Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Organizational Behavior and Human Decision Processes 100 (2006) 60–75
ORGANIZATIONAL BEHAVIOR AND HUMAN DECISION PROCESSES
www.elsevier.com/locate/obhdp

Modeling alliance activity: Opportunity cost effects and manipulations in an iterated prisoner's dilemma with exit option

Darryl A. Seale *, Richard J. Arend, Steven Phelan

Department of Management, College of Business, University of Nevada Las Vegas, 4505 Maryland Parkway, Las Vegas, NV 89154-6009, USA

Received 17 September 2004
Available online 28 February 2006

Abstract

We model the two-firm alliance as an iterated prisoners' dilemma game with an exit option and test several theoretical predictions over two experimental studies. A new major effect on alliance performance arises by including the exit option (i.e., the option to end the alliance and receive a fixed payoff that is less than the payoff for mutual cooperation but greater than the payoff for mutual defection). The opportunity cost levels of the firms either directly or indirectly influence alliance cooperation and alliance payoffs through affecting the alliance strategies that the firms choose. Implications for partner selection, alliance selection and structuring, and strategy choices along the alliance lifespan, build on these results.

Keywords: Prisoners' dilemma; Alliance; Exit option

With more than 10,000 partnerships created annually (Schifrin, 2001), and growth rates projected between 25 and 35% per annum (Bleeke & Ernst, 1995; Harbinson & Pekar, 1998; Pekar & Allio, 1994), strategic alliances represent one of the fastest growing forms of corporate strategy and governance (Das & Teng, 2000). The popularity and growth of strategic alliances, however, does not signify that alliances are any more successful than alternative governance mechanisms like mergers, acquisitions or internal ventures (De Rond & Hamid, 2004). Major consulting firms have estimated failure rates from 50 (McKinsey and Company) to 64% (Pricewaterhouse Coopers) for US ventures, and as high as 70% (Boston Consulting Group) for international alliances (see De Rond, 2003, for a more complete discussion of consulting firms' estimates of alliance failure rates), with failure often attributed to opportunistic behavior by one or both partners (e.g., Parkhe, 1993; Zeng & Chen, 2003). Given the rapid rate of alliance growth and the alarming likelihood of failure, proper modeling of alliance behavior has been a vital topic of organizational research. The continued study of alliances is important for several reasons, including the increasing popularity of the alliance as a form of governance, the relatively high failure rate of alliances, and the multiple lenses that can be used to understand alliance performance.

The main contribution of the paper is in examining how the addition of an exit option stage to the standard prisoners' dilemma (PD) game can be used to model the dynamics of business relationships. Specifically, we test whether the addition of the exit option to the iterated prisoners' dilemma (IPD) game produces the results predicted by the theoretical analysis of Arend and Seale (2005) in their modeling of alliances where sophisticated players understand the interdependent and voluntary nature of the relationship. We use an experimental method to provide a fuller understanding of the

* Corresponding author. Fax: +1 702 895 4370.
E-mail address: dseale@unlv.nevada.edu (D.A. Seale).

0749-5978/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.obhdp.2006.01.002

processes and contingencies of the decisions of real players involved in alliance relationships.

We choose to model the alliance as an iterated prisoners' dilemma game with exit option (IPDEO) for several reasons. First, each stage of an alliance resembles a PD—an interdependent relationship where partners agree to provide certain inputs (which are usually imperfectly enforceable) and for a share of the output, but where individual incentives do not align with joint incentives. Second, most alliances involve iterations with the same partner—i.e., repeated cycles of investment and outcome towards some long-term goal. Third, partners can exit from alliances at any stage; this is a basic tenant of free markets that relationships are voluntary, especially flexible ones like alliances. Fourth, there is support in the literature for each of these assumptions; we are not the first to consider alliances as resembling PDs, IPDs, or IPDEOs.

In the organizational literature, an alliance is typically described as an intermediate form of governance—a method of resource acquisition between the spot market and full ownership (Gulati, 1998; Mitchell & Singh, 1996). Although researchers adopt somewhat different definitions, depending on the focus and purpose of study, they seem to agree on four necessary criteria for defining an alliance accord: (1) the arrangement involves two or more firms that combine their resources outside of the spot and corporate control markets in order to accomplish mutual objectives; (2) there is repeated interaction with the same partner; (3) returns from the alliance are co-dependent—the amount each firm earns depends on the decisions and actions of the partner; and, (4) the arrangement between the firms is voluntary—each firm has the opportunity to terminate the alliance at a future stage. We interpret these criteria in the following manner: Firms remain separate entities throughout the process; the alliance partnership is flexible, includes several separable stages, and is unilaterally terminable; and payoffs in each alliance stage are interdependent, with potential synergies involved when partners cooperate.

Given this characterization, we propose that two-firm alliances are amenable to game theoretic modeling. Several studies (most notably Parkhe, 1993; Parkhe, Rosenthal, & Chandran, 1993; Zeng & Chen, 2003) have used the PD metaphor to explicitly model alliance formations. Parkhe (1993, pp. 794–796) contends, "…*alliances are voluntary interfirm cooperative agreements, often characterized by inherent instability arising from uncertainty regarding a partner's future behavior and the absence of a higher authority to ensure compliance*." And that this uncertainty and the innate incentive to cheat on the arrangement, directly affecting alliance payoffs, can create "…*a situation isomorphous to the game of prisoner's dilemma*." In their nationwide survey of senior executives involved in strategic alliances, Parkhe et al. (1993) conclude that the mixed motive payoff structure, consistent with PD-type payoffs, is a prominent feature of many actual business alliances. Similarly, Zeng and Chen (2003) argue that alliance partners face a cooperative dilemma in their attempt to balance cooperation and competition.

Although these studies adopt a PD metaphor in examining structural factors of alliances, two recent studies (Arend & Seale, 2005; Phelan, Arend, & Seale, 2005) extend this description to repeated or iterated PD games. The IPD has become a popular instrument to study the emergence of cooperative behavior among non-altruistic decision makers. The preponderance of empirical evidence suggests that mutual cooperation tends to emerge when the number of repetitions of the stage game is uncertain, and cooperation in early stages of the game and the likelihood of future interactions with the current partner are sufficiently high (Ashlock, Smucker, Stanley, & Tesfatsion, 1996).

Arend and Seale propose an IPDEO as a potential model of alliance formation. Given certain ordinal relationships between the payoff variables for mutual cooperation and defection, one-sided cooperation and defection, and exiting the game, and given the assumption that each partner has an incentive to remain in the alliance, they solve for the probability of cooperation using basic optimization. Other predictions of incorporating the exit option into PD games are less clear. On the one hand, it is argued that an option to exit the IPD game reduces the dependence players have on each other. The exit option increases the incentive for aggressive players to adopt hit-and-run behaviors (i.e., players choose to defect in round t, then avoid retaliation by exiting the game in round $t + 1$). On the other hand, it is argued that an option to exit an IPD gives more cooperative players an opportunity to flee rather than fight (Boone & Macy, 1999) and may work as a threat, tacitly enforcing higher levels of cooperation.

Besides the theoretical and experimental research that relates to the model adopted here, there are computer simulation and strategy tournament papers that are relevant. The paper by Phelan et al. (2005) extends the IPDEO model of alliance behavior by studying programmed strategies submitted to a round-robin computer tournament. Manipulating both game length and payoff for exiting the game, their results shed light on how one might best respond to an opportunistic partner. Using an evolutionary-style IPDEO tournament, similar to that used by Axelrod (1984), Yamagishi and Hayashi (1996) report that an *out-for-tat* strategy outperformed the popular *tit-for-tat* play. Further, their study predicts higher levels of cooperation in games with exit options than in games without. Using similar methodology, Congleton & Vanberg's (2001) results support the notion that an exit option increases cooperation. Boone and Macy (1999), however, find no significant differences in

IPD games when subjects were allowed to exit and search for other partners.

We complement the Arend & Seale theory paper and the Phelan, Arend & Seale tournament paper by adopting the same IPDEO model and applying it to the study of how real subjects decide how to play the alliance game (versus how fully rational game-theoretic agents do or how sophisticated programs do). In doing so we also complement the IPD experimental literature by analyzing what the exit option does to established results. We use university business students in two sets of experiments that include several variations on the model to test seven related hypotheses. We find that the exit option has significant effects on the IPD results, effects that are often contingent on the belief a player has in a partner's future cooperation. These experimental results are consistent with the theoretical analysis when accounting for the contingencies, but are less consistent with the tournament results. The analysis not only sheds new light on the translation of business-related games to more accurate application where the exit option is appropriate, it also provides several practical implications for managers contemplating alliances. The bases for many of these implications were obtained through several variations on the model to determine how the exit option affects choice of alliance, choice of partner, and structuring the alliance.

The rest of the paper is organized as follows: The next section describes background and rationale for the development of the hypotheses that predict behavior and outcomes in the IPDEO game. The third section describes the experimental methods of Study 1. This is followed by a discussion of the results for Study 1. Alliance choice—the effect of the payoff multiple introduces Studies 2a and 2b, which were conducted to better understand and explain the previous results. Alliance choice—the effect of entry fees describes the results from these two follow-up studies. The final section concludes with a discussion of results, limitations, implications, and suggestions for further research.

Development of hypotheses

The IPDEO model

We model alliance activity as an IPDEO game; players interact repeatedly with the same partner and both have an option to exit at the beginning of each alliance stage. That model is formally established in a separate theoretical paper by Arend and Seale (2005). Essentially, the model is based on three premises established in the literature. The first is that many interdependent relationships among separate parties, including alliances, represent PD-like games (e.g., Axelrod, 1984; Cable & Shane, 1997; Orbell & Dawes, 1993; Parkhe, 1993; Parkhe et al., 1993). The second is that, regardless, the PD metaphor fails to capture an important aspect of many interdependent decision situations—that people can often choose to *not play* (Orbell & Dawes, 1993). The addition of an exit option to the standard strategy choices provides a better model of most business relationships, such as alliances; a fact that Orbell and Dawes (1993, p. 787) note when they observe that decision makers spend considerable time and effort in "*maneuvering out of bad games with unpromising players and into good games with promising players*." The third is that an alliance is a multi-stage relationship that can be separated into discrete steps where commitments and rewards are made and garnered.

Formally, we assume the standard (symmetric) PD payoff inequalities hold in each alliance stage that is entered into: $T > R > P > S$, and $2R > T + S$ (Axelrod, 1984), where T is a firm's payoff for unilateral defection, R is a firm's payoff from mutual cooperation, P is a firm's payoff from mutual defection, and S is a firm's payoff from unilateral cooperation. The standard strategies of cooperate (C) and defect (D) are available; we then introduce an exit option (E), which we model as giving each firm a payoff of its opportunity cost (o) when played. Thus, each stage of the IPDEO game is represented by the following pattern of payoffs:

Firm 1	Firm 2		
	C	D	E
C	*R,R*	*S,T*	*o,o*
D	*T,S*	*P,P*	*o,o*
E	*o,o*	*o,o*	*o,o*

We adopt the usual assumption that the opportunity cost exceeds the payoff for mutual defection but is less than the payoff for mutual cooperation (i.e., $T > R > o > P > S$). This is a standard ordering for PD games with exit options (e.g., Orbell & Dawes, 1993; Tullock, 1985; Vanberg & Congleton, 1992). $R > o$ ensures that firms enter the alliance rationally expecting to benefit if both parties cooperate. $o > P$ makes the game non-trivial; it makes the partners have some risk to their relative wealth when the alliance is based on each firm trying to defect on the other—a situation where it is not unlikely that net losses should occur. The compressed normal form of the game (the 3×3 matrix above) is a convenient representation of the 2-stage game that participants play; where the first stage involves the choice of exiting or playing the PD, and the second stage involves the PD game contingent on both partners having chosen not to exit.

We now consider the expected outcome of such a game; first per iteration, and then across iterations. To do so, we consider the extensive form of the game—where one iteration involves the two stages—and work back from the second stage to the first using the concept of backward induction (i.e., starting at the last stage and working back to the first stage). The second stage is simply the PD. The dominant strategy of defect (i.e., D) is chosen. When we then move to the first stage, however, the players need to consider the probability of the other player choosing to cooperate in the second stage. If any player believes that the partner is of sufficient rationality to know that defection dominates the PD game and that the partner will not tremble in its choice (i.e., choose C by mistake) then exit (i.e., E) is the best strategy in the first stage and, hence, the one-iteration game.

This process is essentially successive elimination of weakly dominated strategies. We acknowledge that this solution concept has been criticized in the game theory literature (e.g., Binmore, 1999). However, the primary focus of the present research is not to determine how players arrive at equilibria, but rather to determine how the exit option affects alliance behavior. Similarly, we also acknowledge that backward induction is a concept that brings with it some additional conditions on players above that required for a PD game or supergame. We tend to agree with Aumann (1995) that simple rationality—i.e., where each player is a utility maximizer—is insufficient for backward induction to hold; common knowledge of rationality is also often required. The purpose of our research, however, is not to enter this debate, or to determine whether players hold beliefs of common knowledge of rationality. We assume that real decision makers involved in alliances, as well as their proxies in this experiment—the business students—fulfill the requirements of the beliefs in common knowledge of rationality.

For the iterated game, we follow the IPD experimental results in expecting a range of outcomes. When the last iteration is known (i.e., the end of the alliance is predictable) then we would expect that an argument similar to that used in PD supergames to apply—i.e., using backward induction the last iteration would entail the same outcome as the one-iteration game and then unravel to the first iteration with the same result, which here is exit. When the last iteration is not known, then we would expect that the arguments of IPD experiments that attempt to explain relatively high levels of cooperation to apply (Axelrod, 1980; Knez & Camerer, 2000; Neyman, 1999; Schuessler, 1989). Players may be able to sustain cooperation in the PD stage (and thus not exit) because defection would trigger exit or defection payoffs that may be unattractive relative to future cooperative payoffs.

Given that these ordinal relationships between the payoff variables hold, Arend and Seale (2005) propose a solution to the IPDEO generated through basic optimization. The optimization technique derives a set of equations for each player that reflect the interdependent payoffs from *cooperate* or *defect* decisions, as well as the opportunity cost (i.e., exit payoff). The solution maximizes the individual payoffs under the constraints that each player has an incentive to remain in the alliance (i.e., the solution is a result of simultaneous optimization of the players' payoff maximizations under linear constraints; the objective is to satisfy the first and second order conditions, given the constraints). The solution represents the level of cooperation (i.e., denoted by p^*—the firm's probability of playing C) that must be maintained for partners to remain in the alliance. Note that this solution is *not* a Nash equilibrium where partners playing at this level have no incentive to deviate. Rather, the solution level delineates both the best mix of C and D for a firm to use, as well as the minimum level of belief the player has to have in the partner's playing C in the alliance stage (i.e., denoted by b^*—the solution belief probability of the partner playing C). When that belief is violated—when the player does not expect the partner to play C with a high enough probability—then the player would have to consider less cooperative strategies of play to optimize her own payoff. This belief level (b^*), that is a function of opportunity costs, will play a major role in the hypotheses and results that follow.

In developing out hypotheses we begin by considering what effect opportunity cost has on alliance cooperation. That is the one exogenous variable a manager has to consider that has not been already covered in the literature (e.g., exogenous payoff level effects have been covered by several authors, including Rapoport & Chammah, 1965). There are many potential endogenous variables—levers that managers can pull—that can also affect alliance cooperation levels in the IPDEO. We consider three main types: we consider one variable relevant to partner choice; we consider three variables relevant to alliance choice; and, we consider two variables relevant to playing the alliance.

The new exogenous variable—The opportunity cost effect

The optimization solution predicts that when the players' beliefs are high enough then increasing the exit payoff, o, within the range $R > o > P$, increases the probability of cooperation. Although no one would expect naive subjects to be able to compute the solution's threshold levels of cooperation (p^*), it is clear from the payoff table that as o approaches R, a higher level of cooperation is required from both parties for the mean earnings to equal or exceed that of the fixed opportunity cost. Thus, a rational player would tend to cooperate more at a higher opportunity cost level given that player chose to enter and stay in the alliance, contingent on beliefs (b) that the partner will also cooperate more. The higher opportunity cost acts as a de facto entry fee that commits an entering player, through forward induction, to more likely cooperate as that level increases (see Cachon & Camerer, 1996 as an example how such an entry fee increases cooperation in a coordination game) given the belief that the same logic is being used by the partner. The first hypothesis follows:

Hypothesis 1. The level of cooperation will increase with opportunity cost level, contingent upon the required belief levels (b^*) being met.

Partner choice—The effect of partner opportunity cost asymmetry

The one main relevant firm characteristic that a manager can use to select an alliance partner is the partner's opportunity cost level. We have hypothesized that opportunity cost levels affect cooperation in our model. We have considered the effect as exogenous from the focal firm's perspective in the first hypothesis, and we now consider the effect as endogenous by assuming that the focal firm can choose its partner based on the partner's opportunity cost level.

Assume that there is an opportunity cost, o^* ($P < o^* < R$), where $b^* = b$—i.e., where the solution belief level (which is the level of belief and level of cooperative play required to generate expected payoffs at the opportunity cost) equals the realized belief level. Below o^* then $b^* < b$, and above o^* then $b^* > b$. Our arguments for H1 imply that cooperation increases with o until o^*, and then decreases. That would suggest then that the focal firm choose a partner with as close to the optimal opportunity cost level as possible in order to maximize cooperation.

Let's consider the general case of asymmetry to understand whether intuition works: we will assume one firm with a higher cooperation level under symmetry than the other (i.e., one firm cooperates more in a symmetric alliance than the other, due to their difference in opportunity cost levels). There are three cases to consider: (1) where both firms have opportunity costs below o^*; (2) where both firms have opportunity costs above o^*; and (3), where one firm is below and the other is above o^*. In case (1), required beliefs are being exceeded (i.e., $b > b^*$), so we would predict (based on the model solution) that asymmetry increases (decreases) cooperation received (i.e., the amount of cooperative actions made by the partner) for the lower (higher) opportunity cost firm. In case (2), required beliefs are not being met, so we would predict that asymmetry decreases (increases) cooperation received for the lower (higher) opportunity cost firm (*n.b.*, this is a less clean prediction because we are considering off-equilibrium paths, and essentially less-than-fully-rational moves). In case (3), only one required belief is being met, so our prediction is more complicated. Unless the high opportunity cost firm has a close-to-optimal opportunity cost, then asymmetry decreases (increases) cooperation received for the lower (higher) opportunity cost firm. In all cases we predict that asymmetry helps (hurts) the firm that would have a lower (higher) cooperation level under symmetry, contingent on beliefs. The second hypothesis follows:

Hypothesis 2. Partner asymmetry in opportunity costs increases cooperation received for the player that would have had a lower cooperation level under symmetry (contingent on beliefs), and decreases cooperation received for the partner.

Alliance choice—The effect of the shadow of the future

A manager can select an alliance based on several characteristics. One characteristic is whether the endpoint of the alliance has been specified or left unspecified—i.e., whether the final alliance stage is known or unknown. It is established in the IPD literature that knowing the endpoint of the iterated game should lead to decreased cooperation (e.g., Axelrod, 1984). One explanation for this theoretical prediction lies in the backward induction logic from game theory (e.g., Binmore, 1992). The experimental results, however, are less clear as variance in the levels of cooperation observed in repeated play are considerable (for example, see Fehr & Gachter, 2000; Fehr & Schmidt, 1999; Rapoport & Chammah, 1965; Yung-An & Day-Yang, 2003). Cooperation is maintained when decision makers see their immediate, short-term gain from defecting as not worth the possible sacrifice of future, long-term gains from mutual cooperation. Players' expectations of reciprocal behavior can form a bond between future earnings and present actions (Parkhe, 1993). This bond is termed the *shadow of the future*. Theorists argue that a crucial element is a determinate end point of the interaction. According to Telser (1980, p. 34), "*Self-enforcing agreements are not feasible if the sequence of occasions of transactions has a definite known last element.*" When the last element is unknown, the Folk Theorem (Fudenberg & Masken, 1986; Rubinstein, 1979) predicts that, in repeated games with no exit option, any individually rational outcome is possible because a threat of a trigger penalty to repetitions of the lowest equilibrium outcome can be used to support any better outcome. Although repeated games that have exit options do not meet the specific requirements of the theorem, the spirit of the argument, that a trigger to penalize partners to endure a low payoff level (e.g., possibly through exit), can possibly support more jointly attractive outcomes. Given theoretical predictions of defection when the end of the game is known, we propose the third hypothesis:

Hypothesis 3. Knowing when the alliance will end decreases the level of cooperation.

Alliance choice—The effect of the payoff multiple

The second characteristic we consider that a manager can use for selection criteria is the alliance's absolute payoff levels. We consider the effect of changing the absolute payoff levels while retaining the relative levels. We do this by multiplying all payoffs, including the opportunity costs. We do this to determine whether*loss aversion*—a decision bias—has an effect on alliance tactics. There are many other payoff manipulations that are available and appear applicable to the IPDEO, but most of these have been covered in previous research on PDs and IPDs (e.g., Rapoport & Chammah, 1965). For example, we would expect that increasing the payoffs to cooperating (R and S) and decreasing the payoffs from defecting (T and P) would increase cooperation levels, as has been established in the literature.

We bring something more interesting and less established to the study. We retain the relative payoff structure (i.e., T, R, o, P, and S all shift equally) while altering the absolute payoff values. We do so in a different way than the previous literature and in a way that is realistic for the context—alliances. The previous literature has used arithmetic shifts of payoffs to alter the perceived decision frame (e.g., McCusker & Carnevale, 1995); we use a multiplier to do the same. Essentially, we are allowing the manager to choose one alliance versus another (or effectively combine or divide similar alliance stages of a given alliance—i.e., which could be done through altering the timing of irreversible commitments). The change in frame allows us to investigate whether an established decision bias can be used to help optimize alliance-level choices.

The loss-avoidance principle (e.g., Kahneman & Tversky, 1979; Mussweiler, 2001; Roxburgh, 2003; Slovic & Lichtenstein, 1973; Tversky & Kahneman, 1986) implies that if people frame their decision using the exit payoff as a reference point, they risk a greater loss in the high exit payoff conditions than in the low exit payoff conditions. The principle predicts that people pick (and expect others to pick) strategies that may result in a gain and avoid (and expect others to avoid) strategies that may result in a loss. Thus, if losses are more salient than gains, players may be more inclined to exit the alliance or play less cooperatively, or both, when losses are increased the same amount as gains in an absolute or even a relative way. This leads to the following hypothesis:

Hypothesis 4. Multiplying (increasing) the payoff values in the alliance decreases the level of cooperation.

Alliance choice—The effect of entry fees

The third characteristic we consider that a manager can use for selection is the entry fee to participate (e.g., up front costs in due diligence, search, and feasibility). We consider how a shift in the frame from using an entry fee affects cooperation. Here, we assume an absolutely significant yet relatively low entry fee to be consistent with the spirit of the context (i.e., potential alliance benefits summed over the multiple stages of the alliance are about a magnitude larger than the entry fee). Similar theoretical arguments regarding loss aversion, as noted above, can be applied here to argue that the result should be lower levels of cooperation. We shift the frame down by the amount of the entry fee, making the evaluation of any potential loss relatively worse. Again, if losses are more salient than gains, players may be more inclined to exit the alliance or play less cooperatively, or both to reduce the chances of realizing the larger potential losses associated with unilateral cooperation. We expect that the fee will have a greater impact on the high opportunity cost players because a mutual cooperation payoff has little effect on covering the fee versus obtaining a unilateral defection payoff.

Hypothesis 5. Imposing an entry fee in the alliance reduces the level of cooperation. This reduction should be greater when opportunity costs increase.

Alliance play—The strategy effects

We now change our focus slightly from trying to increase the cooperation level of the alliance to trying more directly to increase the payoffs to the manager's firm. Instead of looking at choices the manager can make to increase the firm's payoffs through increasing the alliance cooperation level, we look at how the manager chooses to play a given alliance (i.e., after the partner and alliance choices have been made).

Although we have a solution to the IPDEO model, it is based on a certain level of belief the player has in her partner's likelihood of cooperation, and it requires a somewhat sophisticated level of calculation. For managers that may not have those beliefs or mathematical prowess, the optimal strategy to play in an IPDEO is not simple. Even the normally robust tit-for-tat (TFT) strategy may not be an obvious contender. In a population of strategies in a simulated IPDEO game, TFT can be invaded by a defect-and-exit (D&E) strategy (Arend & Seale, 2005). In the IPDEO, strategies that are nice, retaliatory, forgiving and clear—four important characteristics of TFT (Axelrod, 1984; Hofstadter, 1985)—are not obviously superior. Being nice is exploitable by hit-and-run type strategies, such as D&E. Retaliation is mitigated when defectors can opt-out. Forgiveness of defection is not clearly a good tactic given opting out may offer better returns. And, clarity of intensions, especially when advertising niceness and forgiveness may actually work to attract more partners that wish to exploit the PD-type payoffs than partners that wish to enjoy mutual benefits (Hofstadter, 1985).

As outlined in the explanation of the first hypotheses, the optimal strategy to play (i.e., whether to choose C, D or E) is dependent on the firm's beliefs about its partner's probability of playing cooperatively. When that belief differs from the optimization level (i.e., $b \neq b^*$), then so does the strategy that is recommended. Specifically, when the belief is relatively low, then the best strategy is to stay out of the alliance completely (i.e., play E at the first stage). When the belief is slightly higher, then the best strategy is to defect and then exit immediately after (i.e., play D and then play E) given the defection will lower subsequent beliefs. At a higher belief level, it may be worthwhile to play a mix of D and C, but it is not a rationally supportable equilibrium because defecting and exiting should produce a higher payoff. At a higher belief level, where the solution belief level occurs, the best strategy is to play the belief-level of cooperation (i.e., play C and D in a ratio to fulfill the belief). Above that belief, if it can be sustained, more cooperation increases the payoffs as long as the partner matches. This reliance on beliefs has a substantial effect on the cooperation level for the alliance. For example, when the actual belief lies between a given low and a given high opportunity cost level optimal beliefs, then the optimal strategy is defect-and-exit (or exit) for the high opportunity cost firm (i.e., where $b < b^*$) but cooperate for the low opportunity cost firm (i.e., where $b > b^*$). What that means is a possible reversal of the first hypothesis (H1); see Fig. 1 for a graphical representation.

What this analysis means to the manager is the following: Assume that there is little a manager can do to change the partner's beliefs (upward) about the manager's probability of cooperating in the short-run, given each manager is reasonably rational (in the game theoretic sense). Thus, strategies are not chosen to send signals, but instead to optimize medium-run individual payoffs. To do so means that firms with opportunity cost solution belief levels below the expected belief level should cooperate and stay in the alliance to perform better, while firms with opportunity cost solution belief levels above the expected belief level should cooperate less and exit sooner in the face of defection to perform better. The sixth hypothesis follows:

Hypothesis 6. Optimal strategies vary with the opportunity cost for a given alliance and belief level. For players with opportunity cost solution belief levels (b^*) below (above) the belief level (b), more (less) cooperative and more tolerant (less tolerant) strategies will perform relatively better.

Alliance play—Adapting along the alliance stages

On a more detailed level, managers can choose to play the alliance in a way that exploits the learning that takes place as the alliance unfolds. Given that strategies are

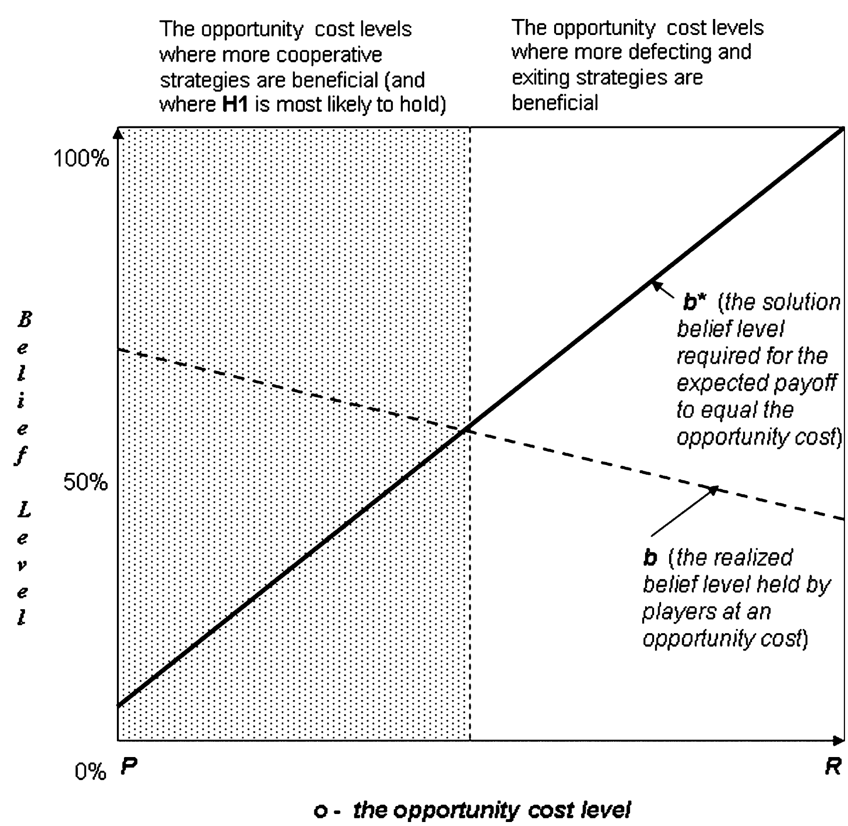

Fig. 1. The implications of opportunity costs and belief levels.

belief-dependent, the managers with the more accurate beliefs should reap better profits for their firms. However, given that partner cooperation levels are likely to change over the term of the alliance and as firms become more experienced from partaking in multiple alliances, accuracy will be dependent on how well managers learn the changes and adapt strategies to them. Accuracy entails several issues, including: (1) making a good estimate of what the initial cooperative level should be for a new alliance; and, (2) understanding the trends of responses to strategic interaction over time in the alliance.

To the first issue, a good estimate of initial cooperation levels is likely to be one that is optimistic relative to the average level for two main reasons. The first reason is that decision-makers have been observed to have an initial niceness bias. One characteristic of successful IPD strategies established in the experimental and simulation literatures (e.g., Axelrod, 1984) is *niceness*—the characteristic of not being the first to defect. There are several rationales for such behavior being observed in subject play. For example, social conditioning, perceived reputational benefits, the importance of first impressions, and fairness all seem to support the initially more cooperative play. Even the con game approach of sharking implies initial cooperative play (to build up trust before exploiting the gullible target). The second reason is that in a finite PD-based interaction, the longer the game continues, the more likely the next stage will be the last—the stage where defection is the only rational strategy.

To the second issue, understanding trends will entail understanding that partners will: likely respond to both gains and losses, doing so with different weights; be myopic; anchor on average payoff; be sensitive to alliance length; etc… That understanding will be indirectly observable by the adjustments in the beliefs, and then in the adjustments in strategies. In fact, the adjustment in beliefs is a signal of the rationality required to correctly adjust strategy. As such, we would expect that players that are more accurate in their beliefs will choose a better strategy and obtain a better payoff. The final hypothesis follows:

Hypothesis 7. More accurate beliefs generate greater performance.

Experimental method—Study 1

Study 1 was designed to measure the effects of: opportunity cost levels (H1); and, knowledge of the alliance's final stage (H3) on levels of cooperation.

Subjects

Eighty subjects, predominantly undergraduate business students enrolled in several senior-level Current Issues in Management or Business Policy/Strategy courses, participated in the experiment. The members of each class were randomly assigned to one of four experimental conditions described below. Subjects were motivated to make thoughtful decisions by basing part of their course grade on how many points they earned during the course of the experiment. In addition, between three and six subjects in each of the four twenty-person groups were randomly selected at the end of each session and paid, in cash, based on their performance. On average these subjects earned $15.25. Most students were between the ages of twenty-one and thirty-five, with an approximate balance between males and females.

Procedure

The experiment was conducted at a computer lab, at a major southwestern university. The lab contained over 40 networked computer stations separated by wide aisles and back-to-back placement. As subjects entered the lab they were randomly assigned to a workstation that was preset to play one of four experimental conditions, were told to read the instructions (see Appendix A for a copy of instructions) that were placed beside each station, and asked not to communicate with one another. The instructions indicated that they would play a simple two-person game over a computer network. The game, which was designed with prisoner's dilemma-type payoffs and included an exit option, was introduced as an *Alliance Formation* exercise. The instructions also indicated that during the course of the experiment subjects would be randomly paired with several different players for a number of repeated trials (stage games).

At the beginning of each trial subjects were required to make one of three decisions: *cooperate*, *compete*, or *opt-out*. Their earnings for a *cooperate* or *compete* decision depended on the decision of the person they were paired with. For example, if both players chose to cooperate (compete), each earned 8 (3) francs for the trial, where francs are an experimental currency that are exchanged for US dollars at the end of the experiment. If one player chose to cooperate (compete) and their partner chose to compete (cooperate), the player earned 0 (10) francs for the trial, while their partner earned 10 (0). If either player chose to opt-out, both players earned the exit payoff, which was either 4 or 7 francs, for the current trial as well as for each trial that the players remained paired. The exit payoff is analogous to opportunity cost in an alliance model—it represents the next best alternative to not joining or not continuing the alliance partnership. The instructions contained the following payoff table to aid subject decision-making.

Your decision	Your partner's decision	
	Cooperate	Compete
Cooperate	8, 8	0, 10
Compete	10, 0	3, 3
Opt-out	4, 4	4, 4

We chose the two exit payoff levels to cover both ends of the spectrum with one near the lower bound, *P*, and the other near the upper bound, *R*. Given the payoff parameters *T*, *R*, *P*, and *S*, when $o = 7$, $p^* = b^* = 0.828$, and when $o = 4$, $p^* = b^* = 0.236$. These values also nicely bracketed the coin-flip belief assessment of partner cooperation of 0.5, meaning that we would likely observe strategic play outside the solution's predictions.

At the end of each trial the computer informed subjects of their partner's decision, and reported both the cumulative and current earnings for the trial. At the end of the fourth pairing, the computer informed subjects that the experiment had ended and provided a re-cap of their cumulative earnings. Several subjects we're randomly selected and paid in cash their cumulative earnings according to the exchange rate 1 franc = $0.05. Following this random selection and payment, subjects were debriefed and dismissed from the lab.

To summarize, Study 1 consisted of a 2 × 2 design (exit payoff × knowledge of length of pairing). In the high (low) exit payoff treatment each subject earned 7 (4) francs per trial following an *opt-out* decision by either the player or her partner (i.e., both players received their opportunity cost as payment for each alliance stage from exit to the last stage of the alliance). Subjects in the full knowledge treatment *knew* exactly how long each paring would last. Subjects in the partial knowledge treatment were *uncertain* as to how long the alliance pairings would last. Consistent with the way unknown game lengths are handled in other studies (e.g., Axelrod, 1984), we told the subjects that there was a significant positive probability of switching alliance partners on the next trial. Regardless of treatment, players were informed each time their pairing changed, and were randomly paired with four different players during the course of the experiment, which lasted for 50 trials.

Results—Study 1

The results from the Study 1 are summarized in Table 1. The level of alliance cooperation, measured as a proportion of cooperative play, is analyzed in an opportunity cost by knowledge of length of pairing (2 × 2) ANOVA. This analysis reveals significant main effects for both exit payoff ($F_{1,79} = 83.88$, $p < 0.01$) and knowledge ($F_{1,79} = 7.57$, $p < 0.01$), and no interaction. However, contrary to the predictions of Arend and Seale's optimization model, the level of cooperation for the high opportunity cost condition was significantly less than for the low opportunity cost condition (i.e., 27.3% versus 70.3%). Thus, H1—cooperation would vary with opportunity cost—was not supported. The result for H3 is clearer; the level of cooperation in alliances with unknown length was

Table 1
Hypotheses tests

Hypothesis		Type of test	Test result	Comment
H1	Level of cooperation increases with opportunity cost (contingent on beliefs)	Means test (*F* test in ANOVA)	$F_{1,79} = 83.88$, $p < .01$	Supported in opposite direction (due to *violation* of belief contingency)
H2	Partner asymmetry increases (decreases) cooperation received for the player that would have had a lower (higher) cooperation level under symmetry	Means tests (*t* test assuming unequal variances)	$t_{118} = 11.63$, $p < .01$ for low opportunity cost players; $t_{116} = 5.71$, $p < 0.01$ for high	Supported
H3	Knowing when the alliance will end decreases cooperation	Means test (*F* test in ANOVA)	$F_{1,79} = 7.57$, $p < .01$	Supported
H4	Multiplying (increasing) the payoff values in the alliance decreases cooperation	Means tests (*t* test assuming unequal variances)	$t_{46} = 3.31$, $p < .01$ for low opportunity cost players; $t_{43} = 2.61$, $p < .05$ for high	Supported
H5	Entry fees decrease alliance cooperation	Means tests (*t* test assuming unequal variances)	$t_{28} = 0.35$, *ns* for low opportunity cost players; $t_{42} = 2.10$, $p < .05$ for high	Partially supported
H6	Optimal strategies vary with the opportunity costs. More cooperative strategies perform better for the low opportunity cost players here	Means tests (*t* test assuming unequal variances)	$t_{36} = 6.53$, $p < .01$ for low opportunity cost players; $t_{38} = -0.80$, *ns* for high	Partially supported
H7	More accurate beliefs generate better performance	Means test (*t* test assuming unequal variances)	$t_{69} = 3.62$, $p < .01$	Supported

Note. ns denotes not significant.

significantly higher than when length was known (i.e., 67.1% versus 51.2%).

Given the unexpected result for H1, we set out to better understand the beliefs that subjects held concerning their partner's expected level of cooperative play, and rationale for exiting the alliance game. The two studies described below add to this understanding.

Experimental method—Studies 2a and 2b

Study 2a

Study 2a was designed to measure the effects of: opportunity cost levels contingent on beliefs (H1); partner opportunity cost asymmetry (H2); strategy choices contingent on beliefs (H6); and, accuracy of beliefs (H7). Study 2a was also designed to collect process data (i.e., subjects' beliefs of partner cooperation) as subjects considered their decisions to cooperate, compete or opt-out. The outcome variables of player cooperation, partner cooperation, and alliance earnings were also recorded.

The population of subjects and procedural methods were identical to Study 1, except: (1) subjects changed alliance partners more frequently—during the course of Study 2a subjects faced ten different alliance partners during (at least) fifty trials of play; (2) before making their decisions, subjects were required to estimate the likelihood that the player they were paired with would cooperate on the current trial by moving a slider button on a scale from 0 to 100% probability of cooperation; (3) data was collected on two groups of twenty players in each of three conditions; and (4) in all conditions subjects were uncertain as to the length of the alliance pairing.

Thus, Study 2a consisted of three conditions: in one condition players were symmetric with an exit payoff of 7 francs; in a second condition players were symmetric with an exit payoff of 4 francs; and in the final condition players were asymmetric—half of the players earned 7 francs as their exit payoff, while the other half earned 4 francs. In the asymmetric condition, players earning 7 (4) francs were always paired with players earning 4 (7) francs.

Study 2b

Study 2b was designed to measure the effects of: multiplying the basic payoffs (H4); and, imposing an entry fee (H5). As well, Study 2b collected additional process data by focusing on why subjects exited alliances.

Study 2b was analogous to the previous two studies, except that the treatments were examined in a within-subjects design over the different opportunity cost levels (i.e., each subject played a treatment, playing half the stages as a low opportunity cost firm and half as a high opportunity cost firm, where the assignment of which half occurred first was random). There were $n = 75$ players, split evenly among three treatments: $n = 25$ played the base game with an unknown end; $n = 25$ played the game with doubled payoffs; and, $n = 25$ played the game with an entry fee of 10 francs imposed at the start of each new alliance. These treatments were properly balanced and their order controlled so that each could be examined with appropriate statistical techniques.

The multiple of two was used to test H4 because it represented the lowest integer multiple that provided a significant absolute change to the payoffs. Intuitively, it would mean that two stages were combined in the alliance.

The entry fee level of 10 francs was chosen to test H5 because it is seemingly large in a per-round sense (i.e., it equals the T payoff) yet somewhat less important when considering the payoff from the alliance as a whole. Stylized facts suggest that up-front costs for alliances and other major transactions lie in that range.

The exit rationales available for subjects to choose when they did exit an alliance were generated from the wrap-up surveys taken by a random sample of subjects from Study 1. The four separate exit rationales were: (1) the partner defected too much; (2) the partner was going to defect; (3) I met my aspiration level for this partnership; and (4) I was not meeting my aspiration level for this partnership.

Results—Studies 2a and 2b

In Study 2a, consistent with Study 1, symmetric players with low opportunity costs cooperated more than symmetric players with high opportunity costs, with a significant main effect for opportunity cost level ($F_{1,119} = 23.79, p < .01$). In Study 2a, low opportunity cost players cooperated at a 65.4% level, held beliefs of partner cooperation at a 64.0% level (with standard deviation 31.1%) and had a b^* of 23.6% whereas high opportunity cost players cooperated at a 31.3% level, held beliefs of partner cooperation at a 45.1% level (with standard deviation 31.3%), and had a b^* of 82.8%. Whereas the belief requirements for a significant cooperation level was met for the low opportunity cost players, the belief requirements for cooperation to occur for the high opportunity cost players was not met (i.e., in only about 10% of the cases could it be met). The result for H1 is thus explained by the violation of the belief level contingency for the high opportunity cost players.

Study 2a also resolved the asymmetry effect. Low opportunity cost players went from receiving 65.4% cooperation under symmetric partnerships to receiving 37.7% cooperation under asymmetric partnerships; high opportunity cost players went in the opposite direction, from 31.3 to 46.3% received cooperation. The means tests ($t_{118} = 11.63$, $t_{116} = 5.71$; $p < .01$ for each) were significant for each change, providing support for H2.

The belief information and the payoff results from Study 2a provided the basis for testing H6 and H7—the strategy-related hypotheses. All low opportunity cost players held average beliefs above the solution belief level (mean player average belief level was 64% versus the solution belief level of 24%). Those who played cooperatively above the mean level (i.e., 66%) experienced significantly better performance than those who played below (means test statistic $t_{36} = 6.53$; $p < .01$). Almost all high opportunity cost players held average beliefs below the solution belief level (mean player average belief level was 45% versus the solution belief level of 83%). Those who played defection or exit above the mean level (i.e., 70%) did not experience better performance than those who played below (means test statistic $t_{38} = -0.80$; *ns*). Thus, the results only partially supported H6.

We collected all symmetric pairing results from Study 2a and calculated for each player how accurately the realized partner strategy was predicted from the beliefs of that stage's strategy. Accuracy ranged from 13 to 100% with a median of 73%. We then compared the payoff performance of those in the top half of accuracy with those in the bottom half using a means test and determined that more accurate beliefs generated better performance ($t_{69} = 3.62$; $p < .01$). Thus, H7 was supported.

Study 2b provided the data to test H4 and H5—two alliance-related hypotheses. We computed the level of cooperation in alliances with the doubled payoffs and compared to the base level of cooperation. We found that multiplying the payoffs significantly decreased cooperation for both the high and low opportunity cost players by roughly the same relative amount (i.e., cooperation decreased 30.6% of the base rate of cooperation in each case). This was a statistically significant change in the means of level of cooperation for both opportunity cost levels (i.e., for low opportunity cost players, $t_{46} = 3.31$ with $p < .01$; for high opportunity cost players, $t_{43} = 2.61$ with $p < .05$). Thus, H4 was supported. We then conducted a similar test for the entry fee treatment, comparing the mean level of cooperation to the base case. We found that the entry fee only statistically significantly reduced the cooperation level for the high opportunity cost players ($t_{42} = 2.10$; $p < .05$), reducing by about 30% from the base case. For the low opportunity cost players, the entry fee had essentially no effect ($t_{28} = 0.35$; *ns*). Thus, only partial support for H5 was found.

Discussion and conclusions

The primary purpose of this paper was to examine how the addition of an exit option stage to the standard IPD game can be used to model the dynamics of certain business relationships. Using the model proposed by Arend and Seale (2005), we developed then tested seven hypotheses concerning alliance behavior. H1—*cooperation will increase with opportunity cost level*, was seemingly supported in the opposite direction—the low opportunity cost alliances realized greater cooperation than the high opportunity cost alliances—but the results also determined that the necessary contingency on belief levels for H1 was violated. The high opportunity cost players did not believe that their partners would cooperate sufficiently to ensure an average alliance payoff at least as attractive as the opportunity cost, so they tended to defect and exit rather than cooperate. The low opportunity cost players more than believed that their partners would cooperate sufficiently to ensure an average alliance payoff at least as attractive as the opportunity cost, so they tended to cooperate and stay in rather than defect or exit. Essentially, when beliefs are at off-the-solution-path levels the first hypothesis has to be recast to: *Alliances involving players with opportunity cost solution belief levels significantly below (above) the belief level realize lower cooperation levels than those involving players with opportunity cost solution belief levels significantly above (below) the belief level.*

For belief levels very near the solution belief level, the outcome is less clear, given the variance in beliefs that is likely to occur. We would still predict, though, that cooperation levels would be higher (lower) for those with realized beliefs levels lying above (below) solution belief levels. As well, when comparing two different opportunity cost level alliances that have solution level beliefs that each fall below (above) the realized belief level, we would expect greater (lesser) cooperation for the one with higher opportunity costs.

There is evidence of the high opportunity cost players acting in ways that reveal their disinterest in cooperating at the solution level (where some tolerance for defections occurs) let alone cooperating at all; such behavior drove the H1 result. The evidence includes process-type data. For example, an analysis of Study 1 revealed that several high opportunity cost players used a defect-and exit strategy (which would be coded as 0% cooperation) whereas none of the low opportunity cost players did. In Study 2b, the rationales for exiting for the high opportunity players revealed a statistically significant difference to those for the low opportunity cost players; the former were much more sensitive to defection (i.e., exit rationales 1 and 2), revealing a loss avoidance bias (apparent in H4 and H5) that would tend to hedge strategies away from cooperation where the largest potential losses can occur (see Fig. 2 for details).

From H1, we learn that firms should expect increases in alliance cooperation as opportunity costs increase only when the solution belief level lies below the expected belief level of a partner's likelihood of coopera-

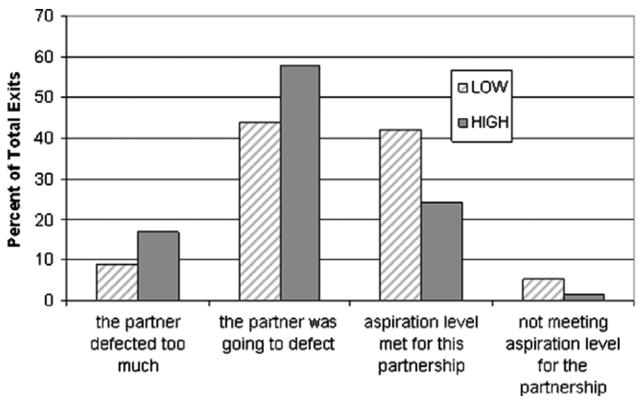

Fig. 2. Exit rationales for low and high opportunity cost players.

tion. The prescription for managers is: participate in alliances with a bias to cooperation and to stay in an alliance when the solution belief level lies below the expected belief level; otherwise refrain from alliance activity when possible. The performance results from Study 2a bear out this advice: 39 of 40 players performed above their opportunity cost level when opportunity costs were low, only 3 of 40 did so when opportunity costs were high.

H2—*partner asymmetry in opportunity costs increases cooperation received for the player that would have had a lower cooperation level under symmetry*, was supported for both types of player; asymmetry tended to decrease (increase) received cooperation for the low (high) opportunity cost players relative to the cooperation realized in symmetric alliances. Further analysis revealed that the decrease for low opportunity cost players was consistent with a decrease in payoffs as well ($t_{54} = 6.65$; $p < 0.01$); however, the high opportunity cost players were no better or worse off. Essentially, the high opportunity cost players appeared to make little adjustment to their bias for defection and quick exiting when confronted with partners that were more willing to cooperate. However, it appeared that the low opportunity cost players made a more significant adjustment in their level of cooperation ($t_{45} = 2.95$; $p < .01$), reducing it given the new partner type they confronted. Perhaps because decision-makers have a difficult time seeing through another's eyes, the adjustment by the low opportunity cost players was neither enough to solicit cooperation from their asymmetric partners nor enough to protect themselves from the defecting that those partners were expected to favor.

From H2, we learn that firms should make careful choices on partners, especially regarding any differences in opportunity costs. Assuming a manager is able to estimate a potential partner's opportunity costs relative to the focal firm's as a standard part of due diligence, the prescription is that managers should seek partners that are: more likely to cooperate (with less to lose and more to gain); more motivated to remain in the alliance; and, more tolerant of being defected upon. More specifically, we recommend that managers seek symmetric partners for low opportunity cost firms rather than partners with much higher opportunity costs.

Results for H3—*knowing when the alliance will end decreases the level of cooperation*, were consistent with established IPD work. The prescription is straightforward: managers should favor an alliance contract that is open to when the final stage is complete over an alliance contract that delineates a specific termination condition (date or stage). When firms write an alliance contract they would be advised to leave it open-ended to increase the likelihood of cooperation for as long as the relationship lasts.

Similarly, results for H4—*increasing payoffs (multiplying all values by a common factor) decreases the level of cooperation*, and H5—*imposing an entry fee reduces the level of cooperation*, were consistent with established studies of the principal of loss avoidance in decision-making. From H4, we prescribe that managers choose or structure alliances with stages that have relatively compact payoffs (i.e., where the payoff range from T to S is quite tight around o). From H5, we prescribe that managers choose alliances with lower entry costs when possible, and especially when they represent a firm with high opportunity costs.

H6—*optimal strategies vary with the opportunity cost*, was only supported for the low opportunity cost players; the high opportunity cost players did not perform better when they favored relatively less cooperative actions. The main reason is that the high opportunity cost alliances ended relatively quickly (i.e., on average at about 42% of the possible alliance length) so that the average payoff was heavily weighted to the opportunity cost level; that essentially washed out the variance from differences in the cooperativeness of strategies. The low opportunity cost alliances lasted significantly longer (i.e., on average at about 72% of the possible alliance length), which provided sufficient performance variance to support H6.

The managerial prescription from H6 is consistent with that of H1. A bias to cooperating and staying in an alliance is prescribed when the expected belief of partner cooperation lies above the solution belief level; staying out of the alliance is prescribed otherwise. Generally, we prescribe that managers adjust strategies to the opportunity costs and belief levels. And that is consistent with the results from H7. We can exploit some of the observed trends in belief level changes and the associated changes in cooperation that were described in the discussion of H7, though, to generate two further related prescriptions. First, we prescribe that the stages that need cooperation more are placed earlier in the alliance when possible. However, this manipulation requires a managerial ability to influence the sequencing of alliance stages. Second, we prescribe, when possible, managers to choose less experienced partners (i.e., those still with high beliefs of cooperation).

H7—*more accurate beliefs generate greater performance*, was supported; more accurate beliefs were correlated with superior performance. Apparently, the players that tracked the changes in partner cooperativeness and adapted to them did better. This meant picking up several trends. One trend was an initial bias to cooperation at the beginning of each alliance, relative to the level of cooperation for the rest of the alliance. A related trend was a decrease in cooperation changes for later alliances, which appeared to show some learning of what the steady-state cooperation level was.

In summary, testing the hypotheses of this study provided several prescriptions for managers. When possible, managers should (i) undertake an alliance with a bias toward cooperation, (ii) remain in an alliance as long as their partner's expected level of cooperation exceeds the solution belief level, (iii) favor alliance partners that have lower opportunity cost levels, (iv) design open-ended alliance contracts, (v) favor alliance contracts with compact payoffs, (vi) adjust alliance strategies based on best-estimates of a partner's PD-type payoffs, (vii) design alliance contracts where the most critical stages are encountered early in the relationship, and (viii) choose less-experienced alliance partners. These prescriptions address each of the three areas common in the study of alliances: choosing an appropriate alliance structure, an appropriate partner, and an appropriate strategy, given the former two selections. However, these three choices are often distinct areas of study in the literature—i.e., one is analyzed as the other two are assumed exogenously set. In our model, one change in opportunity costs affects the relative (but not absolute) payoff levels, the attractiveness of a partner, and the optimal strategy to play. In other words, our model better integrates the connection among all three elements, which perhaps helps managers see that connection more clearly to make better, more holistic decisions. Future work would consider deeper connections among the elements and within the elements. For example, what selection effects exist for firms that apply as potential alliance partners for specific alliance types? what trade-offs exist between alliance structure choices (e.g., do lower entry costs make it more difficult to keep the end date of an alliance open)? will the consideration of different cultures affect the belief contingencies in the model when considering international alliances? how will a limited pool of potential alliance partners affect play (e.g., will a reputation for defecting or for exiting affect an ability to ally)?

From our research, we might expect selection effects to influence alliance outcomes through co-evolving cooperation levels among partners (e.g., the firms that select into marketing alliances may differ from those in R&D alliances on the trust they have and thus explain some differences in outcomes of those alliance types). We might also expect that structural elements of certain alliances are connected (e.g., big entry fees would probably make players want to be more flexible) and hence should all be considered in predicting the alliance outcome. We might also expect that any extra-model factors that are added to affect the payoffs of the game, like reputation factor, would affect how one player evaluates another then decides how to proceed. In general, we would expect more accurate information about a player's cooperative (non-cooperative) reputation to increase (decrease) earnings.

The theoretical implication is more open. If the addition of a pre-stage to a standard iterated game like the PD can generate such different and more complex (i.e., contingent) results, what does that mean to other less studied games that are touted as applicable to real business opportunities that are also voluntary?

The main results of our experiments were consistent with the theoretical results of Arend and Seale (2005), but only when the contingency of beliefs was accounted for. Alliance cooperation actually fell with the opportunity cost level increase, but this was because the contingency was violated—i.e., the belief of partner cooperation fell below the level of cooperation required to match opportunity costs. The experiment brought out the difference between game theory agents and real players in the assessment of those beliefs. Real decision-makers have a heuristic adjustment to beliefs based on the opportunity cost levels and experience, among other factors. This experiment has exposed the importance of that belief level contingency on the results of the IPDEO, which is a valuable complement to the theoretical paper. The importance of what influences beliefs, then, relates to several other theories: Is there a path-dependency story (and is it consistent with Bayesian updating of beliefs) that could explain alliance outcomes of specific firms? Is there an agglomeration effect on beliefs, where the closer non-focal relationships of potential partners influences the beliefs? Is there a intra-industry homogeneity effect on the accuracy of beliefs? Is there a network effect on beliefs where non-business ties also influence alliance outcomes?

The main results differed more substantially from the tournament research by Phelan et al. (2005), mostly due to the sophistication of the strategies used, which is not surprising. The most robust strategies in the tournament where tit-for-tat with triggered exits, where those triggers were adjusted for the opportunity cost level (i.e., with quicker triggers for higher opportunity cost levels). The real decision-makers relative to the programmed ones were perhaps on one hand more emotional and short-sighted because their triggers were even quicker for the high opportunity cost games, but on the other hand more frame-sensitive as they were more forgiving at the low opportunity cost level. There was more variance in strategies with real players; for instance, many players decided a defect-and-exit strategy at the high opportunity cost alliances was a viable strategy. The contrast of the experiment to the tournament generated some valuable insights into how decision-makers with immediate decisions based on immediate feedback can differ from those who must commit to a plan that attempts to account for all possible future outcomes; the latter, at least for this game, generates higher cooperation.

Several shortcomings and limitations of this study are noted. First, the IPDEO model may not apply to all alliances, only those that: involve multiple stages of commitment where the payoffs to the alliance partners in each stage are commonly known (or can be accurately estimated) and resemble a PD situation; and, those where partner opportunity costs can be accurately estimated. Second, several aspects of the experimental design may be questioned, including: no control for gender effects and other possibly relevant subject characteristics; a focus on only two opportunity cost levels and only one payoff matrix; and, the usual budgetary limitations that required us to pay subjects in a probabilistic manner. We have, however, no evidence to suspect that gender, other relevant demographic characteristics, testing more that two levels of opportunity cost, or paying subjects in an alternative manner would change these results.

We set out to model the two-firm alliance as an IPDEO and to test several theoretical predictions in order to generate new prescriptions for managers contemplating such activity. The results of the studies show that several issues have a significant impact on the level of cooperation in, and payoff from, alliance activity. Chief among these issues are: opportunity costs; beliefs of partner cooperation; and, knowledge of the maximum potential alliance length. The first issue is new to the alliance literature because alliances have not been modeled as IPDEOs before, and as such, the exit option has not been considered as an issue to consider in alliance performance. The results of experimental tests reveal that it is a major issue. Firm opportunity costs affect directly and indirectly an alliance's cooperation level, its length, and the payoffs it generates. Opportunity costs directly affect the strategies chosen in the alliance, which affects the alliance outcomes. Opportunity costs also affect the influences of other factors, like partner choice and entry fees, that then affect strategies and outcomes. The main finding is that when the opportunity costs are at a level where the beliefs in a partner's likelihood of cooperation required to sustain an alliance payoff at the opportunity cost lie below the actual belief level of a partner's likelihood of cooperation, then alliance activity is likely to be beneficial and influences of other factors that might be detrimental will be less significant.

Future work should consider several open questions drawn from our results: Do real firms in alliance situations behave like our subjects, or rational game theorists? Are real firms with higher opportunity less cooperative than firms with lower opportunity costs? Do alliance contracts specifying particular stages of investment or endpoints hinder cooperation relative to open-ended alliance agreements? What complications can we expect if noise or reputation effects are present? Empirical studies of these questions, as well as simulations of the interactions under the conditions of these questions, should lead to some valuable insights into why some alliances succeed while others fail.

Appendix A

University of Nevada Las Vegas
Department of Management
Alliance Formation Experiment

Welcome to the experiment on alliance formation. The instructions are simple. If you follow them carefully and make good decisions, you will earn a considerable amount of extra credit points. The number of points you earn will depend on your relative performance in this experiment. In addition to earning extra credit points, several of you will be randomly selected at the end of the experiment and paid a cash bonus. The amount of the bonus will also depend on the number of points you earn. How you earn points will be explained below. If you have any questions before or during the experiment, please raise your hand.

Background

Businesses often form alliances when there are opportunities to share costs, share risks, or gain synergies by working together. Some alliances are short-lived, lasting no longer than a single project, while other alliances may span multiple projects and last as long as both parties are willing to continue the commitment. Typically, an alliance ends when the project is completed, or when one or both partners decide to opt-out. During the course of an alliance, the alliance partners may have multiple chances to either cooperate or compete with one another. While cooperation may result in higher earnings for the alliance as a whole, each individual partner may find that they could earn more by competing. The decision the partners often face is one where they (1) earn a fair return by cooperating, (2) earn considerably more by competing, but may risk losing their alliance partner, or (3) may opt-out of the alliance altogether and earn, on average, the return of their next best alternative. This is the kind of decision we are trying to model in this experiment.

Description of the task

Look around the computer lab. You are one of many people participating in today's experiment. At the beginning of the experiment the computer will randomly pair you with another person. This paring, which is explained in greater detail below, will last for a number of rounds (trials). At the beginning of each trial the computer will display a payoff table, similar to the one below, and ask you to make one of three decisions—*cooperate*, *compete* or *opt-out*. Your payoff will be determined by the decisions that both you and your partner make.

After you have made your decision you will see a screen asking you to *please be patient*. You are waiting for other people to complete their decisions. After everyone has made their decisions, the computer will display the results for the trial—you will see the choices that you and your alliance partner made, your earnings for the trial and your cumulative earnings across all trials. For example (see table below), if both you and your alliance partner choose to cooperate for the trial, each will earn 8 francs (francs are an experimental currency that may be converted to extra credit points at the end of the experiment). If you choose to cooperate and your alliance partner chooses to compete, you will earn nothing, while your alliance partner will earn 10 francs. If you choose to compete and your alliance partner chooses to cooperate, you will earn 10 francs, while your alliance partner will earn nothing for the trial. If both of you choose to compete, each will earn 3 francs. As you examine the payoff table below, note that the columns represent your alliance partner's decision, and the rows represent your decision. *The first number in each cell of the table is your payoff and the second number is your partner's payoff.*

Your decision	Your alliance partner's decision	
	Cooperate	Compete
Cooperate	8, 8	0, 10
Compete	10, 0	3, 3
Opt-out	4, 4	4, 4

On any trial you may also choose to opt-out. If either you or your alliance partner chooses to opt-out, each of you will earn 4 francs for the remainder of trials that you are paired.

Reputation/outcomes

The number of cooperate, compete and opt-decisions that each player has made are tracked by a central computer. Each time you are paired with a new player, several measures summarizing that player's reputation or past outcomes are reported on your decision screen. At the top of this screen you will see a table, similar to the one listed below. In this example, imagine that you have completed 22 trials and are once again paired with a new player. The following items are reported:

Reputation/outcomes of player you are paired with	
Percentage of scheduled trials completed:	50%
Number of *Cooperate* decisions:	8
Number of *Compete* decisions:	3
Number of *Opt-out* decisions:	2

In the example above, the player you were just paired with cooperated on 8 previous trials, competed on 3 trials, and opted out of 2 pairings. Although you do not know when or why this player opted out, you can tell that he/she completed 50% of the scheduled trials [(8 cooperate + 3 compete)/22 completed trials]. Note that the player you are paired with is given similar information on your previous decisions. The number of *cooperate*, *compete*, and *opt-out* decisions, as well as the *percentage of scheduled trials completed* is updated at the start of each trial.

Estimating your partner's decision

After reviewing the reputation/outcome information, you are asked to estimate the probability that the player you are paired with will cooperate on the next trial. You make this estimate, from 0 to 100%, by moving the slider button with the computer mouse. Note that the computer requires you to make an estimate before allowing you to make your decision to cooperate, compete or opt-out.

The experiment will last *at least 50 trials* and you will likely be paired with several different alliance partners during the course of the experiment. The *length of each pairing is determined randomly*, and you will not know in advance how long each pairing may last. However, the computer will inform you each time you are paired with a new alliance partner.

Payment

At the end of the experiment, we will report to your professor the number of francs that you earned during the experiment and assist your professor in converting francs to extra credit points that he/she has offered you for participating in this research. We will also randomly select several of you and pay you in cash, based on your performance. If you are one of the players randomly selected, francs will be converted to US dollars at the rate of 20 francs = $1.

If you have any questions, please raise your hand. We sincerely thank you for your participation.

Good Luck!

(Alliance Instructions SLF.doc)

References

Arend, R., & Seale, D. A. (2005). Modeling Alliance Activity: An iterated prisoners' dilemma with exit option. *Strategic Management Journal, 26*, 1057–1074.

Ashlock, D., Smucker, M. D., Stanley, E. A., & Tesfatsion, L. (1996). Preferential partner selection in an evolutionary study of prisoner's dilemma. *BioSystems, 37*, 99–125.

Aumann, R. J. (1995). Backward induction and common knowledge of rationality. *Games and Economic Behavior, 8*, 6–19.

Axelrod, R. (1980). More effective choice in the prisoner's dilemma. *Journal of Conflict Resolution, 24*, 379–403.

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books Inc.

Binmore, K. (1992). *Fun and games: A text on game theory*. Lexington, MA: D.C. Heath and Company.

Binmore, K. (1999). Why experiment in economics? *The Economic Journal, 109*, f16–f24.

Bleeke, J., & Ernst, D. (1995). Is your strategic alliance really a sale? *Harvard Business Review, 73*(1), 97–105.

Boone, R. T., & Macy, M. W. (1999). Unlocking the doors of the prisoner's dilemma: Dependence, selectivity and cooperation. *Social Psychology Quarterly, 62*, 32–52.

Cable, D. M., & Shane, S. (1997). A prisoner's dilemma approach to entrepreneur-venture capitalist relationships. *Academy of Management Review, 22*(1), 142–176.

Cachon, G. P., & Camerer, C. F. (1996). Loss-avoidance and forward induction in experimental coordination games. *The Quarterly Journal of Economics*, 165–194.

Das, T. K., & Teng, B. S. (2000). Instabilities of strategic alliances: An internal tensions perspective. *Organization Science, 11*(1), 77–101.

De Rond, M. (2003). *Strategic alliances as social facts: Business, biotechnology, and intellectual history*. Cambridge University Press.

De Rond, M., & Hamid, B. (2004). On the dialectics of strategic alliances. *Organization Science, 15*(1), 56–69.

Fehr, E., & Gachter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review, 90*(4), 980–994.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition and cooperation. *The Quarterly Journal of Economics, 114*(3), 817–868.

Fudenberg, D., & Masken, E. (1986). The Folk theorem in repeated games with discounting or with incomplete information. *Econometrica, 54*(3), 533–554.

Gulati, R. (1998). Alliances and networks. *Strategic Management Journal, 19*(4).

Harbinson, J. R., & Pekar, P. (1998). *Smart alliances*. San Francisco: Booz Allen and Hamilton Business Press.

Hofstadter, D. R. (1985). *Metamagical themas*. New York: Basic Books Inc.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–291.

Knez, M., & Camerer, C. (2000). Increasing cooperation in prisoner's dilemmas by establishing a precedent of efficiency in coordination games. *Organizational Behavior & Human Decision Processes, 82*(2), 194–217.

McCusker, C., & Carnevale, P. J. (1995). Framing in resource dilemmas: Loss aversion and the moderating effects of sanctions. *Organizational Behavior & Human Decision Processes, 61*(2), 190–200.

Mitchell, W., & Singh, K. (1996). Survival of business using collaborative relationships to commercialize complex goods. *Strategic Management Journal, 17*(3), 169–195.

Mussweiler, T. (2001). The durability of anchoring effects. *European Journal of Social Psychology, 31*(4), 431–442.

Neyman, A. (1999). Cooperation in repeated games when the number of stages is not commonly known. *Econometrica, 67*(1), 45–65.

Orbell, J. M., & Dawes, R. M. (1993). Social welfare, cooperator's advantage and the option of not playing the game. *American Sociological Review, 58*, 787–800.

Parkhe, A. (1993). Strategic alliance structuring: a game theoretic and transaction cost examination of interfirm cooperation. *Academy of Management Journal, 36*(4), 794–829.

Parkhe, A., Rosenthal, E. C., & Chandran, R. (1993). Prisoner's dilemma payoff structure in interfirm strategic alliances: An empirical test. *Omega, 21*(5), 531–539.

Pekar, P., Jr., & Allio, R. (1994). Making alliances work—guidelines for success. *Long Range Planning, 27*(4), 54–65.

Phelan, S., Arend, R. J., & Seale, D. A. (2005). Using an iterated prisoner's dilemma with exit option to study alliance behavior: Results of a tournament and simulation. *Computational and Mathematical Ogranization Theory, 11*(4), 339–356.

Rapoport, A., & Chammah, A. M. (1965). *Prisoner's dilemma*. University of Michigan Press.

Roxburgh, C. (2003). Hidden flaws in strategy. *McKinsey Quarterly, 2*, 26–39.

Rubinstein, A. (1979). Equilibrium in supergames with the overtaking criterion. *Journal of Economic Theory, 21*, 1–9.

Schuessler, R. (1989). Exit threats and cooperation under anonymity. *Journal of Conflict Resolution, 33*, 728–749.

Schifrin, M. (2001). Partner or perish. *Forbes*, 26–28.

Slovic, P., & Lichtenstein, S. (1973). Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance, 6*, 649–744.

Telser, L. G. (1980). A theory of self-enforcing agreements. *Journal of Business, 53*, 27–41.

Tullock, G. (1985). Adam Smith and the Prisoners' Dilemma. *The Quarterly Journal of Economics, 100*, 1073–1081.

Tversky, A., & Kahneman, D. (1986). Rationality choice and the framing of decisions. *Journal of Business, 59*(4), s251–s278.

Vanberg, V. J., & Congleton, R. D. (1992). Rationality, morality, and exit. *American Political Science Review, 86*(2), 418–431.

Yamagishi, T., & Hayashi, N. (1996). Selective play: Social embeddedness of social dilemmas. In W. B. G. Liebrand & D. M. Messick (Eds.), *Frontiers in Social dilemmas research* (pp. 363–384). Berlin: Springer-Verlag.

Yung-An, H., & Day-Yang, L. (2003). Altruism versus egoism in human behavior of mixed motives. *American Journal of Economics and Sociology, 62*(4), 677–705.

Zeng, M., & Chen, X.-P. (2003). Achieving cooperation in multiparty alliances: A social dilemma approach to partnership management. *Academy of Management Review, 28*(4), 587–605.