Let us begin by making the (extreme) assumption that forecasts of p and d are revised independently of one another. A more sophisticated learning model (explored below) would allow the forecasts, $\hat{p}$, to be conditional on realizations of d, and the forecasts, $\hat{d}$, to be conditional on realizations of p. However, we begin with a simpler updating process to examine the extent to which this updating process can explain the data. The simple forecast revision rule further assumes that:

$$\hat{p}(t) \quad \begin{array}{lll} > & \hat{p}(t-1) & \text{if } \hat{p}(t-1) < p(t-1), \\ = & \hat{p}(t-1) & \text{if } \hat{p}(t-1) = p(t-1), \\ < & \hat{p}(t-1) & \text{if } \hat{p}(t-1) > p(t-1). \end{array} \tag{7}$$

Similarly, in the NG sessions, we assume that:

$$\hat{d}(t) \quad \begin{array}{lll} > & \hat{d}(t-1) & \text{if } \hat{d}(t-1) < d(t-1), \\ = & \hat{d}(t-1) & \text{if } \hat{d}(t-1) = d(t-1), \\ < & \hat{d}(t-1) & \text{if } \hat{d}(t-1) > d(t-1). \end{array} \tag{8}$$

In the G sessions, we would expect that players' forecasts of p for period t are conditional on the agent's choice of $d(t)$ as this choice was announced *before* players made their forecasts of p. Therefore, in the G sessions, we suppose that forecasts of p are revised according to (7) if $d(t) = d(t-1)$.

However, if $d(t) \neq d(t-1)$, then we further assume that forecasts $\hat{p}$ obey

$$\hat{p}(t) \quad \begin{array}{lll} \geq & \hat{p}(t-1) & \text{if } d(t) < d(t-1), \\ \leq & \hat{p}(t-1) & \text{if } d(t) > d(t-1). \end{array} \tag{9}$$

The logic behind the revision process (9) follows from the fact that lower levels of defense encourage greater predation, and higher levels deter predation; this logic was readily apparent from the payoff tables that subjects could consult in making their forecasts and choosing which type of player they wanted to be.

We examined how frequently these simple models of forecast revision characterize the forecasts made by each subject in both the NG and G treatments. Table 5 reports the frequency with which the learning models (7) and (8) predict subjects' forecasts of p and d, respectively, over all but the first round of each NG session.

[Insert Table 5 here.]

We see that our simple learning model can predict only around 50% of subjects' forecasts of p or d in the 3 NG sessions. This finding is likely due to our extreme assumption that the forecasts $\hat{p}$ ($\hat{d}$) are only dependent on realizations of p (d). We will relax this assumption below. Nevertheless, this finding is interesting as it demonstrates that there is some independence in the manner in which individuals revise their forecasts of p and d.

We also examined how well the learning model described by (7) and (9) characterizes behavior in the three G sessions.

[Insert Table 6 here.]

In Table 6, for each G session, we first report the frequency with which model (7) alone predicts subjects's forecasts, $\hat{p}$, over all but the first round of each G session. Next to this frequency we report the frequency with which the model of forecast revision conditional on Δd, (7) and (9), accurately predicts subjects' forecasts, $\hat{p}|\Delta d$. Perhaps not surprisingly, we see that this more sophisticated model (7) and (9) does a better job predicting subjects' forecasts of $\hat{p}$ than does the simpler model (7). This evidence suggests that subjects are responding to changes in the value of d as chosen by the player designated as the agent.

It is also possible to augment the simple forecast revision/learning model used to predict subjects' forecasts in the NG sessions, so that forecasts of p are conditional on changes in d, and forecasts of d are conditional on changes in p. However, since in the NG treatment, p and d were determined simultaneously, it is not possible for subjects to condition their forecasts of p on the contemporaneous realization of d as in the G treatments. Thus in our more sophisticated, conditional learning model for the NG sessions we assume that forecasts of p are revised according to (7) if $d(t-1) = d(t-2)$. However, if $d(t-1) \neq d(t-2)$, then we further assume that forecasts $\hat{p}$ obey

$$\hat{p}(t) \quad \begin{array}{lll} \geq & \hat{p}(t-1) & \text{if } d(t-1) < d(t-2), \\ \leq & \hat{p}(t-1) & \text{if } d(t-1) > d(t-2). \end{array} \tag{10}$$

Similarly, we suppose that forecasts of d are revised according to (8) if $p(t-1) = p(t-2)$. However, if $p(t-1) \neq p(t-2)$, then we further assume that forecasts $\hat{d}$ obey

$$\hat{d}(t) \quad \begin{array}{lll} \geq & \hat{d}(t-1) & \text{if } p(t-1) > p(t-2), \\ \leq & \hat{d}(t-1) & \text{if } p(t-1) < p(t-2). \end{array} \tag{11}$$

Table 7 reports the frequencies with which this more sophisticated, conditional model predicted subjects' forecasts of p and d over all rounds but the first two of each NG session.

[Insert Table 7 here.]

A comparison of the frequencies reported in Table 7 with those reported in Table 5 reveals that the more sophisticated, conditional model of forecast revision is able to explain a greater proportion of subjects' forecasts of p and d. The difference in the reported frequencies between Tables 5 and 7 suggest that individuals are responding to changes in the realizations of both p and d in revising their forecasts of these two variables. While there may be more sophisticated learning models that do a better job of characterizing forecast revision, the differences between the frequencies reported in Tables 5 and 7 indicate that individuals are able to adjust forecasts according to simple best reply dynamics in *two dimensions* rather than in the single dimension that is more typically studied.

6 Conclusion

The experiment reported in this paper tests the predictions of a simple model of "anarchy," where, in equilibrium, some players choose to be producers allocating their endowment of time to both production and the defense of their output while other players choose to be predators, foregoing production and consuming by appropriating some of the producers' output. Predation and defense against predation – "the dark side of economic activity" in the words of Hirshleifer (1994) – have been largely ignored in theoretical and empirical economic analyses despite the obvious importance of such behavior for an understanding of economic behavior and organization.

In contrast to the two existing experimental studies examining "anarchic" equilibria, where subjects were randomly matched in pairs and could play the role of both predator and producer, here we examine whether the equilibrium predictions of our model obtain in larger groups of 10 or 11 players where individuals must choose between being either a predator or a producer. In addition, we asked subjects in each round to make forecasts about aggregate behavior, and we used these forecasts to determine whether subjects were playing best responses to their beliefs. Our experimental design also differs from experimental "market entry" games in that there are no certain payoffs and subjects have to employ multi–dimensional decision–making skills.

We found that the outcomes achieved were generally in the neighborhood of the Nash equilibrium prediction of the model and that subjects were frequently playing best responses to their forecasts of the number of players who would choose to be predators and of the average fraction of time that Type B players (producers) would devote to defense. This finding is most impressive in the NG sessions, where subjects had to make multi–dimensional forecasts and decisions. Indeed, this finding should be of more general interest, (beyond the predator-prey framework we examine) as it suggests that the Nash equilibrium concept continues to have substantive predictive power even in more complex decision making environments. Nevertheless, the success of the Nash equilibrium prediction was not as great as when the forecast problem/decision space was uni–dimensional. Indeed, by the end of our second treatment, where a single player assumed the role of the government agent and chose in advance the level of d for all producers, the 10 "citizen" subjects were able to make their single forecast and action choice nearly always coincide with the Nash equilibrium prediction. In the treatment without a government, subjects had some difficulty making forecasts and decisions so as to achieve the "anarchic," Nash equilibrium. In particular, the producers frequently chose to allocate too little time to defense leading to greater predation than is predicted to occur in equilibrium. While this problem appeared to lessen with time, there continued to be substantial variance in the fraction of time that producers devoted to defense. In one session we observed a few instances of "amorphic" outcomes, where there was little or no production, although such outcomes were never sustained. We also found that subjects' forecasts of these two variables appear to be influenced by ex post realizations of both the number of players choosing to be predators, pN and the average fraction of time that producers devote to defense, d.

Finally, our results underscore the important role the state plays in eliminating the possibility of free-riding by optimally solving a collective action problem with respect to the level of defense. Advocates of minimal government intervention and private militias often overlook this important role. Our findings suggest that the role of the state in making and enforcing a collective choice in the level of defense can indeed lead to a substantial improvement in the welfare of agents relative to the stateless, anarchic equilibrium.

There are many possible avenues for future research that would serve to clarify our findings and expand upon our experimental design. Certainly, further experimental sessions would be useful in understanding how likely we are to observe a breakdown of anarchy in the direction of amorphy as we observed in one of the three NG sessions. Further experimental sessions could also address whether coordination on the anarchic Nash equilibrium is affected by a change in the appropriation technology parameter, θ, that determines the share of output, f, that accrues to producers. Lowering the value of θ from the single value used in this study would increase the fraction of output accruing to producers and thus the equilibrium number of producers in the anarchic equilibrium. With more producers making a choice of how much time to allocate to defense, we might see better coordination on the Nash equilibrium prediction. More extensive changes to the experimental design could yield further insights. As noted earlier, we might also consider changing the distribution of output so that a certain fraction of the population had a slight comparative advantage in predation while the remaining fraction had a slight comparative advantage in production, with this fraction corresponding to the equilibrium fraction of predators to producers. The aim of this change would be to see whether such differences promoted convergence on the Nash equilibrium and led to individuals having "careers" as predators or producers. Another extension would be to modify the government treatment so that the government only *suggests* a level of defense for individual producers to adopt but does not enforce this suggestion. Such a treatment would allow us to understand whether the government's role is to aid in solving a coordination problem among agents who are unable to communicate with one another (i.e. the use of "moral suasion") or whether it is to prevent free–riding by making a binding collective action decision for all producers.[16] Finally, we could explore alternatives to the government institution we consider. Instead of having a single agent play the role of the government, we could have several agents play this role, making decisions as a committee, or taking turns "in office". We could explore the possible emergence of the state's collective action choice, e.g., by allowing producers to communicate with one another, or to vote on a collective choice (tax) for defense, or even to elect one of their own to play the role of the government agent. There are indeed many further possibilities to consider, and we hope to pursue these in future research.

[16] Van Huyck et al. (1992) show that in certain types of coordination games, players will often follow the non-binding advice of a disinterested third party as to how to play the game. Similarly, Schotter and Sopher (2001) report that players in coordination games will often follow the non-binding advice offered by players who have previously played the same coordination game.

References

Anderton, C.H., Anderton, R.A., and Carter, J.R., 1999. Economic activity in the shadow of conflict. Economic Inquiry 37, 166–179.

Camerer, C. and Lovallo, D., 1999. Overconfidence and excess entry: An experimental approach. American Economic Review 89, 306–318.

Carter, J.R. and Anderton, C.H., 2001. An experimental test of a predator–prey model of appropriation. Journal of Economic Behavior and Organization 45, 83–97.

Durham, Y., Hirshleifer, J., and Smith, V.L., 1998. Do the rich get richer and the poor poorer? Experimental tests of a model of power. American Economic Review 88, 970–983, reprinted in Hirshleifer (2001).

Foss, N.J., 2001. Leadership, beliefs and coordination: An explorative discussion. Industrial and Corporate Change 10, 357-388.

Grossman, H.I., 2002. Make us a king: Anarchy, predation and the state. European Journal of Political Economy 18, 31–46.

Grossman, H.I., and Kim, M., 1995. Swords or plowshares? A theory of the security of claims to property. Journal of Political Economy 103, 1275–1288.

Grossman, H.I., and Kim, M., 1996. Predation and accumulation. Journal of Economic Growth 1, 333–350.

Grossman, H.I. and Kim, M., 2000. Predators, moral decay, and moral revivals. European Journal of Political Economy 16, 173–187.

Hirshleifer, J., 1994. The dark side of the force. Economic Inquiry 32, 1–10, reprinted in Hirshleifer (2001).

Hirshleifer, J., 1995. Anarchy and its breakdown. Journal of Political Economy 103, 26–52, reprinted in Hirshleifer (2001).

Hirshleifer, J., 2001. The dark side of the force: Economic foundations of conflict theory. (Cambridge, Cambridge University Press).

Rapoport, A., 1995. Individual strategies in a market–entry game. Group Decision and Negotiation 4, 117–133.

Rapoport, A., Seale, D.A., Erev, I., and Sundali, J.A., 1998. Equilibrium play in large group market entry games. Management Science 44, 119–141.

Selten, R., and Stoecker, R., 1986. End behavior in sequences of finite prisoner's dilemma supergames: A learning theory approach. Journal of Economic Behavior and Organization 7, 47–70.

Sened, I., 1997. The political institution of private property. (Cambridge: Cambridge University Press).

Schotter, A. and Sopher, B., 2001. Social learning and coordination conventions in inter-generational games: An experimental study. Working paper, NYU and Rutgers University.

Skaperdas, S., 1992. Cooperation, conflict, and power in the absence of property rights. American Economic Review 82, 720–739.

Usher, D., 1987. Theft as a paradigm for departures from efficiency. Oxford Economic Papers 39, 235–252.

Usher, D., 1992. The welfare economics of markets, voting and predation. (Ann Arbor, University of Michigan Press).

Van Huyck, J.B., Battalio, R.C., and Beil, R.O., 1990. Tacit coordination games, strategic uncertainty, and coordination failure. American Economic Review 80, 234–248.

Van Huyck, J.B., Gillette, A.B., and Battalio, R.C., 1992. Credible assignments in coordination games. Games and Economic Behavior 4, 606–626.

Van Huyck, J.B., Battalio, R.C., and Walters, M., 1995. Commitment versus discretion in the peasant dictator game. Games and Economic Behavior 10, 143–170.

Appendix: Instructions and Payoff Table Used in the Experiment

Below are the instructions used in the NG sessions. The instructions for the G sessions are similar. Also included in this appendix is the payoff table that was used in all sessions (both NG and G treatments) of the experiment reported in this paper.

Instructions

Overview

You are about to participate in an experiment in the economics of decision–making. Funding for this experiment has been provided by the University of Pittsburgh Central Research Development Fund. If you follow these instructions carefully and make good decisions you might earn a considerable amount of money that will be paid to you in cash at the end of the session. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with one another during the experiment.

This experimental session involves a total of 10 players. Each of you will play a game with the other 9 players for a number of rounds. In each round, you will be asked to make some choices. Your choices together with the choices made by the other 9 players determine your payoff in each round as will be explained below. While your choices will be reflected in information that is revealed to all players, you will never be identified by name, so others will not be able to associate you with any of the choices you make, even at the end of the session.

In addition to these instructions, you should also have 1) a payoff table, 2) a record sheet and 3) a receipt. Please write down your player ID number at the top of your record sheet.

The Choices You Make in Each Round of the Game

In each round of the game, all 10 players must choose whether to be a “Type A” or “Type B” player. The payoff to you from being a Type A or Type B player will depend on your choice and the choices made by the other 9 players, and are given in the *payoff table* which will be described momentarily. Those who choose in a round to be Type B players each have one unit of “time” which they can devote to one of two actions, “Action 1” or “Action 2.” Those choosing to be Type B players must choose the fraction of time (in 10% increments), 0%, 10%, 20%...100%, they will devote to Action 1; the remaining fraction of time (if any) is assumed to be devoted to Action 2.

The Payoff Table

The payoff table – see Table A.1 – reveals the payoff you earn from choosing to be a Type A or Type B player conditional on the choices made by all players in the game. The payoffs in dollars to being a Type A or Type B are given in each box or "cell" of the table, with the payoff to being a Type A given first, and the payoff to being a Type B given second: Type A payoff, Type B payoff.

[Insert Table A.1 here.]

The rows of the payoff table indicate the number of the other 9 players (players other than yourself) choosing to be Type A players in a round; the players not choosing to be Type A players must have chosen to be Type B players, since all players must choose either to be a Type A or a Type B player in every round. For example, if 0 of the other 9 players choose to be Type A players, then it must be that all 9 of these other players have chosen to be Type B players. In this case, your payoff possibilities would be given by the first row of payoff pairs in the payoff table. For another example, if 2 of the other 9 players have chosen to be Type A players, then it must be that 7 of the other 9 players have chosen to be Type B players. In this case, your possible payoffs are given by the third row of payoff pairs in the payoff table. Notice that, for any given column of the payoff table, the payoff to being either a Type A or Type B player decreases, or does not change as the number of the other 9 players choosing to be Type A players increases.

The columns of the payoff table (which alternate in color between white and grey) reveal the *average fraction of time* that all Type B players have chosen to devote to Action 1. For example, if four players choose to be Type B players in a round, and their choices of the fraction of time to devote to Action 1 are 60%, 20%, 70% and 10%, the *average* fraction of time devoted to Action 1 by all four Type B players is 40%, and this is the fraction that is used to determine all players' payoffs for that round, which in this case would correspond to the fifth column of the table of payoff values. If necessary, the average fraction of time that all Type B players devote to Action 1 is rounded to the nearest 10% increment. If no player chooses to be a Type B, then the fraction of time that Type B players devote to Action 1 is 0. Notice that for any given row of the payoff table, the payoff to being a Type A player decreases or does not change as the average fraction of time that Type B players devote to Action 1 increases, while the payoff to being a Type B player may increase or decrease as the average fraction of time that Type B players devote to Action 1 increases.

The *actual* payoff you receive in a round depends on your decision to be a Type A or Type B player, as well as on the number of the other 9 players choosing to be Type A players and the average fraction of time that all Type B players devote to Action 1. To find this payoff from the payoff table, you look at the row that corresponds to the number of the other 9 players who actually chose to be Type A, and the column that corresponds to the average fraction of time that all Type B players actually devoted to Action 1. The intersection of this row and column results in a single cell of the payoff table, which contains two numbers. The first number is the payoff you receive if you chose to be a Type A player and the second number is the payoff you receive if you chose to be a Type B player

Practice Test # 1

Before proceeding further with the instructions we ask each of you to answer some practice questions to make sure you understand how to read the payoff table. Write down your answers to the following questions in the spaces provided. Your answers will be checked and any mistakes you make will be corrected.

Suppose that in a round, 3 of the 10 players chose to be type A players, and the remaining 7 chose to be type B players. These type B players devoted an average of 40% of their time to Action 1. Suppose you chose to be a Type A player.

Q1. How many of the other 9 players (excluding yourself) chose to be Type A players? _____

Q2. What payoff did you earn as a Type A player? ______________

Q3. What payoff would you have earned if you had instead chosen to be a Type B player assuming that the other 9 players did not change their choice of type? ______________

Please wait while we review the answers to questions Q1–Q3.

Other Choices You Make

Before choosing whether to be a Type A or Type B player in each round, you are asked to make choices in response to two questions which are presented on the first computer screen you see in each round. The first question you are asked is:

`How many of the other 9 players, excluding yourself, do you think will choose to be Type A players in this round?`

You answer this question by choosing one of the 10 possibilities, given on your computer screen: 0,1,2...9. Use the mouse to click on the radio button next to your choice.

The second question you are asked is:

`What do you think will be the average fraction of time that all Type B players devote to Action 1 in this round?`

You answer this question by choosing one of the 11 possibilities, given on your computer screen: 0%, 10%, 20% ... 100%. Again, you use the mouse to click on the radio button next to your choice.

NOTE: If you choose to be a Type B player, the fraction of time YOU choose to devote to Action 1 is set equal to your expectation of the average fraction of time that all Type B players devote to Action 1 (your answer to question 2). So, if you choose to be a Type B player your response to this second question and the responses of the other players choosing to be Type B players are summed together, and this sum is divided by the total number of Type B players to determine the average fraction of time that all Type B players devote to Action 1 in the round.

If you want to change your answer to either the first or second question, just move your mouse to the choice you want to change to and click on the radio button next to this choice. You must answer both questions. We ask these questions in order to understand your expectations prior to the play of each round of the game. Also, as we have explained, if you choose to be a Type B player, your answer to the second question determines the fraction of time you, as a Type B player, devote to Action 1. If you choose to be a Type A player, your forecast of the average fraction of time that Type B players devote to Action 1 will not figure in the calculation of this average fraction.

After you have answered the two questions, you are asked to decide whether you want to be a Type A player or a Type B player. You are asked:

`Which type of player do you want to be in this round?`

Use your mouse to click on the radio button next to your choice, Type A or Type B. You are free to choose to be either type of player in every round of the game. Again, if you want to change your choice, move your mouse to the choice you would prefer and click on the radio button next to this choice.

When you are satisfied with your answer to the first two questions and your decision to be either a Type A or Type B player, record all three of your choices for the round on your record sheet. Then click on the Submit Choices button at the bottom of your screen to submit your choices for the round. Once your choices have been submitted, please wait for the results of the round. If you want to clear all of your answer choices and start over, click on the Reset Choices button. If you click on this button you have to resupply answers to the first two questions and choose again whether the be a Type A or Type B player.

The Results of a Round

Once all 10 players have submitted their choices, the results for the round will be reported on a second computer screen. You will see both your forecast and the actual number of the other 9 players who chose to be Type A. You will also see your forecast and the actual average fraction of time that all Type B players devoted to Action 1. Finally, you will be reminded of your choice of type and informed of your payoff for the round. Please record your payoff for the round in the final column of your record sheet. Your history of play will be updated every round and will appear at the bottom of this second "results screen" for your personal reference.

Practice Test # 2

Again, before proceeding, we ask that you answer some test questions to make sure you understand how your choices can affect your payoffs. Write your answers to the following questions in the space below. Your answers will be checked and any mistakes you make will be corrected.

Suppose that in this round you forecast that 3 of the other 9 players will choose to be Type A players and you also forecast that the average fraction of time that all Type B players devote to Action 1 will be 80%. Suppose also that you choose to be a Type B player.

Q4. What fraction of time have you chosen to devote to Action 1? ______________

Q5. Using the payoff table, what payoff would you receive if both of your forecasts were correct?

____

Continuing with this same example, suppose that after all players have submitted their choices you learn that 5 of the other 9 players actually chose to be Type A players and that the average fraction of time that all Type B players actually devoted to Action 1 was 40%.

Q6. Using the payoff table, what is your actual payoff for this round? ______________

Please wait while we review the answers to questions Q4–Q6.

Payments

At the end of the session we will ask you to total up your earnings from all rounds played and record the sum at the bottom of your record sheet (we have a calculator available to perform this task).

Add to this sum the \$5 participation payment and record this total amount on your receipt. All earnings will be paid in cash at the end of the session.

ARE THERE ANY QUESTIONS BEFORE WE BEGIN?

Table 1: Means (Standard Deviations) of d and pN
Over all 30 rounds and Over the Last 10 Rounds of Each session

Session Number(s)	Average Value or Choice of d		Number of Type As, pN	
	All 30 Rounds	Last 10 Rounds	All 30 Rounds	Last 10 Rounds
NG1	0.5600 (0.2143)	0.7000 (0.1333)	7.00 (1.36)	7.10 (1.10)
NG2	0.3033 (0.1650)	0.2900 (0.2079)	8.03 (1.19)	8.50 (0.85)
NG3	0.6333 (0.1845)	0.6800 (0.1135)	7.03 (1.43)	6.80 (1.23)
All 3 NG	0.4989 (0.2349)	0.5567 (0.2445)	7.36 (1.40)	7.47 (1.28)
Nash Eq. Prediction	0.7000	0.7000	7.00	7.00
G1	0.7100 (0.1322)	0.7900 (0.0316)	3.13 (3.75)	0.40 (1.26)
G2	0.7300 (0.1442)	0.8000 (0.0000)	2.57 (3.21)	0.10 (0.32)
G3	0.7533 (0.0900)	0.7800 (0.0422)	2.57 (3.32)	1.20 (2.30)
All 3 G	0.7311 (0.1242)	0.7900 (0.0305)	2.76 (3.41)	0.57 (1.55)
Nash Eq. Prediction	0.8000	0.8000	0.00	0.00

Table 2: Mean Squared Deviations of d and pN from predicted equilibrium values. Over three non-overlapping 10 round periods

Session Number	MSD of d from eq. value			MSD of pN from eq. value		
	First 10	Second 10	Last 10	First 10	Second 10	Last 10
NG1	0.141	0.035	0.016	2.8	1.5	1.1
NG2	0.164	0.18	0.207	2.6	1.8	2.9
NG3	0.088	0.012	0.012	3.2	1.3	1.4
All 3 NG	0.131	0.076	0.078	2.9	1.5	1.8
G1	0.069	0.005	0.001	57.2	11.4	1.6
G2	0.074	0.001	0.000	45.2	4.4	0.1
G3	0.027	0.001	0.002	38.4	7.1	6.2
All 3 G	0.057	0.002	0.001	46.9	7.6	2.6

Table 3: Percentage of Total Nash Equilibrium Payoff Earned by Each Group: Mean and (Standard Deviation) Over Last 10 Rounds of Each Session

Session Number	Mean (St. Deviation)
NG1	0.962 (0.471)
NG2	1.149 (0.785)
NG3	1.172 (0.625)
G1	0.989 (0.034)
G2	0.990 (0.029)
G3	0.923 (0.137)

Table 4: Frequency of Best Response Play by Individual Players
Over All 30 rounds and Over the Last 10 Rounds of Each Session

Subject Number	Session NG1		Session NG2		Session NG3	
	All Rounds	Last 10	All Rounds	Last 10	All Rounds	Last 10
1	0.67	0.20	0.77	0.80	0.87	0.90
2	0.43	0.30	0.87	0.90	0.83	0.70
3	0.87	1.00	0.57	0.60	0.43	0.40
4	0.60	1.00	0.77	0.80	0.93	1.00
5	0.90	0.90	0.90	0.90	0.63	0.70
6	0.80	1.00	0.93	0.90	0.80	0.70
7	0.80	0.80	1.00	1.00	0.70	0.60
8	0.90	0.90	0.53	0.50	0.57	0.50
9	0.83	1.00	0.93	1.00	0.77	0.40
10	0.93	1.00	0.83	1.00	0.73	0.60
All	0.77	0.81	0.81	0.84	0.73	0.65
Type A	0.93	0.92	0.98	0.96	0.82	0.78
Type B	0.41	0.55	0.12	0.13	0.49	0.38
Subject Number	Session G1		Session G2		Session G3	
	All Rounds	Last 10	All Rounds	Last 10	All Rounds	Last 10
1	0.93	0.90	0.77	1.00	0.70	0.70
2	0.70	0.90	0.97	1.00	0.87	0.80
3	0.83	1.00	0.80	1.00	0.83	0.80
4	0.93	0.90	0.63	0.90	0.77	1.00
5	0.93	0.90	0.90	1.00	0.97	1.00
6	0.83	0.90	0.83	1.00	1.00	1.00
7	0.87	1.00	0.80	1.00	0.83	0.80
8	0.83	0.90	0.83	1.00	0.93	1.00
9	0.90	1.00	0.97	1.00	0.97	1.00
10	0.90	1.00	0.97	1.00	0.80	0.80
All	0.87	0.94	0.85	0.99	0.87	0.89

Table 5: The Frequency With Which Individual Subjects' Forecasts of p or d Were Predicted by the Learning Model (7) or (8) in Each NG Session

Subject	Session NG1		Session NG2		Session NG3	
Number	$\hat{p}$ predicted	$\hat{d}$ predicted	$\hat{p}$ predicted	$\hat{d}$ predicted	$\hat{p}$ predicted	$\hat{d}$ predicted
1	0.45	0.45	0.59	0.55	0.41	0.31
2	0.59	0.45	0.38	0.45	0.66	0.48
3	0.48	0.48	0.62	0.62	0.62	0.52
4	0.59	0.38	0.59	0.66	0.34	0.14
5	0.31	0.41	0.48	0.45	0.76	0.38
6	0.55	0.48	0.34	0.52	0.48	0.52
7	0.48	0.52	0.34	0.45	0.55	0.45
8	0.52	0.66	0.41	0.62	0.48	0.48
9	0.34	0.38	0.45	0.52	0.38	0.48
10	0.55	0.69	0.41	0.41	0.52	0.66
All	0.49	0.49	0.46	0.52	0.52	0.44

Table 6: The Frequency With Which Individual Subjects' Forecasts of p
Were Predicted by the Learning Model (7) or (7) and (9) in Each G Session

Subject	Session G1		Session G2		Session G3	
Number	$\hat{p}$ predicted	$\hat{p}\|\Delta d$ predicted	$\hat{p}$ predicted	$\hat{p}\|\Delta d$ predicted	$\hat{p}$ predicted	$\hat{p}\|\Delta d$ predicted
1	0.62	0.86	0.59	0.83	0.45	0.52
2	0.83	0.76	0.66	0.72	0.45	0.83
3	0.76	0.83	0.59	0.59	0.52	0.79
4	0.79	0.90	0.59	0.79	0.45	0.76
5	0.66	0.90	0.59	0.76	0.55	0.59
6	0.72	0.86	0.62	0.79	0.48	0.76
7	0.66	0.79	0.66	0.79	0.41	0.72
8	0.66	0.83	0.62	0.72	0.52	0.76
9	0.69	0.76	0.59	0.83	0.52	0.90
10	0.76	0.79	0.48	0.72	0.59	0.76
All	0.71	0.83	0.60	0.76	0.49	0.74

Table 7: The Frequency With Which Individual Subjects' Forecasts of p or d
Were Predicted by the Learning Model (7) and (10) or (8) and (11) in Each NG Session

Subject	Session NG1		Session NG2		Session NG3	
Number	$\hat{p}\|\Delta d$ predicted	$\hat{d}\|\Delta p$ predicted	$\hat{p}\|\Delta d$ predicted	$\hat{d}\|\Delta p$ predicted	$\hat{p}\|\Delta d$ predicted	$\hat{d}\|\Delta p$ predicted
1	0.75	0.64	0.46	0.57	0.50	0.57
2	0.71	0.57	0.75	0.68	0.46	0.46
3	0.68	0.57	0.64	0.43	0.75	0.61
4	0.71	0.75	0.68	0.64	0.82	0.75
5	0.71	0.68	0.54	0.61	0.75	0.79
6	0.50	0.57	0.57	0.57	0.46	0.46
7	0.57	0.50	0.61	0.68	0.57	0.79
8	0.54	0.46	0.54	0.71	0.68	0.57
9	0.61	0.61	0.68	0.75	0.57	0.71
10	0.71	0.75	0.61	0.50	0.71	0.75
All	0.65	0.61	0.61	0.61	0.63	0.65

Table A.1: Payoff Table Used in the Experiment

No. of Other 9 Players Who Choose to be Type A:	Fraction of time that all Type B players devote to Action 1: 0%		10%		20%		30%		40%		50%		60%		70%		80%		90%		100%	
0	54.00	6.00	37.55	5.40	26.00	4.80	17.71	4.20	11.72	3.60	7.40	3.00	4.35	2.40	2.26	1.80	0.93	1.20	0.22	0.60	0.00	0.00
1	24.00	0.00	19.10	1.23	14.84	1.91	11.17	2.23	8.07	2.30	5.51	2.18	3.47	1.92	1.92	1.55	0.84	1.10	0.21	0.58	0.00	0.00
2	14.00	0.00	11.71	0.62	9.56	1.09	7.57	1.41	5.76	1.58	4.15	1.62	2.76	1.53	1.61	1.32	0.75	0.99	0.20	0.55	0.00	0.00
3	9.00	0.00	7.72	0.38	6.48	0.70	5.30	0.95	4.17	1.13	3.12	1.22	2.17	1.22	1.33	1.11	0.65	0.88	0.18	0.52	0.00	0.00
4	6.00	0.00	5.23	0.25	4.47	0.48	3.73	0.67	3.01	0.82	2.32	0.92	1.67	0.96	1.07	0.91	0.55	0.77	0.16	0.48	0.00	0.00
5	4.00	0.00	3.52	0.17	3.05	0.33	2.58	0.47	2.12	0.59	1.67	0.68	1.24	0.73	0.82	0.73	0.45	0.65	0.14	0.44	0.00	0.00
6	2.57	0.00	2.28	0.12	1.99	0.22	1.71	0.33	1.42	0.42	1.14	0.49	0.87	0.55	0.60	0.57	0.34	0.53	0.12	0.38	0.00	0.00
7	1.50	0.00	1.34	0.07	1.18	0.15	1.02	0.22	0.86	0.28	0.70	0.34	0.54	0.38	0.38	0.41	0.23	0.40	0.09	0.32	0.00	0.00
8	0.67	0.00	0.60	0.04	0.53	0.09	0.46	0.13	0.39	0.17	0.32	0.21	0.25	0.24	0.19	0.26	0.12	0.27	0.05	0.24	0.00	0.00
9	0.00	0.00	0.00	0.02	0.00	0.04	0.00	0.06	0.00	0.08	0.00	0.09	0.00	0.11	0.00	0.13	0.00	0.14	0.00	0.14	0.00	0.00

Note: The first number in each cell is the payoff to being a Type A player, while the second number in each cell is the payoff to being a Type B player.

Fig 1.1a Average Fraction of Time Devoted to Action 1 by Type Bs
NG Session # 1

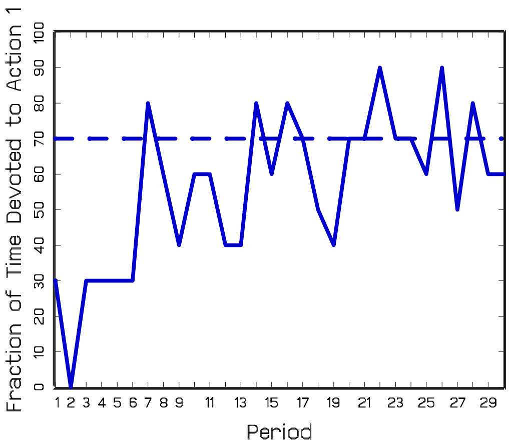

Fig 1.2a Average Fraction of Time Devoted to Action 1 by Type Bs
NG Session # 2

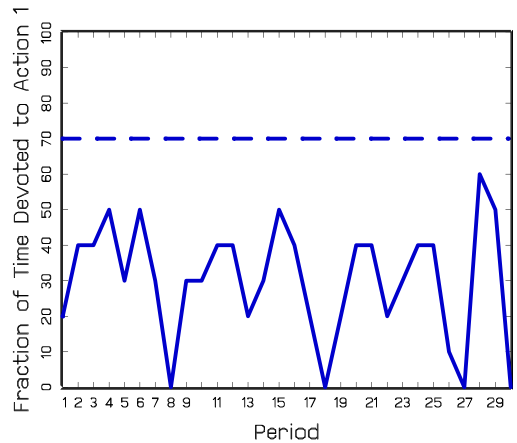

Fig 1.3a Average Fraction of Time Devoted to Action 1 by Type Bs
NG Session # 3

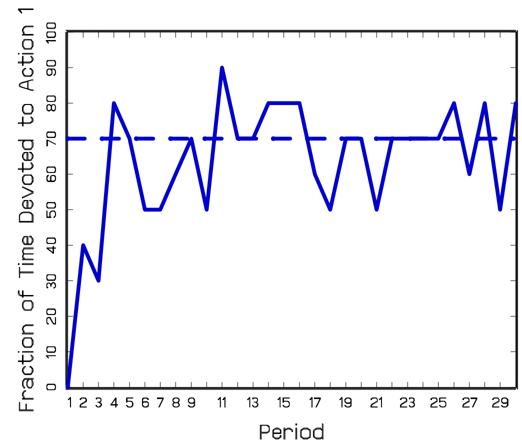

Fig 1.1b Number of Type As
NG Session # 1

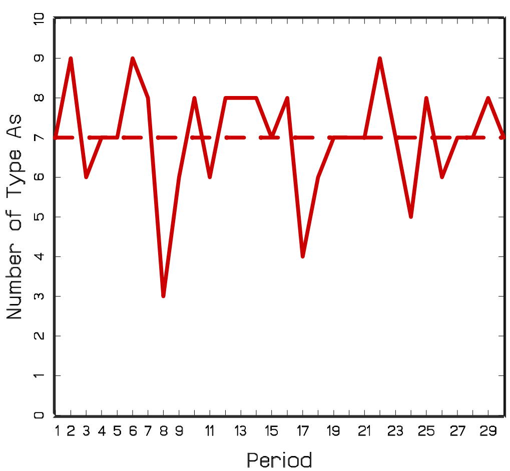

Fig 1.2b Number of Type As
NG Session # 2

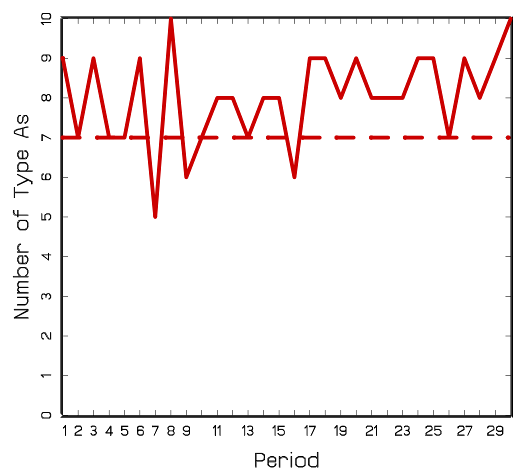

Fig 1.3b Number of Type As
NG Session # 3

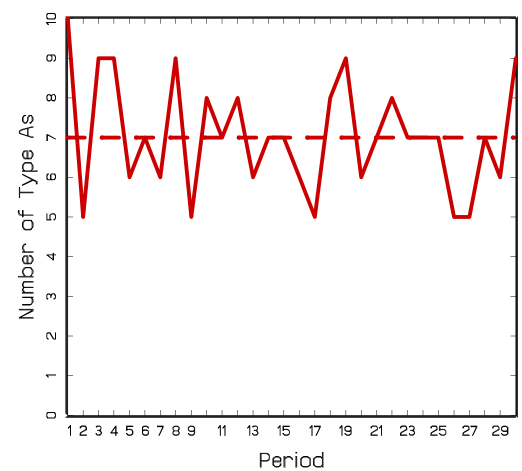

Fig 2.1a Agent's Choice of the Fraction of TIme Devoted to Action 1 G Session # 1

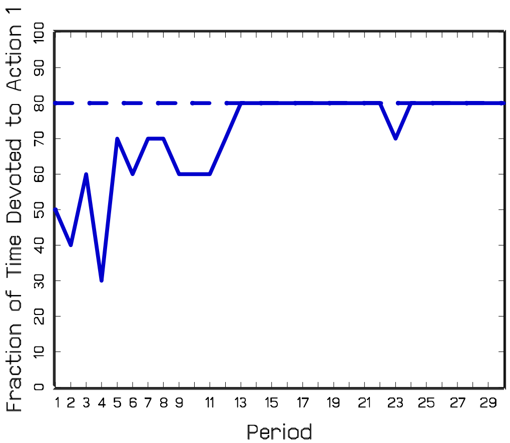

Fig 2.2a Agent's Choice of the Fraction of Time Devoted to Action 1 G Session # 2

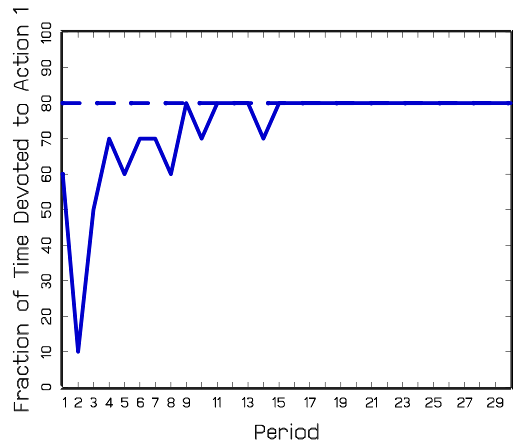

Fig 2.3a Agent's Choice of the Fraction of Time Devoted to Action 1 G Session # 3

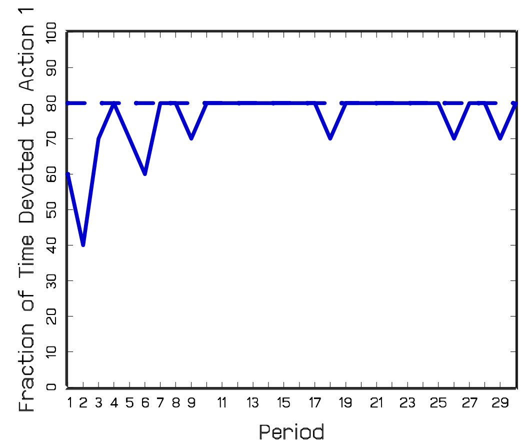

Fig 2.1b Number of Type As G Session # 1

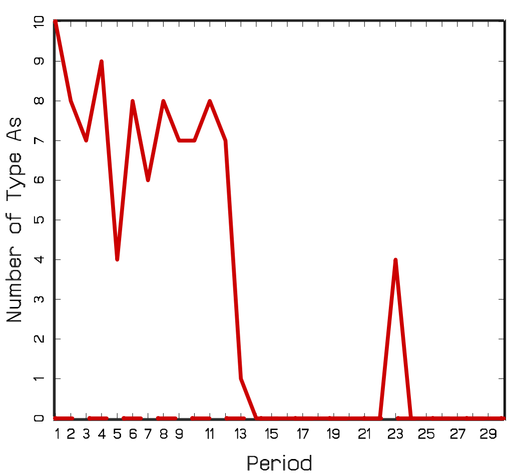

Fig 2.2b Number of Type As G Session # 2

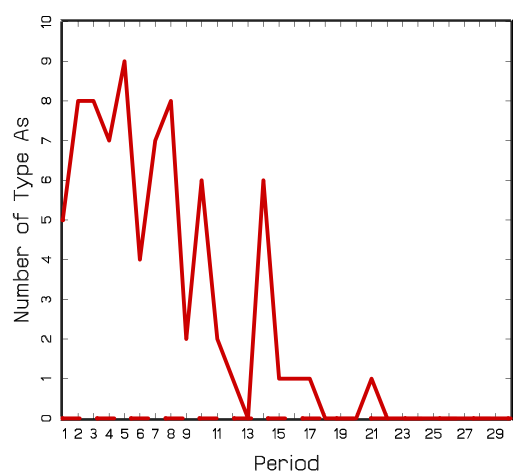

Fig 2.3b Number of Type As G Session # 3

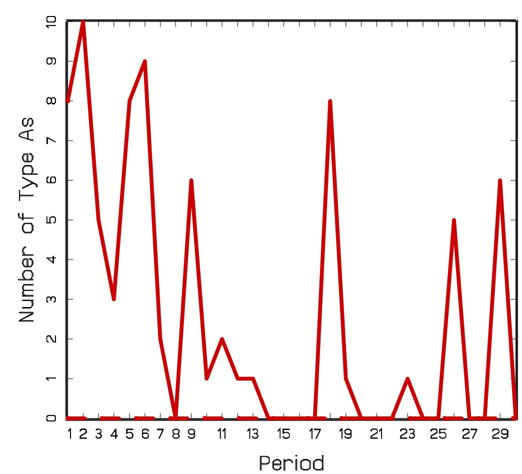

**Fig 3.1 Frequency With Which Each Player Chose to Be Type A
and Their Average Forecast of the Value of d
Over the Last 10 Rounds of Session NG1**

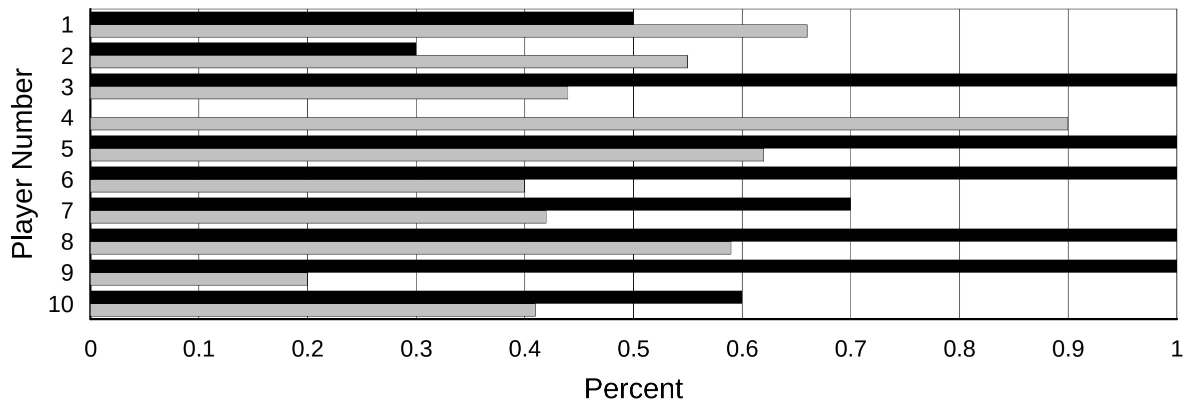

**Fig 3.2 Frequency With Which Each Player Chose to Be Type A
and Their Average Forecast of the Value of d
Over the Last 10 Rounds of Session NG2**

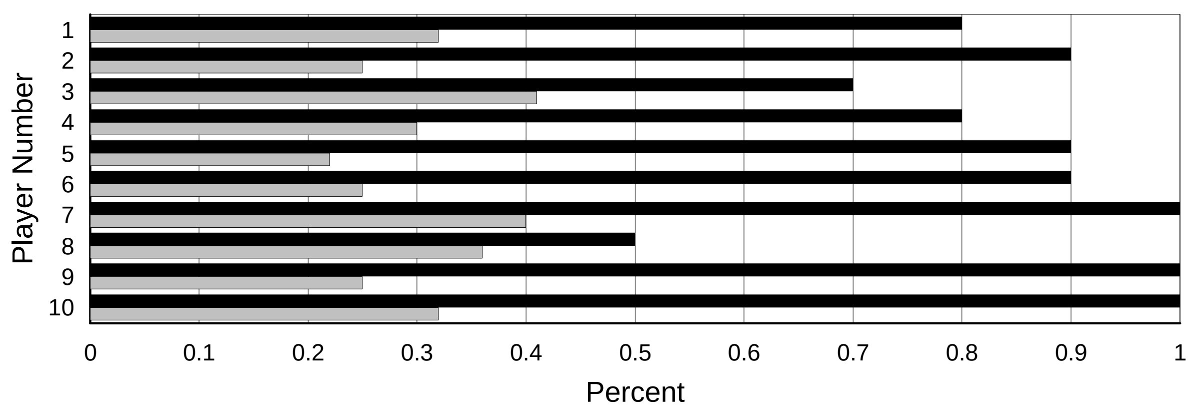

**Fig 3.3 Frequency With Which Each Player Chose to Be Type A
and Their Average Forecast of the Value of d
Over the Last 10 Rounds of Session NG3**