In Support of Trigger Strategies: Experimental Evidence from Two-Person Noncooperative Games

Charles F. Mason and Owen R. Phillips

University of Wyoming
Laramie, WY 82071-3985

Cooperative equilibria can be supported in a repeated game when players use trigger strategies. This paper tests how well trigger strategies explain behavior in two-person experimental games. Reducing payoffs for choices larger than the Cournot level induces smaller average outputs, behavior generally consistent with trigger strategy models. Reducing payoffs for choices well above the Cournot level will not affect behavior if actions are consistent with a trigger strategy involving longer-lived, less intense punishment phases (the grim-reaper strategy), but would matter for trigger strategies with short-lived but intense punishment phases. Results show that behavior is most consistent with the former.

1. Introduction

In many one-shot games the Nash equilibrium is a Pareto-inferior outcome. Private incentives, however, prevent agents from committing to actions that would support an improved outcome. The dilemma faced by agents is plain. By agreeing to Pareto-superior actions, rivals collectively generate higher payoffs, but each agent knows that a substantial reward can be reaped by unilaterally deviating from this agreement. With repeated interaction agents can orchestrate cooperation if they use trigger strategies. A trigger strategy promises cooperative choices so long as no party has defected from the implicit agreement in past play. Following a defection of any sort, all players enter a punishment phase. This punishment phase can be a permanent reversion to

This work has been supported by the College of Business, University of Wyoming. The findings and conclusions expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agency. Helpful comments were received on earlier versions of this paper from Jerry Dwyer; participants at the Latin American Econometric Society meetings, the Public Choice meetings, and the Southern Economic Association meetings; seminar participants from the Department of Economics at Texas A&M University; and two anonymous referees and a coeditor of this journal. Remaining errors are the responsibility of the authors.

Journal of Economics & Management Strategy, Volume 11, Number 4, Winter 2002, 685–716

the one-shot Nash equilibrium (Friedman, 1967), or it can be short-lived but intense (Abreu, 1986, 1988; Segerstrom, 1988). For a trigger strategy to support a Pareto-superior outcome, it must create sufficiently large long-term gains to overcome the short-term incentive to defect.

Models that use a trigger strategy to support a preferred outcome have been successfully applied to a large range of topics. These include the regulation of industry (Salant and Woroch, 1992), efficiency wages (Akerlof and Yellen, 1986; Shapiro and Stiglitz, 1984), human-capital acquisition (Prendergast, 1993), international pollution abatement (Folmer et al., 1993), the formation of joint ventures (Marx and Matthews, 2000), monetary policy (Barro and Gordon, 1983), and the repayment of sovereign debt (Grossman and Van Huyck, 1988). A particularly large literature uses trigger strategies to model cartel formation among oligopolists [see Fudenberg and Tirole (1989) for a survey]. Examples include analysis of collusion in multiple markets (Bernheim and Whinston, 1990), the influence of business cycles upon collusive strategies (Rotemberg and Saloner, 1986; Haltiwanger and Harrington, 1991), and collusion under conditions of uncertain demand (Green and Porter, 1984). Many of the above analyses use a simple form of trigger strategy, wherein defection triggers perennial reversion to the one-shot Nash equilibrium; we refer to this as the *grim-reaper* strategy below. In contrast, more elaborate trigger strategy models that use shorter but more intense punishment phases have been shown to support higher levels of cooperation (Abreu, 1986, 1988).

The ubiquity of trigger-strategy models underscores their theoretical importance. Even so, there have been few attempts to empirically test the hypothesis that agents actually employ such strategies, or behave as if they do. In principle, identifying the use of trigger strategies is impeded by their very nature: if the strategy is successful, punishments should not be observed.[1] Only when exogenous events

1. There are some examples of explicit agreements with codified rules for retaliation. Under the rules of the General Agreement on Trade and Tariffs, if a country raises tariffs above their negotiated level, the other countries are authorized to punish the offender by raising tariffs on the imports from the offending country. The rules state that the punishment must be "substantially equivalent" to the lost value of exports caused by the original tariff deviation. In 1984, the European Community punished the United States for increasing tariffs on European steel by raising tariffs on a long list of US exports, including steel. Thus, trigger strategies were encouraged, and it appears that they were used by members of the Agreement. As a second example, since the 1970s there have been numerous instances in the production of crude oil where members of the OPEC cartel raised their production in response to one member exceeding its production quota. This behavior is consistent with models of cooperation enforced via a trigger strategy. The punishment phase of the strategy usually ended when the violator returned to its assigned quota.

alter payoff functions would we anticipate changes in actions. Since any observed actions in the field can be caused by a combination of events, drawing inferences about strategies from behavior is difficult. Such problems aside, Porter (1983) and Ellison (1994) have undertaken careful analyses of the prices charged by the Joint Executive Committee, a railroad cartel during the late 1800s. Interpreting periods of high prices as cooperative phases and periods of low prices as punishment phases, these authors draw conclusions about the length and nature of punishment phases in the cartel. For members of the cartel the punishment for defecting had a finite length, and usually was short-lived.

An alternative perspective is that sophisticated repeated-game strategies are not empirically important because they require agents to employ hyperrational analysis that exceeds most cognitive abilities. As a general rule, learning takes place in a noisy environment where it is difficult to interpret the actions of a rival (Green and Porter, 1984). There are limits to learning; the ability of managers to understand experiences and put them in perspective is bounded (Levinthal and March, 1993). Under this more realistic view, the play of boundedly rational agents evolves to some behavioral rules of thumb. Even so, over time agents may play strategies that appear quite similar to sophisticated repeated-game strategies, and because of this we believe that connections between actions and intentions must be interpreted with caution.

We therefore are led to a consideration of the comparative statics of cooperation. We compare the outcomes of games in order to assess their consistency with trigger strategies. No claims are made about agents using certain trigger strategies; our purpose is to determine if play can be explained by this game-theoretic framework. To avoid some of the inferential difficulties alluded to above, we obtain our data using laboratory markets.[2] We analyze the choices of subjects

2. Palfrey and Rosenthal (1994) use two experimental public-good contribution games to analyze the role of repeated play in facilitating cooperation. One of these is mathematically equivalent to a one-shot game, while the second is equivalent to an infinitely repeated game. They find that subjects' contributions cannot be distinguished from the Nash equilibrium levels in the first game, and that there is an increase in contributions in the second. Nevertheless, this increase is minor. At best the authors interpret these results as weak support that repeated play leads to more cooperative behavior. Sell and Wilson (1996) attempt to identify the degree to which subject behavior is consistent with the grim-reaper strategy using experiments. They use a four-player infinitely repeated prisoner's-dilemma game. In one of their designs, subjects can freely choose between the cooperative and noncooperative actions in each period. In their second design, anyone who chooses the noncooperative action in some period τ is bound to that choice in all periods $t > \tau$. They find slightly more cooperation in the second treatment, and so conclude that subject behavior is not completely consistent with the grim-reaper strategy in the first treatment.

in several two-person noncooperative games. Subjects make choices from payoff tables that are common knowledge. The setup is like duopolists making quantity choices. To make our games mathematically equivalent to infinitely repeated games with discounting, we invoke a random stopping rule; a fixed continuation rule is applied at the end of each period after a prespecified number of periods is reached (Fudenberg and Tirole, 1989; Rasmusen, 1994).

We consider three versions of the stage game, so that we can distinguish between two repeated-game strategies that are popular in the literature. Our primary design acts as a control. We then devise two other games that manipulate the gains from deviation. In the second game, we reduce the payoffs from all deviation outputs above the one-shot Cournot-Nash equilibrium level, including the one-period best response to the average choice we observed in the control experiment. If subject behavior is consistent with the use of trigger strategies, choices in this treatment should be more cooperative than choices in the control game. In the third game, we reduce payoffs for choices that are greater than the one-period best response to the estimated equilibrium choice from the control treatment. This treatment could influence behavior for trigger strategies that have limited but more intense punishment phases. Reducing the payoffs of severe punishment choices reduces the profits earned during the punishment phase. Accordingly, the threatened punishment is more intimidating. This intimidation should facilitate enhanced cooperation if subjects' behavior is consistent with the use of limited but intense punishment phases (Abreu, 1986). On the other hand, if behavior is consistent with the grim-reaper trigger strategy, which has a punishment phase that permanently reverts to the one-shot Nash choice at the time of defection, outcomes in this game should not differ from the control game.

The results from our analysis indicate that subjects' behavior is significantly more cooperative in the second treatment than in the primary treatment. We therefore conclude that subjects' behavior is consistent with the use of trigger strategies generally. On the other hand, behavior in the third treatment is not markedly different from behavior in the primary treatment, which is less consistent with the use of a strategy that employs shorter and more intense punishments than with a longer-lived, less intense punishment phase.

We also consider two other structural changes to the game design. In the first variation, we reduce the continuation probability toward the end of the experiment. This alteration lowers the discount factor (Rasmusen, 1994, pp. 108–110) and so should reduce the level of cooperation, if subject behavior is consistent with the use of trigger strategies. Our analysis indicates that choices are significantly larger in this treatment than in the primary treatment. In the second variation, we replicate the results from the first set of experiments in an experimental market design that is a stationary repeated game. Payoffs in these experiments are discounted from period to period at a fixed rate until a prespecified period is reached, at which time a random stopping rule is engaged. The discount rate and termination probability are chosen carefully to preserve stationarity.[3] The results from this treatment corroborate the findings from the original set of experiments.

The next section gives a theoretical discussion of trigger strategies; here we provide more details on the effect our proposed payoff restrictions should have on the level of cooperation. Section 3 describes the laboratory markets we construct and the resulting data set. These data are analyzed in Section 4. In Section 5 we discuss the two additional structural changes. We conclude in Section 6 with a brief discussion.

2. The Trigger Strategy Model in a Two-Person Game

We consider a two-person infinitely repeated game with payoffs similar to a quantity-choosing duopoly market. Two players (1 and 2) simultaneously select outputs q_{1t} and q_{2t} in each period $t = 1, 2, \ldots$. In any period t, player i's payoffs are generated by the stage-game payoff function $\pi(q_{it}, q_{jt})$ for $i, j = 1, 2$. The weight placed upon next period's payoffs, i.e., the discount factor, is δ for each player. For simplicity, we restrict our attention to stationary, symmetric subgame-perfect equilibria of the repeated game. It can be argued that these are focal, and so are most likely to be selected from the set of subgame-perfect equilibria.

With a grim-reaper strategy, each agent chooses q^c so long as both players have done so in all previous periods; however, if either party defects in some period τ, both play the one-shot Nash-equilibrium choice q^n in all future periods $t = \tau + 1, \tau + 2, \ldots$.

3. There is a trade-off between the mathematical consistency of this design and the added complexity in the presentation to subjects. Rather than prejudge the relative merits of the original design and the more complicated but stationary design, we analyze both.

Let $q^d(q^c)$ be the optimal one-period defection to the cooperative-looking choice, i.e., the Nash best response. Then we may write the cooperative payoffs as $\pi^c \equiv \pi(q^c, q^c)$ and the defection payoffs as $\pi^d \equiv \pi(q^d(q^c), q^c)$. It is evident that both payoffs are uniquely determined by q^c, and so henceforth we shall write them only as functions of q^c. We assume that π^c is a concave function of q^c. Finally, we denote the one-shot Cournot-Nash equilibrium payoffs as π^n. For the choice q^c to be supported as part of a subgame perfect equilibrium path, that is, in order for this strategy to be part of a subgame-perfect equilibrium, these payoffs must satisfy the incentive constraint:

$$\pi^d(q^c) + \frac{\delta\pi^n}{1-\delta} \leq \frac{\pi^c(q^c)}{1-\delta}. \tag{1}$$

The interpretation of this constraint is that players can realize the cooperative payoffs perennially if no one defects. Following any defection, however, play reverts to an infinite sequence of one-shot Cournot-Nash equilibrium choices. Thus, each player must compare the present discounted value of continued cooperation, the right side of (1), against the present discounted value of defection, the left side of (1), which includes the one-shot gain from defection and the discounted flow of one-shot Cournot-Nash equilibrium profits.

It is well known that there are many values of q^c that satisfy this constraint. Indeed, there is a compact interval $[\underline{q}, q^n]$ of values that satisfy the constraint. The lower bound, $\underline{q}$, satisfies (1) as an equality. We denote the most profitable of the values in this interval as q^*; this value yields strategies that payoff-dominate all alternative (stationary symmetric) subgame-perfect equilibria. If the perfectly collusive value q^m satisfies (1), then $\underline{q} \leq q^m = q^*$. If not, then $q^m < \underline{q} = q^*$. The results of the experiments we report below seem more likely to correspond to this second case, so for the remainder of the section we shall focus on the case where $q^m < q^*$. We note that q^* is then implicitly defined by

$$\pi^c(q^*) - (1-\delta)\pi^d(q^*) - \delta\pi^n = 0. \tag{2}$$

Now suppose the payoff function is altered by the incorporation of an additional unit cost for choices above some $\bar{q} > q^*$. We shall discuss two levels of $\bar{q}$: (i) the one-shot Cournot-Nash equilibrium choice q^n, and (ii) a choice slightly above the best-response level $q^d(q^*)$. In the first case, for sufficiently large extra cost, $\bar{q}$ becomes the new optimal defection output for values of $q^c < q^n$. In this scenario, payoffs at choices at or below q^n are not changed, and so π^c and π^n remain unaltered. But π^d must be reduced, since the optimal defection output has a reduced payoff. It follows that the left side of (2) is now positive at q^*, so that the new optimal cooperative output level must be less than q^*. In the second case, none of the payoffs are changed for choices in the relevant range, so (2) remains unchanged and the optimal cooperative output level remains at q^*.[4] We observe that the possible range of cooperative output levels expands in the first case, but does not change in the second case. All else equal, we would therefore expect a reduction in the observed output from a randomly drawn firm in the first case, but no change in the observed output from a randomly drawn firm in the second case.[5]

While we have cast our discussion in terms of a grim-reaper trigger strategy, we can speculate on the effect these two adjustments to payoffs will have on actions under alternative trigger strategies. Consider the optimal penal-code strategy, where agents play strategies with a one-period but more intense punishment phase. Abreu (1986, Lemma 17) shows that this type of strategy will be characterized by two values, which we refer to below as a carrot output (q^c) and a stick output (q^s). Each agent starts off by choosing q^s in period 1. In any period $t > 1$, both choose q^c so long as neither party deviated from the strategy in period $t-1$. If one or both deviated in $t-1$, either by not playing q^c when they were supposed to or by not playing q^s when they were supposed to, then the agent plays q^s in period t. The two outputs q^c and q^s satisfy the conditions

$$\pi^{dc} - \pi^c \le \delta[\pi^c - \pi^s], \tag{3}$$

$$\pi^{ds} - \pi^s \le \delta[\pi^c - \pi^s], \tag{4}$$

where we have used the notation $\pi^c = \pi(q^c, q^c)$, $\pi^{dc} = \pi(q^d(q^c), q^c)$, $\pi^s = \pi(q^s, q^s)$, and $\pi^{ds} = \pi(q^d(q^s), q^s)$. The idea is that the carrot output, q^c, is supported by an aggressive stick output, q^s. The most profitable pair (q^{c*}, q^{s*}) satisfies (3) and (4) as equalities. Accordingly,

4. This second change alters π^d at larger outputs, and so has the potential to create a new optimal output if $\bar{q}$ exceeds $q^d(q^*)$ by only a small amount.

5. We are assuming that the expansion of the possible range of cooperative choices does not alter the relative frequency of two outputs larger than the original q^*. Writing the original distribution as p_0 and the new distribution as p_1, our assumption is $p_0(q_a)/p_0(q_b) = p_1(q_a)/p_1(q_b)$ for any $q_a, q_b \in [q^*, q^n]$. Expansion of the interval must then reduce the expected value of an observed q^c.

we have $q^{c*} < q^n < q^{s*}$ and so $q^d(q^{s*}) < q^n \leq q^d(q^{c*})$. In the linear-quadratic version of the model we use for our experimental design, it turns out that $q^n - q^{c*} = q^{s*} - q^n$ and $q^n - q^d(q^{s*}) = q^d(q^{c*}) - q^n$ $(= [q^n - q^{c*}]/2)$.[6] The hypothetical adjustments that we are considering raise the costs for outputs above $\bar{q}$, which will tend to lower both $q^d(q^{c*})$ and q^{s*}. It is tedious but straightforward to show that the adjustment also reduces both π^{dc} and π^s at the new optimal combination of carrot and stick outputs. Both changes tend to accommodate more cooperative-looking carrot outputs, so we would expect each change to lower q^{c*}.

It also can be argued that imposing costs on larger choices makes smaller choices more focal, perhaps because selected payoff differences are greater. The idea is that if agents making choices from a payoff table see their current or a more cooperative outcome surrounded by very low payoff options, then the high-paying choices become focal to the players. If this is so, behavior should be more cooperative in the two alternative games, which reduce a relatively large number of payoff options, and less cooperative in the original game. Some argument can be made for the second treatment generating the most cooperation, because it reduces more of the off-equilibrium choices than the third treatment.

In sum, identifying the type of strategy that best explains behavior turns on resolving the accuracy of the following two hypotheses:

Hypothesis A: Reducing payoffs for choices at or above the one-shot Nash-equilibrium level will result in more cooperative actions.

Hypothesis B: Reducing payoffs for choices above the optimal defection to the cooperative level will not alter behavior.

If play is consistent with the grim-reaper trigger strategy, both hypotheses should be true. If play is consistent with a trigger strategy that has short-lived and intense punishment phases, Hypothesis A should be true but Hypothesis B should be false, with more cooperative behavior observed in the alternative treatment. Likewise, if subjects tend to focus on more profitable combinations after selected

6. Let $\pi_i(q_i, q_j) = (a - bQ)q_i$, where $Q = q_i + q_j$. The one-shot best-response function is $q^d(q_k) = a/2b - q_k/2$ for $k = i$ or j. It is easy to show that $\pi^{dg} = (a - bq^g)^2/4b$, $g = c$ or s, and that $\pi^g = (a - 2bq^g)q^g$, $g = c$ or s. Algebraic manipulation then yields $(q^n - q^g)^2 = 4(\pi^{dg} - \pi^g)/9b$ for $g = c$ or s, where $q^n = a/3b$. Writing $\Delta = \delta[\pi^c(q^c) - \pi^s(q^s)]$, we have $(q^n - q^g)^2 = 4\Delta/9b$. It follows immediately that $q^n - q^c = q^s - q^n$. Since $q^n = q^d(q^n)$, it also follows that $q^{dc} - q^n = (q^n - q^c)/2$.

payoffs are reduced, Hypothesis A should be true and Hypothesis B should be false, with more cooperative behavior observed in the alternative treatment.

The precise form of payoffs we utilize in the next section is linear-quadratic. For expositional clarity, we represent our experimental environment as a duopoly market for a homogeneous good with a linear demand curve and constant marginal costs. In our primary treatment, marginal costs are zero; then in the two other treatments, marginal costs are zero below some output, but positive for greater outputs. This positive cost is constructed to influence the best response under various conditions.

3. Experimental Design and Data

Subjects in our experimental design make choices from a payoff table. The row choice made by one subject becomes the column value of a counterpart. The intersection of the row and column in the payoff table shows earnings for the period. The payoff table represents the normal form of the stage game; this design is motivated by the game-theory model of oligopoly behavior (Friedman, 1983).[7]

We recruited subjects for each of the three experimental market structures from upper-level undergraduate economics classes. They reported to a reserved classroom with a personal computer at each seat. At the beginning of a session, instructions were read aloud as subjects followed along on their own copy. Questions were taken, and one practice period with a sample payoff table was held. In the practice period, a monitor randomly chose a column value while subjects, at the same time, chose a row value from the sample payoff table. Payoffs from the intersection of a row and the monitor's column were recorded by every subject. Each person was checked during the practice period to ensure that everyone understood payoff tables and that they were correctly recording their choices and earnings. Earnings were recorded in a fictitious currency called tokens. At the end of the experiment, tokens were exchanged for cash at the rate of $1.00 = 1000 tokens.

7. The use of payoff tables in experiments, however, has a history that predates their description of oligopoly markets. Rapoport et al. (1976) and Colman (1982), for instance, provide extensive surveys of literally hundreds of experiments that use payoff tables to generally learn more about rivalry and bargaining behavior. Surveys of how researchers have used payoff tables to study noncooperative behavior are provided by Davis and Holt (1993, Chapter 2), Friedman and Sunder (1994, Chapter 9), Kagel and Roth (1995), and Plott (1989).

In each market period subjects were instructed to type their row choices into the personal computer. The PCs were linked together and networked by the University's VAX cluster. Subjects were anonymously paired for the duration of the experiment, and paired individuals were not in proximity to each other. Once everyone had made a choice, the computer screen reported back to each subject his or her choices, earnings, and balance. Subjects wrote this information on a record sheet; they could always check the computer's calculations from the payoff tables provided to them. Subjects also were informed of the rival's choice and earnings. Finally, all participants knew that the experiment would last at least 35 periods. Subjects were told that at the end of period 35, and at the end of each period thereafter, the computer would randomly generate a number between 0 and 100, and that the experiment would end in the first period the random number did not exceed 20.[8] Sessions generally ended between period 35 and 40, and took a little longer than 1 hour. Earnings averaged about $20 per subject.

Each of the three experimental designs is a two-person repeated game, where the payoff table can be regarded as a reduced form that is derived from linear demand and cost conditions. In the unrestricted control design, which we shall refer to as "treatment 1" below, each agent faced fixed costs of $\frac{1300}{19}$, but no variable costs. The inverse demand function, for all treatments, was $P = \frac{150}{19} - 5(q_i + q_j)/76$. Every entry in the payoff table for the primary treatment, therefore, came from the payoff function

$$\pi_i = \left(\frac{150}{19} - \frac{5(q_i + q_j)}{76}\right)q_i - \frac{1300}{19}. \tag{5}$$

The function in equation (5) gives payoffs in cents; the payoff tables report payoffs in tokens.

In the payoff table seen by the subjects, quantity choices were rescaled to values between 1 and 22, where a choice of 1 corresponded to an output of 28, and so on. A copy of that payoff table is provided in the appendix as Table VI(A). It is labeled "no kink"; this

8. These experimental games are not stationary. However, the subgame-perfect equilibria that satisfy the incentive constraints for a stationary game also satisfy the incentive constraints in each period of nonstationary games such as these. There are additional subgame-perfect equilibria; these typically entail more collusive choices in early rounds, with play converging to the stationary equilibria by period 35. As we indicate below, our subjects' choices generally did not rise over time, so they evidently did not play these more exotic strategies. We address the concern that our nonstationary structure altered play in Section 5, where we analyze results from a treatment designed to eliminate the nonstationarities.

label was not on the tables provided to subjects. Subjects were never told they were picking outputs, just that they were choosing values from a payoff table where the intersection of their row choice and their rival's column choice determined their earnings from that choice period. With the parametrization used in our design, the Cournot-Nash equilibrium output is 40 for each player, while the symmetric joint-profit-maximizing output is 30. With the rescaling of output choices, the Cournot-Nash choice in the payoff table was 13, and the symmetric joint-profit-maximizing choice was 3.

Treatment 2 adjusted payoffs to make 13 the one-shot best response to any choice more cooperative than Cournot. Costs remained equal to $\frac{1300}{19}$ for outputs at or below 40 (i.e., choices at or below 13); for larger outputs, costs were adjusted to $\frac{1300}{19} + \frac{65}{38}(q_i - 40)$. Thus, the cost function can be represented as

$$C(q_i) = \begin{cases} \dfrac{1300}{19} & \text{for } q_i \leq 40, \\ \dfrac{65}{38} q_i & \text{for } q_i \geq 40. \end{cases} \tag{6}$$

This cost function contains a kink at the Cournot output of 40. A kink in the cost curve could simply represent a piecewise linear approximation to a nonlinear cost curve. In the context of oligopolists, it could reflect a surcharge on large output levels, as with a quota system, or it could reflect the extra cost of expanding capacity. It is straightforward to check that this function is continuous, that marginal costs change discretely from 0 to $\frac{65}{38}$ at an output of 40 (i.e., a choice of 13), and that the one-shot best response to any output between the symmetric collusive and Cournot values (i.e., between 30 and 40) becomes the Cournot output (40). We see from the corresponding table [Table VI(B)] that the optimal deviation is reduced to a choice of 13 for all choices between 3 and 13, and that deviation gains are not as large as they were in the control design. Thus, the kink in costs we have introduced in this treatment reduces deviation gains. This should increase cooperation between agents, if play is consistent with the use of a trigger strategy.

Treatment 3 altered the payoff matrix in a similar fashion. Here, costs remained equal to $\frac{1300}{19}$ for outputs at or below 42 (i.e., choices at or below 15); for larger outputs, costs were $\frac{65}{38}(q_i - 42) + \frac{1235}{19}$. The total cost function can be represented as

$$C(q_i) = \begin{cases} \dfrac{1300}{19} & \text{for } q_i \leq 42, \\ \dfrac{65}{38} q_i - \dfrac{65}{19} & \text{for } q_i \geq 42. \end{cases} \tag{7}$$

It is straightforward to check that this function is continuous, that marginal costs change discretely from 0 to $\frac{65}{38}$ at an output of 42, that the one-shot best response to any output between 36 and 40 (i.e., choice between 9 and 13 on the payoff table) is unaltered, and that the one-shot best response to any output between 30 and 36 becomes 42. This payoff table is given in Table VI(C) and is labeled as "kink at 15."

Suppose subjects are making symmetric choices of 10 in the payoff table (so that the typical pair chooses 20), which is very close to the estimated steady-state behavior in the primary treatment. The one-shot best response to this choice is 14 or 15. Accordingly, a change in payoffs beginning at choices greater than 15 should have no effect on play if behavior is consistent with a grim-reaper strategy, since the one-shot deviation gains are unaffected. However, play could be altered if subjects wanted to play choices larger than 15 to punish defection more severely. This alteration in the payoff table reduces the profits associated with punishment and therefore may support larger cooperation profits, i.e., more cooperation. Generally, if an equilibrium is based on more intense punishment than grim reaper, choices should be smaller under this alternative treatment than under the primary treatment design.[9]

Figure 1 summarizes subject behavior in the three treatments. We conducted two experimental sessions for each design. There were 14 market pairs in treatment 1, the primary treatment; 11 in treatment 2 (with costs kinked at a choice of 13); and 14 in treatment 3 (with costs kinked at a choice of 15). The solid line, marked AV1, shows average individual choices for treatment 1. The short-dashed line, marked AV2, shows average individual choices for treatment 2. For 70% of the periods, the plot AV2 lies below AV1. Hence, these data are supportive of our argument that reducing deviation gains will increase cooperation. The dashed line, AV3, shows average individual choices from treatment 3. This plot tends to lie above both AV1 and AV2 in early periods, but the differences dissipate after period 18. There is, however, substantial variation in the plots from period to period, so these differences may not be significant.

Selected summary statistics for our experiments are contained in Table I. Here we provide information on the number of observations, average choice, and standard error of the average choice for each treatment. There were a total of 1120 observations in treatment

9. Closer review of the payoff tables shows that putting a kink in costs at 15 will also affect the deviation gains if subjects are choosing values of 6 or less. At this level of cooperation, the alteration in the payoff table should cause subject pairs to make smaller choices. This would reinforce the effects discussed in Section 2 for trigger strategies with short-lived and intense punishment phases.

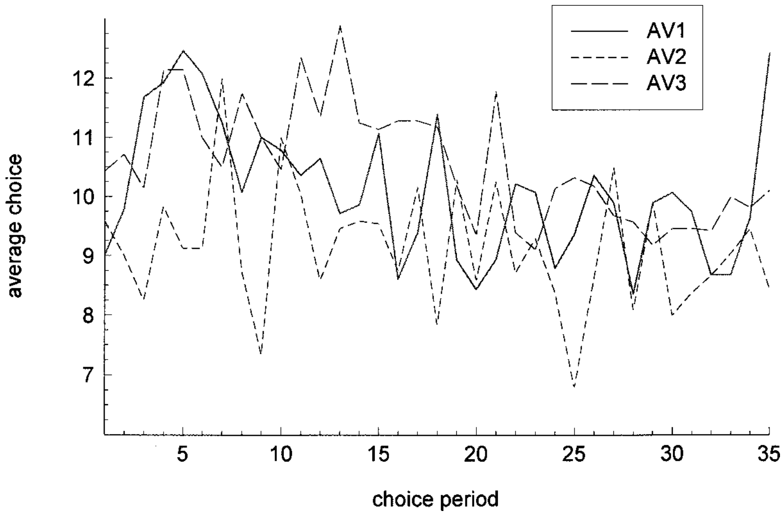

FIGURE 1. AVERAGE INDIVIDUAL CHOICES

1, 882 observations in treatment 2, and 1048 observations in treatment 3. These values are listed in column (1). The associated overall average individual choices are presented in column (4), with standard errors given in parentheses below. The immediate impression given by these statistics is that choices tend to be smallest in treatment 2, and that choices do not appear to differ much between treatments 1 and 3.

In addition to providing information on average choices for the entire session, we also are interested in subject behavior in periods where a subject observes an increase in the rival's choice. To this end,

TABLE I.
AGENT CHOICE PATTERNS

	Number of Observations			Mean Choice (Standard Error of Mean)		
Treatment	(1) All Choices	(2) Rival Increased	(3) Rival Did Not Increase	(4) All Choices	(5) Rival Increased	(6) Rival Did Not Increase
1 (no kink)	1120	358	762	10.382 (0.1903)	12.869 (0.3073)	9.214 (0.2278)
2 (kink at 13)	882	211	671	8.751 (0.1686)	8.801 (0.3437)	8.735 (0.1937)
3 (kink at 15)	1048	369	679	10.595 (0.1720)	12.201 (0.2551)	9.722 (0.2196)

we denote subject i's choice in period t by x_{it}. (Recall that choices are obtained from outputs by subtracting 27, i.e., $x_{it} = q_{it} - 27$.) Since subjects only observe lagged rival choices, subject i knows in any period t that the rival (subject j) has increased his or her choice when $x_{j\,t-1} > x_{j\,t-2}$. The number of observations where a subject has just observed an increase in the rival's choice is listed in column (2); the number of observations where the rival did not raise his or her choice is listed in column (3). In all three treatments, at least two-thirds of all observations did not follow an increase in the rival's choice, although in every treatment there are a substantial number of observations where an increase did occur. For these two subsamples, we present the average choice and the associated standard error.[10] These values are given in columns (5) and (6). In treatments 1 and 3, the average choice was markedly larger after the rival raised his or her choice, as we would expect if subjects retaliate when their rival increases his or her choice. On the other hand, there does not seem to be much difference in behavior between the two choice subsets for treatment 2. The similarity in average subject behavior when the rival had raised his or her choice and when no such increase had just occurred is somewhat puzzling, particularly in view of the fact that choices were noticeably smaller in treatment 2 than in the other two treatments. This feature led us to analyze the magnitude of rival increases and the frequency of punishments.

Table II(A) lists the number of observations where a subject's rival has raised his or her choice by S units, for values of S equal to $1, \ldots, 5$ and for $S \geq 6$. For each treatment, the total number of observations where the rival has increased his or her choice is given in parentheses following the treatment number. In part (B) we convert these observations into fractions. For example, there are 358 observations in treatment 1 where the rival raised his or her choice; of these, 61 correspond to an increase of 1 unit, which translates into a fraction 0.1704 (= 61/358). We note that the fraction of increases falls gradually as S rises from 1 to 5 in treatments 1 and 3, with less than one in five observations occurring at $S = 1$. By contrast, well over half the observed increases in the second treatment entailed an increase of just one unit. We believe it is conceivable that subjects tended to ignore such small increases, perhaps regarding them as trembles.

10. Let A represent the set of values of subject identity and period number where the subject observes the rival has just increased his or her choice, i.e., $A = \{(i,t): x_{j\,t-1} > x_{j\,t-2}$, and j is paired with i in the experiment$\}$. Let N_{RI} represent the number of observations in A, and N equal the total number of observations. Then the two averages we tabulate are $\bar{x}_{\text{RI}} = \sum_{(i,t)\in A} x_{it}/N_{\text{RI}}$ (average choice after the rival has just increased choice) and $\bar{x}_{\text{RNI}} = \sum_{(i,t)\notin A} x_{it}/(N - N_{\text{RI}})$ (average choice after the rival has not just increased choice). Standard errors are determined analogously.

TABLE II.
DESCRIPTIVE STATISTICS OF RIVAL INCREASES AND RETALIATIONS

(A)

Treatment	Number of Observations When Rival Increases By:					
(Total Obs.)	1	2	3	4	5	≥6
1 (358)	61	47	28	37	28	157
2 (211)	138	9	5	5	4	50
3 (369)	66	57	41	33	19	153

(B)

	Fraction of Observations When Rival Increases By:					
Treatment	1	2	3	4	5	≥6
1	0.1704	0.1313	0.0782	0.1034	0.0782	0.4385
2	0.6540	0.0427	0.0237	0.0237	0.0190	0.2370
3	0.1789	0.1545	0.1111	0.0894	0.0515	0.3604

(C)

Treatment	Number of Times Subject Retaliates When Rival Increases By:					
(Total Obs.)	1	2	3	4	5	≥6
1 (168)	23	18	14	20	14	79
2 (71)	47	3	3	2	1	15
3 (192)	32	30	14	18	14	84

(D)

	Fraction of Times Subject Retaliates When Rival Increases By:					
Treatment	1	2	3	4	5	≥6
1	0.3770	0.3830	0.5000	0.5405	0.5000	0.4929
2	0.3406	0.3333	0.6000	0.4000	0.2500	0.3061
3	0.4848	0.5263	0.3415	0.5455	0.7368	0.5564

To pursue this theme, we also tabulate the number of times an increase in the rival's choice between periods $t-2$ and $t-1$ triggered an increase in the subject's choice in period t.[11] This information is contained in Table II(C) and (D). As with the information in the top

11. These statistics include all observations where the rival increased his or her choice by S units. Some of these increases may represent attempts by subjects to explore the payoff possibilities, as opposed to deviations from an equilibrium regime. It is natural to ask if a sequence of stable choices followed by a defection by one subject leads to a punishment phase. Defining stability as choices for at least five consecutive periods that do not shift by more than two units, or a pattern of choices (like alternating row choices between 1 and 10) that is repeated for five or more periods, 20 of the 39 pairs exhibit stable behavior over some part of an experimental session. We view punishment as having occurred if a large choice is made in response to cheating on a previously

half, the total number of observations where the subject responded to an increase in the rival's choice by raising his or her choice is given in parentheses following the treatment number. In Table II(D), these raw data are transformed into percentages. Here, we give the ratio of the number of observations where a subject retaliated to the number of observations where the rival raised his or her choice by S units. For example, in the first treatment the retaliations to an increase of 1 represent a fraction of $\frac{23}{61}$, or 0.3770. The main point from these data is that subjects appeared more willing to retaliate in response to a larger increase than a smaller one. Broadly speaking, less than half of rival increases of 1 or 2 units drew retaliation, whereas over half of larger increases led to retaliation. While these numbers are not especially dramatic, we believe they are consistent with the view that subjects were more likely to ignore smaller increases. This argument explains the similarity of average choices in treatment 2 with and without increases in rival choice. Relative to the other treatments there are fewer deviations. Moreover, deviations in treatment 2 tend to be considerably smaller than in the other treatments, and subjects tend to ignore small changes when considering punishments.

4. Econometric Analysis of Data

We now turn to a rigorous analysis of choices. The goal is to estimate the equilibrium choice for a typical agent. We analyze the experimental data by treating each session as a pooled cross section time series.[12] In the sample, the cross-sectional element is given by the number of participants, with the number of observations per subject determined by the length of the session in which he or she participated.

stable pattern of choices. This is where one agent responds to deviation with a matching choice, or a Cournot or possibly larger response. Of the 39 subject pairs, 20 reach a stable choice pattern. There are six pairs that reach stability with no punishment phases and remain stable for the duration of the experiment. For seven other subject pairs there are instances of cheating and then a short-term punishment phase with some attempt to restore cooperation before the end of the experiment. For the remaining seven pairs we observe a punishment phase that appears permanent; subjects go to the Cournot equilibrium or a larger choice and stay at these levels. We hesitate to give much weight to this cataloging of behavior, since our definitions of stability and punishment are arbitrary. Nevertheless, choice patterns show that defection does generally lead to a punishment phase. In no case do we observe cheating on a stable choice pattern that goes without retaliation.

12. Experimental economists have often analyzed mean choices over a subset of the play in similar settings. Such an approach is inferior to our procedure because it neglects learning and dynamic adjustments, which can distort the results (Alger, 1987).

To model a subject's current choice we assume that period-t choices are based on period-$t-1$ choices for each agent, i.e., we assume subjects use Markov strategies. The spirit of a trigger strategy is that subjects respond more aggressively to output increases than to output reductions. Specifically, subjects who play in a manner consistent with the grim-reaper strategy ought to be disinclined to play cooperatively if the rival has raised his or her output in a past period. The literal interpretation of this scheme is that any output increase would trigger such a reaction. However, as discussed above, subjects might be willing to ignore small output increases—perhaps regarding them as trembles. Under this view, reversion would only occur if the increase exceeds some threshold level. Accordingly, we define a dummy variable D_{it} that equals 1 if the difference between i's rival's period-τ and period-$\tau - 1$ choices exceeds some nonnegative critical level S, and 0 otherwise, for any value of $\tau < t$. We then build a switching regression model (Cason and Mason, 1999) that allows for different linear Markov responses when $D_{it} = 1$ as compared to 0. Allowing for a disturbance term e_{it}, our switching regression model can then be written as[13]

$$x_{it} = (a_0 + b_0 x_{it-1} + c_0 x_{jt-1})(1 - D_{it}) + (a_1 + b_1 x_{it-1} + c_1 x_{jt-1}) D_{it} + e_{it}. \qquad (8)$$

In the interpretation we suggest above, the threshold value S is a free parameter to be estimated. In addition, we would expect the parenthetical component multiplied by D_{it} to be larger than the component multiplied by $1 - D_{it}$ if output increases trigger punishment responses.

An interpretation consistent with our model is that subject choices are converging to some long-run equilibrium, or steady-state level. Using this interpretation, we can develop a consistent estimate

13. The function in equation (8) may be interpreted as a dynamic reaction function or as a repeated game strategy. Alternatively, the relation can be motivated by such features as signaling and learning. Any attempts at signaling a desire to collude hinge on an intertemporal connection (Shapiro, 1980). Similarly, any learning implies a connection between current and preceding choices (Cason and Friedman, 1999; Cheung and Friedman, 1997; Mason and Phillips, 1997). This would be true for either adaptive learning or rational learning (Kalai and Lehrer, 1993; Mason and Phillips, 2001). An alternative is to use a weighted average of past choices by subject i's rivals. While Cason and Friedman (1999) find some support for this more elaborate model, Mason and Phillips (2001) provide evidence that the simpler one-lag specification provides a satisfactory characterization of agents' beliefs. Actually, the specific regression equation we employ is equivalent to that of equation (8), except that we break out the constant term from the first part and replace it with a series of dummy variables to capture individual-specific effects. This has no effect on the results we report below, since the tables list the average value of this intercept. We also analyzed a version of the switching regression where the switching dummy equaled one if the rival's choice increased between periods $t - 2$ and $t - 1$ (as opposed to any period $\tau \leq t - 1$). Those results, which are available upon request, are quite similar to the ones reported in the text.

of the equilibrium choices for a typical subject in each design. Consider a subject pair where both agents play this long-run equilibrium choice for several consecutive periods. In this case, $D_{it} = 0$. Then the deterministic version of (8) may be used to derive

$$x_m^e = \frac{a_{0m}}{1 - b_{0m} - c_{0m}}, \tag{9}$$

where a_{0m}, b_{0m}, and c_{0m} are the values of a_0, b_0, and c_0 for treatment m, where $m = 1$ for the primary treatment, $m = 2$ for the design where there is a kink in the cost function at a choice of 13, and $m = 3$ for the design where there is a kink in the cost function at a choice of 15. Since the function in equation (8) is continuous in the parameters, so long as b_{0m} and c_{0m} do not sum to one, Slutsky's theorem implies that x_m^e can be consistently estimated by inserting consistent estimators of the parameters a_{0m}, b_{0m}, and c_{0m} into the right side of equation (9). Covariance information from the maximum-likelihood parameter estimates can be used to construct consistent estimates of the covariance structure for the steady-state values x_m^e (Fomby et al., 1988, Corollary 4.2.2).

With the lagged structure in (8), we forfeit the first observation for each agent, leaving us with $T - 1$ observations for each individual (where T is the number of observations in the experimental session). This gives us 952 data points from the first treatment, 816 data points from the second treatment, and 952 data points from the third treatment. In estimating the parameters in equation (8), we used a fixed-effects approach. The main virtue of this approach is that it allows for idiosyncratic differences across individuals. In particular, the error component is modeled as $e_{it} = \eta_i + \varepsilon_t$, so that individual differences can be thought of as a shift in the intercept term. Under this interpretation, the estimated value for the intercept is the average effect across all subjects. With conventional assumptions, the fixed-effects approach yields consistent estimates of the parameters a_0, b_0, c_0, a_1, b_1, and c_1. These estimates, the implied steady-state choice, and the associated standard errors, R^2 values, and log-likelihood values are reported in Table III. In addition, we report Durbin's h-statistic (the appropriate test statistic for the hypothesis of serially uncorrelated errors in the presence of lagged explanatory variables); under the null hypothesis of no serial correlation this statistic follows a Student's t-distribution. All of this information is reported for each of the three designs.

The issue of interest has to do with the impact of changing payoffs on the estimated equilibrium choice. There are two hypotheses, from Section 2, that we wish to analyze. Placed in the context of our experimental treatments, Hypothesis A is that subjects' behavior is

TABLE III.
PARAMETER ESTIMATES UNDER DIFFERENT TREATMENTS WITH MAXIMUM-LIKELIHOOD SWITCHING CRITERIA

Parameter	Estimate	Standard Error
Treatment 1 (1120 observations, no kink)		
a_{01}	4.7324	0.4234
b_{01}	0.2766	0.0324
c_{01}	0.2647	0.0357
a_{11}	12.2392	1.2820
b_{11}	0.0978	0.0570
c_{11}	0.0563	0.0742
x_1^e	10.3186	0.7111
Maximum-likelihood estimate of S	2	
Log-likelihood function	−3378.42	
R^2	0.3970	
Durbin's h-statistic:	−0.6615	
Treatment 2 (882 observations, kink at 13)		
a_{02}	9.7453	0.4556
b_{02}	−0.0728	0.0344
c_{02}	−0.0420	0.0368
a_{12}	5.5823	2.3683
b_{12}	0.0821	0.0972
c_{12}	0.1699	0.1373
x_2^e	8.7417	0.1329
Maximum-likelihood estimate of S	3	
Log-likelihood function	−2367.27	
R^2	0.4980	
Durbin's h-statistic	0.5727	
Treatment 3 (1048 observations, kink at 15)		
a_{03}	6.3125	0.4680
b_{03}	0.0650	0.0316
c_{03}	0.3257	0.0347
a_{13}	13.4221	1.4814
b_{13}	−0.3505	0.0572
c_{13}	0.1659	0.0922
x_3^e	10.3599	0.4897
Maximum-likelihood estimate of S	3	
Log-likelihood function	−2983.74	
R^2	0.4381	
Durbin's h-statistic	−0.4223	

The t-statistic for $H_0 : x_1^e = x_2^e$ is 2.180, significant at better than 5% level; that for $H_0 : x_1^e = x_3^e$ is −0.049, insignificant at conventional levels.

not different between treatments 1 and 2, i.e., $x_1^e = x_2^e$. If the behavior is consistent with trigger-strategy play, choices should be more cooperative in treatment 2, where payoffs at outputs larger than Cournot are reduced, i.e., $x_1^e > x_2^e$. Hypothesis B is that subjects' behavior is not different between treatments 1 and 2, i.e., $x_1^e = x_3^e$. This hypothesis will be correct if subjects play in a manner that is consistent with a grim-reaper strategy. On the other hand, if behavior is consistent with a short-lived but intense punishment phase, or if play is influenced by such focal effects as the reduction in the range of choices with positive payoffs, then behavior should be more cooperative in treatment 3 than in the control treatment. Combining these possibilities, the alternative to Hypotheses B is $x_1^e > x_3^e$. Thus, both Hypotheses A and B are tested by means of a one-sided t-test.

Both hypotheses may be tested by comparing the t-statistic for the relevant difference ($x_1^e - x_2^e$ or $x_1^e - x_3^e$) against the critical point on a Student's t-distribution. These test statistics are reported at the bottom of Table III. The test statistic for Hypothesis A indicates that the null hypothesis is easily rejected at conventional levels. That is, reducing the deviation gains facilitates cooperation. The conclusion to be drawn is that subject behavior matches the predictions of trigger-strategy models and models that rely on subjects recognizing focal points. The test statistic for Hypothesis B indicates that the null hypothesis cannot be rejected. Hence the conclusion we draw from both of these tests is that subject behavior is not consistent with the use of a trigger strategy that involves a short-lived, intense punishment phase. It is also not well explained by a theory that suggests subjects become more cooperative following any change to payoff functions that expands the range of choices with negative payoffs. The strategy that is consistent with our observed results is the grim-reaper strategy, with its long-lived, less intense punishment phase.

In addition to shedding light on the two key hypotheses about trigger strategies that we outlined in Section 2, our econometric results allow us to discuss the reaction of a typical agent to an increased choice by the rival. If the trigger model is correct, we would expect to see larger choices by a subject following choice increases by the rival. These actions should affect the estimated coefficients in the switching model. In particular, we would expect that the choice implied by the model following a switch would exceed the choice implied with no switch. The easiest way to investigate this prediction is to compare the estimated steady state (which assumes no defections for the preceding two periods) against the response that would be triggered by a past increase in the rival's output that exceeds the critical value of S. For concreteness, we suppose that the rival defects by selecting the best response to player *i*'s (steady state) choice. In terms of the results in Table III, this increase by the rival would imply choices of 14.33, 13, and 14.32 for treatments 1, 2, and 3, respectively. (The associated steady-state choices are 10.32, 8.74 and 10.36, for treatments 1, 2, and 3, respectively.) The induced retaliation in treatments 1 and 3 (14.06 and 12.17, respectively) are significant, though the implied reaction in treatment 2 (8.87) does not differ significantly from the steady state. In addition, inspection of the estimated coefficients a_0 and a_1 reveals that defections are met with sharp responses in treatments 1 and 3, in that the estimated intercept is substantially larger following defection. This sort of response does not seem to occur in treatment 2 (although in fairness we should note that the estimated intercept following defection is dramatically noisier in that treatment). Instead, in that treatment it appears that subjects respond more gradually to defections, continually raising their choices toward the one-shot Nash equilibrium level.

Our econometric results can be used to assess the empirical importance of two alternative repeated-game strategies that have been proposed in the literature. The first, due to Stanford (1978), suggests that dynamic reaction functions will treat the rival's increases and decreases in the same way, and that these reactions are based solely on the rival's lagged choice (as opposed to the subject's own lagged choice). In the context of our regression model, the implied restriction is $b_0 = 0$, $b_1 = 0$, and $c_0 = c_1$. This restriction can be analyzed using a chi-squared test; since there are three parameter restrictions, the test statistic will follow a central chi-squared distribution with three degrees of freedom if the null hypothesis is true. In our application, the test statistics are 95.04, 6.82, and 52.8 for treatments 1, 2, and 3, respectively. The first and last of these test statistics is well in excess of the 5% critical value (7.82), indicating that we may reject the null hypothesis with considerable confidence. On the other hand, the test statistic is (barely) insignificant for treatment 2, indicating that we cannot reject the null hypothesis at the 5% level. One interpretation of these results would be that the Stanford strategy might explain behavior in the second treatment, where we observe the most cooperative-looking sessions.[14]

The second alternative strategy we consider is the so-called *tit-for-tat* strategy. Under tit-for-tat, each agent's period-t choice equals

14. We suggest this explanation for the sake of completeness. While it can be supported if one is militant about using a 5% confidence level, we note that the test statistic is significant at approximately the 7.8% level. Accordingly, we believe that the support for the Stanford strategy here is preliminary at best.

the rival's period-$t-1$ choice. In our econometric model, this corresponds to the parameter restrictions $b_0 = 0$, $b_1 = 0$, $c_0 = 1$, and $c_1 = 1$. With four parameter restrictions, the test statistic would follow a central chi-squared distribution with four degrees of freedom here. In our application, the test statistics are 587.38, 646.28, and 455.68 for treatment 1, 2, and 3, respectively, which are dramatically larger than the 5% cutoff value (9.49). We conclude that the subjects' behavior is inconsistent with tit-for-tat in our experiments.

5. Further Evidence

In this section, we consider two additional comparative effects to corroborate the results discussed in the preceding section. For concreteness, we refer to these as "variation 1" and "variation 2" in the pursuant discussion. In variation 1, we reduce the continuation probability for treatment 1 (no kink in the cost function). This change will reduce the range of cooperative choices that are consistent with either form of trigger strategy, so we would expect to see larger long-run choices in this design than in treatment 1 in the baseline experiments. Variation 2 is motivated by the observation that invoking a random endpoint in the game after period 35 creates a nonstationarity in the design. In principle, agents could exploit this nonstationarity by adjusting their choices over time.[15] To the extent that such behavior occurs, it may cloud the comparisons we sought to make above. To address these two concerns we analyze two further experimental designs.

5.1 Reducing the End-Period Probability

In variation 1, we reduce the continuation probability that is applied after period 35. It was explained to subjects that at the end of period 35, and at the end of every period thereafter, the computer would randomly choose a number between 1 and 100. If this number was bigger than 35, the game was played for the period; if 35 or less, the session immediately ended. Thus, the continuation probability is 0.65 in this new treatment (as compared to 0.80 in the primary treatment). All other aspects of this design are identical to the primary treatment discussed above.

15. In particular, the most cooperative output that can be selected in a given period rises slightly up to the period where the random stopping rule is invoked (Phillips and Mason, 1996). In the experimental sessions described in Section 4, subjects' choices tended to fall, not rise, over time. Nevertheless, we consider this variation for the sake of completeness.

TABLE IV.
PARAMETER ESTIMATES UNDER DIFFERENT CONTINUATION PROBABILITIES

Parameter	Estimate	Standard Error
Variation 1 (884 observations, $p = 0.65$)		
a_{04}	11.5617	1.6045
b_{04}	−0.0872	0.0613
c_{04}	0.1511	0.0846
a_{14}	7.8186	0.5776
b_{14}	0.1063	0.0367
c_{14}	0.2147	0.0414
x_4^e	11.5134	0.3743
Maximum-likelihood estimate of S	2	
Log-likelihood function	−2674.94	
R^2	0.3971	
Durbin's h-statistic	−0.4176	
Primary Treatment (1120 observations, $p = 0.8$)		
a_{01}	4.7324	0.4234
b_{01}	0.2766	0.0324
c_{01}	0.2647	0.0357
a_{11}	12.2392	1.2820
b_{11}	0.0978	0.0570
c_{11}	0.0563	0.0742
x_1^e	10.3186	0.7111
Maximum-likelihood estimate of S	2	
Log-likelihood function	−3378.42	
R^2	0.3970	
Durbin's h-statistic	−0.6615	

The t-statistic for $H_0 : x_1^e = x_4^e$ is 1.901, significant at better than 5% level.

We conducted two sessions of this variation; five subject pairs made choices for 39 periods in the first session, and seven subject pairs made choices for 37 periods in the second. Accordingly, we have 884 observations available here. Most subjects earned between \$15 and \$20, with sessions lasting about one hour.

Analysis of the data proceeds along the same lines as described above. In particular, we estimated the switching regression model described in equation (8), identified the maximum-likelihood value of S, and used the parameter estimates to calculate the steady-state choice using equation (9). These estimates, subscripted by 4 to distinguish them from earlier estimates, are presented in Table IV. For reference, we reproduce the estimates for the primary treatment (treatment 1) from Table III. The hypothesis of interest is that $x_1^e = x_4^e$; the obvious alternative hypothesis is that behavior is less cooperative with the smaller continuation probability, so that $x_1^e < x_4^e$. The results in Table IV indicate that x_4^e is larger than x_1^e, and that the t-statistic on this difference is statistically significant at better than the 5% level (using a one-tailed test). We conclude that lowering the continuation probability induced subjects to behave less cooperatively, in accordance with a binding incentive constraint in the trigger-strategy model.

5.2 A Stationary Design

As we noted above, subjects could have exploited the nonstationarities in the design reported in Section 4 by adjusting their behavior over time. While we do not believe our subjects exhibited such behavior, for the sake of completeness we consider an adaptation of the earlier basic design that constructs a stationary game. We do this in variation 2 by creating a moving exchange rate between tokens and cash. The exchange rate is set at $M_t = 200/(0.9)^{t-1}$, where t is the period of the experiment. Subjects were told that tokens would be exchanged for cash at the end of the experiment, and that the payments from each period depended on the exchange rate for that period, with M_t tokens worth \$1.00 for period t. The moving exchange rate is in force for the first 20 periods. Thereafter it is fixed (at M_{20}) and we invoke a random stopping rule with termination probability equal to 0.1 in each period.[16]

The purpose of this design is to create an environment that resembles an infinitely repeated game with a fixed discount factor. We are unaware of other experimental games that have combined a set δ with a random endpoint in this way, although Roth (1995) surveys many experimental bargaining studies that impose discounting by decreasing the size of the pie over successive periods. In most sequential bargaining experiments the endpoint is known, but some designs leave agents uninformed about the number of periods in a game. To assess any influences of the nonstationarity in the experimental environment discussed above upon our results, we repeat the

16. During the discounting phase, the payoffs in any period $t+1 \leq 20$ are 90% of the payoffs in period t. During the phase with probabilistic continuation, the expected values of payoffs in period $t+1$ are 90% of payoffs in period t, for $t \geq 20$ (as there is a 10% chance period $t+1$ will not be reached).

three treatments using this stationary environment. We conduct two sessions for each treatment, including 26 subjects in the treatment with no kink in the cost function, 24 subjects with kink at 13, and 24 subjects with kink at 15. Taking into account the lagged structure of our model, we have 780 observations from treatment 1, 668 observations from treatment 2, and 720 observations from treatment 3. The typical subject earned between \$20 and \$25 for a $1\frac{1}{2}$-hour session.

The econometric approach we take in analyzing data from these sessions is again based on the model described above, in equations (8) and (9). As in the earlier analysis, we identify the maximum-likelihood value of S as part of the regression analysis. Our estimation results are collected in Table V. The qualitative results reported in Table III are replicated: (1) subjects tend to react aggressively to substantial past increases in the rival's choice; (2) the average estimated equilibrium choice in the control treatment is significantly larger than when the cost curve is kinked at a choice of 13; and (3) the difference between behavior in the control treatment and behavior when costs are kinked at a choice of 15 is not statistically significant. Reducing the deviation gains encourages more cooperation, while reductions in payoffs above the one-period best response have no impact on behavior. Altogether, these observations are again consistent with subjects acting as if they employ a grim-reaper strategy.

6. Discussion

Our analysis of subject choices in our experimental two-person noncooperative games indicates that a systematic reduction in deviation gains leads to more cooperative outcomes. While this supports the hypothesis that play is consistent with use of trigger strategies, an alternative explanation is that such reductions make smaller choices more focal, so that the behavior becomes more cooperative. But this alternative explanation ought to apply equally well whether the kink in costs is imposed at a choice of 13 or at a choice of 15. Since there was no statistical difference between the behavior in the treatment with the kink at 15 and the behavior in the control treatment, we must reject this alternative explanation of subjects' behavior. Our econometric estimates provide robust evidence that behavior is more consistent with the predictions from the trigger-strategy model.

In a strict sense, our subjects do not play equilibrium strategies, since subjects rarely make the same choices for the entire game. Therefore, observed actions do not satisfy a literal interpretation of one

TABLE V.
PARAMETER ESTIMATES IN THE STATIONARY DESIGN

Parameter	Estimate	Standard Error
Treatment 1 (780 observations, no kink)		
a_{01}	7.4461	0.6271
b_{01}	0.1628	0.0371
c_{01}	0.2192	0.0418
a_{11}	11.0228	2.1521
b_{11}	−0.1774	0.0794
c_{11}	0.2106	0.1141
x_1^e	11.8388	0.3781
Maximum-likelihood estimate of S	5	
Log-likelihood function	−2310.70	
R^2	0.3584	
Durbin's h-statistic	−0.5437	
Treatment 2 (668 observations, kink at 13)		
a_{02}	8.7899	0.5541
b_{02}	−0.0800	0.0399
c_{02}	0.2599	0.0429
a_{12}	16.2934	1.9378
b_{12}	−0.3802	0.0917
c_{12}	−0.1420	0.1344
x_2^e	10.7175	0.3763
Maximum-likelihood estimate of S	6	
Log-likelihood function	−1719.13	
R^2	0.3272	
Durbin's h-statistic	−1.7421	
Treatment 3 (720 observations, kink at 15)		
a_{03}	4.0222	0.5038
b_{03}	0.2746	0.0344
c_{03}	0.3576	0.0389
a_{13}	18.8350	2.7209
b_{13}	0.0505	0.0923
c_{13}	−0.3734	0.1568
x_3^e	10.9378	1.5714
Maximum-likelihood estimate of S	6	
Log-likelihood function	−1993.46	
R^2	0.4836	
Durbin's h-statistic	0.0736	

The t-statistic for $H_0 : x_1^e = x_2^e$ is 1.820, significant at better than 5% level; that for $H_0 : x_1^e = x_3^e$ is 0.657, insignificant at conventional levels.

equilibrium supported by a trigger strategy. Our point is not that subjects use trigger strategies *per se*, but that their play is *consistent* with the use of such strategies and can be explained by trigger-strategy models. There is some indication, nevertheless, that many subjects do form crude versions of a trigger strategy. The raw data indicate this, as described in Tables I and II. Also, after the game is over and subjects have been paid, informal interviews indicate that agents punish cheaters by selecting big numbers. The described purpose of the punishment is to discourage cheating and/or encourage other agents to cooperate. The length of a punishment phase, however, seems to be a vague notion in most people's minds. Subjects may prefer a short but more intense phase, but after defection it is more difficult to signal intentions because an element of trust is lost. Hence, *de facto* the punishment phase goes on indefinitely.

Frequently, in the above discussion we have described agents in the games as two identical firms making output choices in a homogeneous product market. Deviation gains were reduced in the same way a capacity constraint in the technology might be imposed. At high outputs the cost of production increases. We point out that the capacity constraint is not binding in the static sense. In all of our laboratory markets, agents were producing less than the output at which the cost kink existed. Nevertheless, our results show that these off-equilibrium changes do affect the equilibrium outcome. What one might label a weak capacity restriction is, therefore, a collusive influence in the market. It is widely recognized that binding capacity restraints help rivals to restrict output and keep prices high. The evidence in this paper shows that the constraint need not be binding in order to facilitate the collusive effort. In a broader sense, any number of weak regulatory controls could have a collusive influence on behavior. For example, nonbinding quota restrictions in international trade or pollution charges that generate extra costs for output levels larger than those firms currently produce may promote anticompetitive behavior in markets.

Appendix

The payoff tables described in Section 3 are presented in Table VI.

TABLE VI.
PAYOFF TABLES FOR SUBJECTS

(A) No Kink
Horizontal Axis: Value of Strategy Selected by Other Participant
Vertical Axis: Value of Strategy Selected by Subject

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
1	495	476	458	439	421	403	384	366	347	329	311	292	274	255	237	218	200	182	163	145	126	108
2	518	499	480	461	441	422	403	384	365	346	327	308	289	270	251	232	213	193	174	155	136	117
3	539	520	500	480	461	441	421	401	382	362	342	322	303	283	263	243	224	204	184	164	145	125
4	560	539	519	499	478	458	438	417	397	376	356	336	315	295	274	254	234	213	193	172	152	132
5	579	558	537	516	495	474	453	432	411	389	368	347	326	305	284	263	242	221	200	179	158	137
6	597	575	553	532	510	488	466	445	423	401	380	358	336	314	293	271	249	228	206	184	163	141
7	613	591	568	546	524	501	479	457	434	412	389	367	345	322	300	78	255	233	211	188	166	143
8	628	605	582	559	536	513	490	467	444	421	398	375	352	329	306	282	260	237	214	191	168	145
9	642	618	595	571	547	524	500	476	453	429	405	382	358	334	311	287	262	239	216	192	168	145
10	655	630	606	580	557	533	509	484	460	436	411	387	363	338	314	289	265	241	216	191	166	141
11	666	641	616	591	566	541	516	491	466	441	416	391	366	341	316	291	266	241	216	191	166	141
12	676	650	624	599	573	547	522	496	470	445	419	393	368	342	316	291	265	239	214	188	163	137
13	684	658	632	605	579	553	526	500	474	447	421	395	368	342	316	289	263	237	211	184	158	132
14	691	664	638	611	584	557	530	503	476	449	422	395	368	341	314	287	260	233	206	179	152	125
15	697	670	642	614	587	559	532	504	476	449	421	393	366	338	311	283	255	228	200	172	145	117
16	702	674	645	617	589	561	532	504	476	447	419	391	363	334	306	278	249	221	193	164	136	108
17	705	676	647	618	589	561	532	503	474	445	416	387	358	329	300	271	242	413	184	155	126	97
18	707	678	648	618	589	559	530	500	470	441	411	382	352	322	293	263	234	204	174	145	115	86
19	708	678	647	617	587	557	526	496	466	436	405	375	345	314	284	254	224	193	163	133	103	72
20	707	676	645	614	584	553	522	491	460	429	398	367	336	305	274	243	213	182	151	120	89	58
21	705	674	642	611	579	547	516	484	453	421	389	358	326	295	263	232	200	168	137	105	74	42
22	702	670	638	605	573	541	509	476	444	412	380	347	315	283	251	218	186	154	122	89	57	25

Continued

(B) Kink at 13
Horizontal Axis: Value of Strategy Selected by Other Participant
Vertical Axis: Value of Strategy Selected by Subject

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
1	495	476	458	439	421	403	384	366	347	329	311	292	274	255	237	218	200	182	163	145	126	108
2	518	499	480	461	441	422	403	384	365	346	327	308	289	270	251	232	213	193	174	155	136	117
3	539	520	500	480	461	441	421	401	382	362	342	322	303	283	263	243	224	204	184	164	145	125
4	560	539	519	499	478	458	438	417	397	376	356	336	315	295	274	254	234	213	193	172	152	132
5	579	558	537	516	495	474	453	432	411	389	368	347	326	305	284	263	242	221	200	179	158	137
6	597	575	553	532	510	488	466	445	423	401	380	358	336	314	293	271	249	228	206	184	163	141
7	613	591	568	546	524	501	479	457	434	412	389	367	345	322	300	78	255	233	211	188	166	143
8	628	605	582	559	536	513	490	467	444	421	398	375	352	329	306	282	260	237	214	191	168	145
9	642	618	595	571	547	524	500	476	453	429	405	382	358	334	311	287	262	239	216	192	168	145
10	655	630	606	580	557	533	509	484	460	436	411	387	363	338	314	289	265	241	216	191	166	141
11	666	641	616	591	566	541	516	491	466	441	416	391	366	341	316	291	266	241	216	191	166	141
12	676	650	624	599	573	547	522	496	470	445	419	393	368	342	316	291	265	239	214	188	163	137
13	684	658	632	605	579	553	526	500	474	447	421	395	368	342	316	289	263	237	211	184	158	132
14	674	647	620	593	566	539	513	486	459	432	405	378	351	324	297	270	243	216	189	162	135	101
15	663	636	608	580	553	525	497	470	442	414	387	359	332	304	276	249	221	193	166	138	111	83
16	651	622	594	566	538	509	481	453	424	396	368	339	311	283	255	226	198	170	141	113	85	57
17	637	608	579	550	521	492	463	434	405	376	347	318	289	261	232	203	174	145	116	87	58	29
18	622	592	563	533	503	474	444	414	385	355	326	293	266	237	207	178	148	118	89	59	30	0
19	605	575	545	514	484	454	424	393	363	333	303	272	242	212	182	151	121	91	61	30	0	−30
20	588	557	527	495	464	433	402	371	340	309	278	247	216	186	155	124	93	62	31	0	−31	−62
21	568	537	505	474	442	411	379	347	316	284	253	221	189	158	126	95	63	32	0	−32	−63	−95
22	548	516	484	451	419	387	355	322	290	258	226	193	161	129	97	64	32	0	−32	−64	−97	−129

Continued

(C) Kink at 15
Horizontal Axis: Value of Strategy Selected by Other Participant
Vertical Axis: Value of Strategy Selected by Subject

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
1	495	476	458	439	421	403	384	366	347	329	311	292	274	255	237	218	200	182	163	145	126	108
2	518	499	480	461	441	422	403	384	365	346	327	308	289	270	251	232	213	193	174	155	136	117
3	539	520	500	480	461	441	421	401	382	362	342	322	303	283	263	243	224	204	184	164	145	125
4	560	539	519	499	478	458	438	417	397	376	356	336	315	295	274	254	234	213	193	172	152	132
5	579	558	537	516	495	474	453	432	411	389	368	347	326	305	284	263	242	221	200	179	158	137
6	597	575	553	532	510	488	466	445	423	401	380	358	336	314	293	271	249	228	206	184	163	141
7	613	591	568	546	524	501	479	457	434	412	389	367	345	322	300	78	255	233	211	188	166	143
8	628	605	582	559	536	513	490	467	444	421	398	375	352	329	306	282	260	237	214	191	168	145
9	642	618	595	571	547	524	500	476	453	429	405	382	358	334	311	287	262	239	216	192	168	145
10	655	630	606	580	557	533	509	484	460	436	411	387	363	338	314	289	265	241	216	191	166	141
11	666	641	616	591	566	541	516	491	466	441	416	391	366	341	316	291	266	241	216	191	166	141
12	676	650	624	599	573	547	522	496	470	445	419	393	368	342	316	291	265	239	214	188	163	137
13	684	658	632	605	579	553	526	500	474	447	421	395	368	342	316	289	263	237	211	184	158	132
14	691	664	638	611	584	557	530	503	476	449	422	395	368	341	314	287	260	233	206	179	152	125
15	697	670	642	614	587	559	532	504	476	449	421	393	366	338	311	283	255	228	200	172	145	117
16	685	657	628	600	572	543	515	587	459	430	402	374	345	317	289	261	232	204	176	147	119	91
17	571	642	613	584	555	526	497	468	439	411	382	353	324	295	266	237	208	179	150	121	92	63
18	656	626	597	567	537	508	478	449	419	389	360	330	301	271	241	212	182	153	123	93	64	34
19	639	609	579	549	518	488	458	428	397	367	337	307	276	246	216	186	155	125	95	64	34	4
20	622	591	560	529	498	467	436	405	374	343	312	282	251	220	189	158	127	96	65	34	3	−28
21	603	571	539	508	476	445	413	382	350	318	287	255	224	192	161	129	97	66	34	3	−29	−61
22	582	550	518	486	453	421	289	357	324	292	260	228	195	163	131	99	66	34	2	−30	−62	−95

References

Abreu, D., 1986, "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, 39(1), 191–225.

——, 1988, "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56, 383–396.

Akerlof, G. and J. Yellen, eds., 1986, *Efficiency Wage Models of the Labor Market*, Cambridge, England: Cambridge University Press.

Alger, D., 1987, "Laboratory Tests of Equilibrium Predictions with Disequilibrium Data," *Review of Economic Studies*, 54, 105–145.

Barro, R. and D. Gordon, 1983, "Rules, Discretion, and Reputation in a Model of Monetary Policy," *Journal of Monetary Policy*, 12, 101–121.

Bernheim, B.D. and M.D. Whinston, 1990, "Multimarket Contact and Collusive Behavior," *RAND Journal of Economics*, 21, 1–26.

Cason, T.N. and D. Friedman, 1999, "Learning in a Laboratory Market with Random Supply and Demand," *Experimental Economics*, 2, 77–98.

—— and C.F. Mason, 1999, "Uncertainty, Information Sharing and Collusion in Laboratory Duopoly Markets," *Economic Inquiry*, 37, 258–281.

Cheung, Y.W. and D. Friedman, 1997, "Individual Learning in Games: Some Laboratory Results," *Games and Economic Behavior*, 19, 46–76.

Colman, A., 1982, *Game Theory and Experimental Games*, New York: Pergamon Press.

Davis, D. and C. Holt, 1993, *Experimental Economics*, Princeton, NJ: Princeton University Press.

Ellison, G., 1994, "Theories of Cartel Stability and the Joint Executive Committee," *RAND Journal of Economics*, 25, 37–57.

Folmer, H., P.V. Mouche, and S. Ragland, 1993, "Interconnected Games and International Environmental Problems," *Environmental and Resource Economics*, 3, 313–335.

Fomby, T., R. Hill, and S. Johnson, 1988, *Advanced Econometric Methods*, New York: Springer-Verlag.

Friedman, D. and S. Sunder, 1994, *Experimental Methods*, New York: Cambridge University Press.

Friedman, J., 1967, "An Experimental Study of Cooperative Duopoly," *Econometrica*, 35(3–4), 379–397.

——, 1983, *Oligopoly Theory*, New York: Cambridge University Press.

Fudenberg, D. and J. Tirole, 1989, "Non-cooperative Game Theory for Industrial Organization: An Introduction and Overview," in R. Schmalensee and R. Willig, eds., *Handbook of Industrial Organization*, Vol. I, Amsterdam: Elsevier Science.

Green, E. and R. Porter, 1984, "Non-cooperative Collusion under Imperfect Information," *Econometrica*, 52, 87–100.

Grossman, H.I. and J.B. Van Huyck, 1988, "Sovereign Debt as a Contingent Claim: Excusable Default, Repudiation, and Reputation," *American Economic Review*, 78, 1088–1097.

Haltiwanger, J. and J. Harrington, 1991, "The Impact of Cyclical Demand on Collusive Behavior," *RAND Journal of Economics*, 22, 89–106.

Kagel, J.H. and A.E. Roth, 1995, *The Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press.

Kalai, E. and E. Lehrer, 1993, "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61, 1019–1045.

Levinthal, D. and J. March, 1993, "The Myopia of Learning," *Strategic Management Journal*, 14, 95–112.

Marx, L.M. and S. Matthews, 2000, "Dynamic Voluntary Contribution to a Public Project," *Review of Economic Studies*, 67, 327–358.

Mason, C.F. and O.R. Phillips, 1997, "Information and Cost Asymmetry in Experimental Duopoly Markets," *The Review of Economics and Statistics*, 79, 290–299.

——and ——, 2001, "Dynamic Learning in a Two-Person Experimental Game," *Journal of Economic Dynamics and Control*, 25, 1305–1344.

Palfrey, T. and H. Rosenthal, 1994, "Repeated Play, Cooperation, and Coordination: An Experimental Study," *Review of Economic Studies*, 61, 545–565.

Phillips, O.R. and C.F. Mason, 1996, "Market Regulation and Multimarket Rivalry," *RAND Journal of Economics*, 27, 596–617.

Plott, C., 1989, "An Updated Review of Industrial Organization: Applications of Experimental Methods," in R. Schmalensee and R. Willig, eds., *Handbook of Industrial Organization*, Vol. II, Amsterdam: Elsevier Science.

Porter, R., 1983, "A Study in Cartel Stability," *Bell Journal of Economics*, 14, 301–314.

Prendergast, C., 1993, "The Role of Promotion in Inducing Specific Human Capital Acquisition," *Quarterly Journal of Economics*, 108, 523–534.

Rapoport, A., M. Guyer, and D. Gordon, 1976, *The* 2×2 *Game*, Ann Arbor, MI: University of Michigan Press.

Rasmusen, E., 1994, *Games and Information: An Introduction to Game Theory*, 2nd ed., Cambridge, MA: Basil Blackwell.

Rotemberg, J. and G. Saloner, 1986, "A Supergame-Theoretic Model of Price Wars during Booms," *American Economic Review*, 76, 390–407.

Roth, A.E., 1995, "Bargaining Experiments," in J.H. Kagel and A.E. Roth, eds., *The Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press.

Salant, D.J. and G.A. Woroch, 1992, "Trigger Price Regulation," *RAND Journal of Economics*, 23, 29–51.

Segerstrom, P.S., 1988, "Demons and Repentance," *Journal of Economic Theory*, 45, 32–52.

Sell, J. and R. Wilson, 1996, "Can Grim Trigger Strategies Lead to Grim Cooperation?" Working Paper, Rice University.

Shapiro, C. and J. Stiglitz, 1984, "Equilibrium Unemployment as a Discipline Device," *American Economic Review*, 74, 433–444.

Shapiro, L., 1980, "Decentralized Dynamics in Duopoly with Pareto Optimal Outcomes," *Bell Journal of Economics*, 11(2), 730–744.

Stanford, W.G., 1986, "On Continuous Reaction Function Equilibria in Duopoly Supergames with Mean Payoffs," *Journal of Economic Theory*, 39(1), 233–250.