

Defense Against an Opportunistic Challenger: Theory and Experiments*

Cary Deck Joshua Foster Hongwei Song

Economics Department
University of Arkansas

May 27, 2014

Abstract

This paper considers a contest setting in which a challenger chooses between one of two contests to enter after observing the level of defense at each. Despite the challenger's chance of success being determined by a proportional contest success function, the defenders effectively find themselves in an all-pay auction that largely dissipates the value of the defended resources because the challenger will target the weaker defender. However, if the defenders form a protective alliance then their expected profits increase despite the fact that a successful challenge is theoretically more likely, given the overall reduction in defense. Controlled laboratory experiments designed to test the model's predictions are also reported. Observed behavior is generally consistent with the comparative static predictions although challengers exhibit the familiar overbidding pattern. Defenders appear to anticipate this reaction and adjust their behavior accordingly.

Keywords: Game Theory, Contests, Attack, Defense Alliances, Experiments, Psychology of Decisions

JEL Classification: C72, C91, D72, D74

*The authors wish to thank Brendan Joyce, Dan Kovenok, Yiwei Qian, Jared Reber, Sudipta Sarangi, Roman Sheremeta, Christian Vossler, seminar participants at Appalachian State University, Baylor University, Georgia State University, the University of Memphis, Simon Fraser University, the University of Tennessee, Texas Christian University, and attendees of the Southern Economic Association Annual Meetings, and the Conference on Tournaments, Contests, and Relative Performance Evaluation for helpful comments and suggestions at various stages of the development of this project.

1 Introduction

In many situations that can be described as a contest, one of the participants (a challenger) can decide which contest to enter after observing the behavior of the other contestants (the defenders). For example, the challenger could be a terrorist who has a single bomb and multiple possible targets such as planes owned by different airlines. The terrorist has the advantage of being able to observe the relative strength of each target's defense and respond accordingly. Intuition suggests that the terrorist would prefer to attack the weaker target thereby increasing the chance of success. Since the more strongly defended target does not get attacked, each defender has an incentive to be slightly better protected than its rival resulting in an all-pay auction among the defenders. Such a situation arises in other settings as well. *Ceteris paribus*, a criminal prefers to burgle the least protected house in a neighborhood, explaining the popularity of home security and "Beware of Dog" signs. An employee hoping to become a regional manager only needs to outshine the weakest current person in that position, just as a new politician can gain office by unseating the weakest incumbent.¹ An entrepreneur looking to start a new retail store would prefer to operate where the competitor is the weakest. A young male animal would prefer to usurp the feeblest established male to claim mating rights. This situation also arises in the old joke about two people going hiking in an area inhabited by bears when one points out that they cannot outrun a bear, and the other says, "I just have to outrun you."²

Rather than providing separate defenses, in some settings the defenders could band together and form an alliance. For example, airplane security is done at the airport level rather than the airline level. Residential communities often form neighborhood watches. Incumbent firms may seek a zoning ordinance to keep potential entrants out. In fact, alliances are common throughout society and psychologists have argued that people favor the formation of an alliance when facing conflicts due to the competitive disadvantage of the lone individual confronting a group (Baumeister and Leary 1995).

In this paper, we construct a formal model to analyze these two strategic situations and test the model using controlled laboratory experiments. The theoretical results confirm that the challenger will prefer to attack the weaker defender when targets are protected independently leading defenders to invest heavily. A challenger who targets the weak link has the flair of

¹Ryvkin (2010) discusses various types of tournaments for trying to identify high quality employees when talent is heterogeneous, although the structure of uncertainty is different in his setup.

²Of course, any of these situations could involve multiple challengers: two burglars could operate in the same neighborhood; other employees may be seeking a promotion; etc. Allowing for more players may introduce multiple and/or asymmetric equilibria. However, this paper restricts attention to the case of two defenders and one challenger.

previous research on the attack and defense of a network (e.g., Major, 2002; Woo, 2002, 2003; O’Hanlon et al., 2002, Levitin, 2003a,b). In contrast, when the defenders work together in an alliance, the aggregate level of defense is much lower resulting in both a greater likelihood of a successful challenge and simultaneously higher expected profits for defenders.

The normal intuition for an alliance is that the joint defense is greater than each individual defense and thus the alliance is better able to deter or handle a challenger. Research by Sheremeta and Zhang (2010) suggests alliances make better decisions in contests than individuals. Specifically, in lottery contests when team members are able to communicate, groups are found to make more rational decisions than individuals. While most of the literature on contests has not focused on alliances, there has been some work considering the impact of how the alliance shares the spoils of its success (Katz and Tokalidu 1996; Esteban and Sakovics 2003; Muller and Warneryd 2001; Warneryd 1998; and Konrad 2004).³ In these models there is typically a single prize to be allocated among members of the alliance.

In general these models find that the internal conflict diminishes the contribution of alliance members. This outcome is also found when there are spillovers between independently defended targets in a network such as in Kunreuther and Heal (2003). However, Ke et al. (2010) conduct an experimental analysis of alliances and show that the future internal conflict does not prevent alliance members from fighting shoulder-to-shoulder. On average, they find that allies in a contest against an outside opponent devote the same contest effort irrespective of how they will share the spoils of victory. In addition, the collaboration in alliances is reasonably good, leading to higher success against lone challengers than predicted. Garfinkel (2004) develops a positive analysis of alliance formation, building on a simple economic model that features a “winner-take-all” contest for control of some resource. When an alliance forms, members pool their efforts in that contest and, if successful, apply the resource to a joint production process. Due to the familiar free-rider problem, the formation of alliances tends to reduce the severity of the conflict over the contestable resource. Despite the conflict that arises among the winning alliance’s members over the distribution of their joint product, under reasonable conditions, this effect alone is sufficient to support stable alliance formation in a non-cooperative equilibrium.

Our model is distinct from these papers in that each member of the alliance values its own item so that the alliance is about common protection rather than an arrangement for increasing

³An alliance in our setting reduces the game to a single battle, which is distinct from the setting in which the entire defense of a network is defended by a single decision maker as in Bier and Abhichandani (2002), Bier et al. (2005), Azaiez and Bier (2007) and Hausken (2008) where defense remains target specific. Kovenock and Roberson (2012a,b) provide the necessary conditions for and discuss the misunderstandings in Hausken’s (2008) results.

the chance of claiming a shareable prize of a given size.⁴ Thus, in our setting there is no distributional conflict within the alliance resulting from a successful defense. Further, in our model the challenger cannot claim more than one prize regardless of whether or not the defenders opt to form an alliance. Returning to the example of a terrorist with a single bomb attempting to attack a plane, if the terrorist is unsuccessful both airlines retain their respective planes but if the terrorist is successful only one airline incurs the entire loss while the other incurs no harm. In the example of an employee vying for a regional manager job, an incumbent who keeps her job is not harmed when someone else is let go.

The paper most closely related to ours in structure is Dighe, et al. (2009), which considers an attack and defense game with two possible targets and one challenger. In their game, defense is a binary choice and the outcome is deterministic, as an attack is only successful if launched against an undefended target. They compare a decentralized defense where different decision makers defend each target and a centralized defense where a single decision maker makes both defense decisions jointly thereby internalizing the externality associated with defense. In their setup, defense is unobservable and they find that centralized decision making is optimal since deterrence can be achieved in some scenarios by protecting only one target. Our paper is also similar to Hausken and Bier (2011) who consider a single defender with multiple attackers using a similar conflict success function. They show how the move order and relative values of the attackers will encourage some attackers to abandon the contest.

We also report the results of controlled laboratory experiments designed to test the empirical validity of our model. In our laboratory experiments, defenders are observed to bid less when in an alliance as predicted by the model. However, the difference in the bids is not as dramatic as predicted. There are now several experimental papers on contests (cf. Sheremeta, et al. 2012 for a thorough survey) and one of the common findings is that people overbid to the point that the equilibrium surplus is often fully depleted (cf. Davis and Reilly 1998; Potters et al. 1998; Gneezy and Smorodinsky 2006; and Lugovskyy and Puzello 2008). Noussair and Silver (2006) addresses the effect of experience, showing that experience helps decrease over-bidding but

⁴Our results are unchanged if alliance members agree to equally share all prizes that any member claims. The alliance structure in our paper is also related to the literature on group contests (cf. Muenster 2009 who extends the axiomatic characterization of contest success functions of Skaperdas 1996 and Clark and Riis 1998 to contests between groups). Baik (2008) examines the equilibrium effort levels of individual players and groups in contests in which n groups compete to win a group-specific public-good prize. In the basic model the chance of success depends on total effort and only the highest-valuation players expend positive effort leading to under-investment in the contest for the group as a whole. Lee (2012) considers the situation in which the probability of winning follows a weakest-link rule so that it is the lowest-valuation players in each group who play the decisive roles. Ryvkin (2011) studies how aggregate effort exerted in contests between groups of heterogeneous players depends on the sorting of players into groups. Abbink et al. (2010) examines the impact of group members being able to punish each other.

does not eliminate it. Contrary to these previous contest experiments, we find that defenders under bid when defending separately, perhaps because the theoretical predictions are relatively greater in our setting. Our results are also driven in part by the fact that the alliance members do not internalize the benefits of their defense investments for the other alliance members. This aspect of alliance behavior was pointed out at least as far back as Olson and Zechhauser (1966). However, Ke et al. (2010) observe group members overbidding in a setting where the group shares a common bid against another party and equally split the proceeds from a successful bid. Recently, Nitzan and Ueda (2008) examine the effect of group size on performance in a collective contest and find that larger groups tend to be less effective at pursuing the collective interest.

2 Theoretical Model

Consider a situation in which a single challenger has two possible targets, T_1 and T_2 , each valued at $P \geq 0$. T_i is valued at $V \geq 0$ by defender i . Allocation of the targets is determined by the outcome of a contest resolved with a proportional success function (Tullock 1980) based upon the level of investment by the relevant contestants. Let the investment (bid) by defender i be denoted by $b_i \geq 0$ and the investment (bid) of the challenger be denoted by $b_C \geq 0$. Before investing in defense, the defenders have the option to defend the targets independently or form an alliance. After observing the defense structure and the investments of the defenders, the challenger decides how to proceed. A key feature of this set-up is that a challenger can enter at most one contest and can claim at most one target.⁵ Our model is meant to capture the situation in which the second mover terrorist has a single bomb or the second mover burglar only has enough time to enter a single home.

Formally, the game proceeds through 3 stages.

Stage 1: *Defenders decide to form an alliance or not.*

Stage 2: *Conditional on the outcome of Stage 1, defenders independently select their investment levels.*

Stage 3: *Conditional on the outcomes of Stages 1 and 2, challengers decide which contest to enter and how much to invest.*

⁵If the challenger could enter both contests in the event that the defenders opted to defend separately, the result would be two independent standard Tullock contests. Allowing the challenger to claim both targets when defeating an alliance is discussed in footnotes 8 and 9.

The appropriate solution concept is subgame perfection; therefore, we consider the stages in reverse order.

2.1 Stage 3: Response of the Challenger

How the challenger should respond depends on whether or not the defenders formed an alliance.

2.1.1 Case 1: Response to Independent Defenses

At this stage, the challenger observes the defense investment at each target. The challenger's problem is to decide which target if any to pursue and how much to invest. From pursuing target T_i , the challenger's expected profit is

$$\Pi_C = \begin{cases} \frac{b_C}{b_C + b_i} P - b_C & \text{if } b_C + b_i > 0 \\ \frac{P}{2} & \text{else} \end{cases} . \quad (1)$$

where the second term is due to the fact that the winner is determined randomly if $b_C + b_i = 0$. The optimal response by the challenger, derived from the first order condition for (1), is

$$b_C^* = \begin{cases} \sqrt{P b_i} - b_i & \text{if } b_i < P \\ 0 & \text{else.} \end{cases} \quad (2)$$

When it is optimal to attack, substituting (2) into (1) yields $\Pi_C^* = (\sqrt{P} - \sqrt{b_i})^2$, which is decreasing in b_i . Therefore, the challenger finds it more profitable to pursue the less defended target and is indifferent between the targets if they are equally defended.

2.1.2 Case 2: Response to an Alliance

In the alliance, the challenger faces the combined defense of the alliance members and if successful then randomly selects one of the targets to claim.⁶ At this stage, the challenger maximizes

⁶An alternative cooperative arrangement is for the defenders to communicate and coordinate their activity, essentially merging into a single decision making entity and thus internalizing the positive externality associated with investing. There are two possible implementations of this arrangement mirroring our independent and alliance set-ups. The parallel to the independent defense is such that the defender will choose to invest the same amount at each target because the challenger will still prefer to contest the weaker target and thus any additional investment on one target is wasted. Hence, the objective function of this single defender would be $\Pi_D = \left(1 + \sqrt{\frac{b_D}{P}}\right) V - 2b_D$ where D denotes the single defender and b_D is the defender's level of investment for each target. Notice that in this case the defender

$\Pi_C = \frac{b_C}{b_C + b_1 + b_2} P - b_C$. The first order condition yields the optimal challenge given by

$$b_C^* = \begin{cases} \sqrt{P(b_1 + b_2)} - b_1 - b_2 & \text{if } b_1 + b_2 < P \\ 0 & \text{else.} \end{cases} \quad (3)$$

2.2 Stage 2: Investments of the Defenders

How the defenders should behave at this stage depends on whether or not they formed an alliance in the first stage and takes into account how the challenger will respond in the next stage.

2.2.1 Case 1: Independent Defense

At this stage, the defenders know that the challenger will focus on the weaker target. The implication is that the stronger defender will claim the prize V with certainty while the weaker defender will earn V only if the ultimate contest with the challenger is unsuccessful. Equation (2) informs the weaker defender of how the challenger will react. Letting b_w denote the level of the weaker defense and using (2), the weak defender expects to earn

$$\Pi_w = \begin{cases} V\sqrt{\frac{b_w}{P}} - b_w & \text{if } b_w < P \\ V - b_w & \text{else.} \end{cases} \quad (4)$$

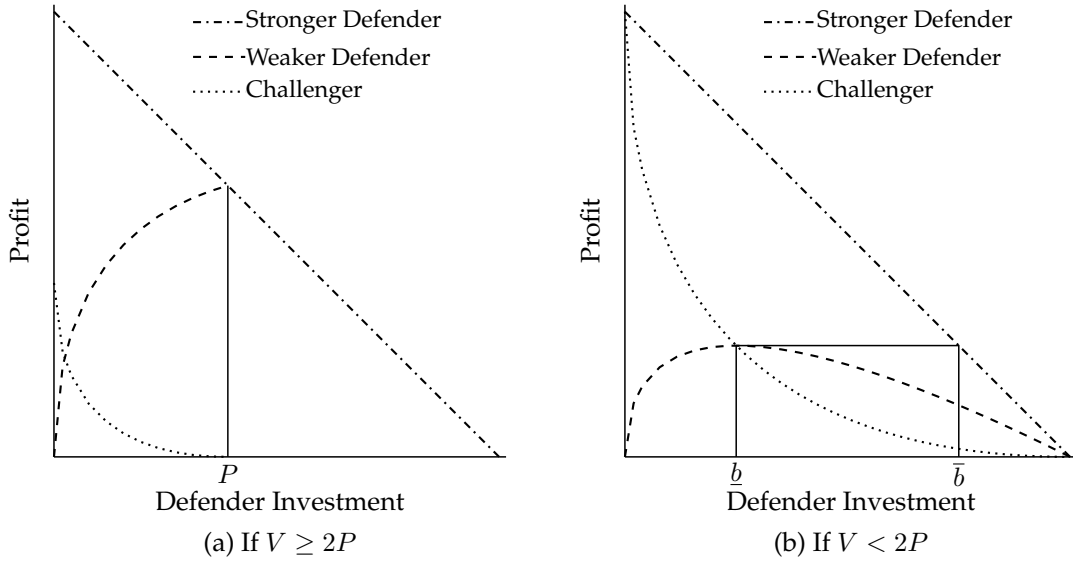
Thus, the defender of target T_i earns $V - b_i$ if $b_i > b_j$, the profit given by (4) with $b_w = b_i$ if $b_i < b_j$, or one of these two amounts selected randomly if $b_i = b_j$, for $j \neq i$.⁷ Notice that the weak defender's profit, if challenged, is maximized when $b_w = \frac{V^2}{4P}$. Thus there are two cases to consider depending on how $\frac{V^2}{4P}$ compares to P , the defense bid above which the challenger drops out, as in Hausken and Bier (2011).

If $V \geq 2P$ then the profit of the defenders is shown by Panel (a) of Figure 1. In this setting,

is assured of receiving V as one target will not be attacked. In this case $b_D^* = \frac{V^2}{16P}$ and the optimal response by the challenger is $b_C^* = \frac{V}{4} \left(1 - \frac{V}{4P}\right)$ which is the same as the level of investment by the challenger in the alliance. The defender's expected profit in this case would be $V + \frac{V^2}{8P}$ and the challenger's expected profit would be $(1 - \frac{V}{4P})(P - \frac{V}{4})$. The other set up would allow the one defender to jointly protect both targets. Here the objective function of this single defender would be $\Pi_D = \left(1 + \sqrt{\frac{b_D}{P}}\right) V - b_D$ and the resulting expected profit to the defender would be $V + \frac{V^2}{4P}$. Clearly, of these two choices a single decision maker would prefer to jointly protect the two targets. Spolaore (2010) goes through a similar exercise when looking at various alliances and political unions in geopolitical contests.

⁷The interaction between the independent defenders has some similarity with the model of Baye et al. (2012); however our case is distinct in that the loser's profit is not monotonically decreasing in the amount bid.

Figure 1: Profit by Player



the equilibrium is for both defenders to set a bid of P . It is clear from the figure that no defender would ever want to bid above P with any positive probability because her profit would be increased by increasing the probability of bidding P by that amount. It is also the case that no defender would bid below P with any positive probability in equilibrium. Obviously, neither defender would ever want to bid strictly below the lowest possible bid that the other defender might make. Hence, the two defenders would have to have the same lower bound for their bidding behavior, but then a defender could increase her expected profit by shifting weight from close to this lower bound to P .

If $V \leq 2P$ then the profit of the defenders is shown by Panel (b) of Figure 1. Here the defenders find themselves in an all-pay auction in which the low bidder's profit is a function of her bid. The maximum profit a defender can assure himself is the same as the profit the defender would expect to earn if there was a single target (or he knew he would be contested with certainty). In that case, the defender would choose to defend at the level $b = \frac{V^2}{4P}$, which is the bid that maximizes (4). This level of investment, \underline{b} , identifies the security profit for a defender, $\frac{V^2}{4P}$ (found by plugging $b_w = \frac{V^2}{4P}$ into equation 4). Based on a similar argument as before, a bidder will never find it optimal to bid below \underline{b} with any positive probability. Regardless of whether or not the target would be contested, a defender would never find it optimal to bid strictly above

$\bar{b} = V - \frac{V^2}{4P}$ because doing so would yield a profit strictly less than the $\frac{V^2}{4P}$ profit that can be assured by bidding \underline{b} . It is straightforward to show that $V < 2P$ implies $\bar{b} < P$ and hence in the region of interest, defender bids must be in the interval $[\underline{b}, \bar{b}]$. Following Baye et al. (1996), this all-pay auction with complete information will have a unique symmetric mixed strategy Nash equilibrium, $f(b)$. Informally, no bidder can find it optimal to play a strategy that has mass points on some bids or gaps in the support of the distribution because the rival would react in such a way to disadvantage the bidder. This means that the players are using a continuous distribution, which is uniquely identified by generating an expected profit equal to the security profit. Let $F(b)$ be the cumulative density function associated with $f(b)$. If the other defender is playing according to $f(b)$, then defender i 's problem is to maximize

$$\Pi_i = \left(\frac{b_i}{b_i + b_C^*} V - b_i \right) [1 - F(b_i)] + F(b_i)(V - b_i). \quad (5)$$

Substituting b_C^* from (2) into (5) yields

$$\Pi_i = \left(\sqrt{\frac{b_i}{P}} V - b_i \right) [1 - F(b_i)] + F(b_i)(V - b_i).$$

Since each strategy generates the same expected profit in equilibrium, it must be that

$$\frac{V^2}{4P} = \left(\sqrt{\frac{b_i}{P}} V - b_i \right) [1 - F(b_i)] + F(b_i)(V - b_i). \quad (6)$$

Solving (6) for $F(b_i)$ yields

$$F(b_i) = \frac{\frac{V^2}{4P} - \sqrt{\frac{b_i}{P}} V + b_i}{\left(1 - \sqrt{\frac{b_i}{P}}\right) V}. \quad (7)$$

It is straightforward to show that $F(\underline{b}) = 0$, $F(\bar{b}) = 1$, and (7) is increasing in b_i . As $f(b_i) > 0$ over the interval $[\underline{b}, \bar{b}]$, equation (7) defines the unique symmetric Nash equilibrium.

2.2.2 Case 2: Alliance Defense

At this stage, defender i maximizes her expected profit, given by

$$\Pi_i = \frac{b_i + b_j}{b_C^* + b_i + b_j} V + \left(1 - \frac{b_i + b_j}{b_C^* + b_i + b_j}\right) \frac{V}{2} - b_i$$

which, taking (3) into account, simplifies to

$$\Pi_i = \left(1 + \sqrt{\frac{b_i + b_j}{P}}\right) \frac{V}{2} - b_i \text{ where } i \in \{1, 2\} \text{ and } j = 3 - i. \quad (8)$$

The first order condition of (8) leads to a best response function $b_i^*(b_j) = \frac{V^2}{16P} - b_j$, which implies that any pair of non-negative defender bids that sum to $\frac{V^2}{16P}$ is an equilibrium if $\frac{V^2}{16P} < P$ or $V < 4P$.⁸ While there are multiple equilibria, the per capita equilibrium defender bid is unique (cf. Nti 1998). The challenger responding to the total defense will bid $b_C^* = \frac{V}{4} \left(1 - \frac{V}{4P}\right)$. The average expected profit of the defender is $\left(1 + \frac{3V}{16P}\right) \frac{V}{2}$ while the challenger expects to earn $\left(1 - \frac{V}{4P}\right)^2 P$. If on the other hand, $V \geq 4P$ then the two defenders would prefer to bid a total of P and take the challenger out of the game. Again there are multiple equilibria, but in any equilibrium the average profit of a defender will be $V - \frac{P}{2}$ and the challenger will earn 0.

2.3 Stage 1: Alliance Formation

In the first stage of the game, defenders independently decide to form an alliance or not. If both defenders opt to form the alliance, an alliance is formed. Otherwise, it is not. The expected profits to defenders, regardless of forming an alliance or not, depend on the relative values of V and P as summarized in Table 1.

Table 1: Bid Amounts and Expected Profits for Defenders

<i>Parameters</i>	<i>Bid Amounts</i>		<i>Expected Profits</i>		
	<i>Independent</i>	<i>Alliance</i>	<i>Independent</i>	<i>Alliance</i> ^A	<i>Decision</i>
$V \geq 4P$	$b_i, b_j = P$	$b_i + b_j = P$	$V - P$	$V - P$	Alliance
$4P > V \geq 2P$	$b_i, b_j = P$	$b_i + b_j = \frac{V^2}{16P}$	$V - P$	$\frac{V}{2} + \frac{V^2}{16P}$	Alliance
$V < 2P$	$b_i, b_j \in \left[\frac{V^2}{4P}, V - \frac{V^2}{4P}\right]$	$b_i + b_j = \frac{V^2}{16P}$	$\frac{V^2}{4P}$	$\frac{V}{2} + \frac{V^2}{16P}$	Alliance

^A: Alliance profit gives the minimum equilibrium profit for an alliance member. This is the profit an alliance member would earn if the other defender bid zero.

If $V \geq 4P$ defenses will be set such that the challenger drops out regardless of whether or not the defenders form an alliance. Because an independent defender would invest P and alliance members would invest at most P , defenders weakly prefer to form an alliance and per capita defender profits are strictly higher with an alliance.

⁸If the challenger was able to claim both prizes upon winning the contest the resulting sum of defender bids would be $\frac{V^2}{8P}$, which is still low compared to independent defense levels.

If $V \in [2P, 4P)$ then the challenger will end up placing a positive bid when facing an alliance but not when facing independently defended targets. In this range, the appropriate comparison is the profit of $V - P$ each defender receives from separate defenses and the $(1 + \frac{3V}{16P}) \frac{V}{2} - \frac{V^2}{32P}$ (which reduces to $\frac{V}{2} + \frac{V^2}{16P}$) that an alliance member who provided the entire defense investment expects to earn. The $-\frac{V^2}{32P}$ term captures the difference in profit from the symmetric equilibrium where a defender pays half of the $\frac{V^2}{16P}$ and the most inequitable equilibrium where the bidder pays the full $\frac{V^2}{16P}$. It is straightforward to show that $\frac{V}{2} + \frac{V^2}{16P} > V - P$ and thus the defenders would always prefer to form an alliance for values of V in this range.

Finally, for $V \in (0, 2P)$ the challenger will bid regardless of the defense structure. In this region, an alliance generates greater expected returns, even to a defender that fully finances the alliance if $\frac{V}{2} + \frac{V^2}{16P} > \frac{V^2}{4P}$, which holds when $V < 2P$.

Taking these comparisons together shows that defenders always prefer to form an alliance rather than engage in independent defenses.⁹ Because the formation of an alliance requires unanimity, both players opting to form an alliance and both opting to defend independently can occur in equilibrium. However, only both players opting to form an alliance satisfies the refinement of trembling hand perfection.

As it turns out, the challenger never strictly prefers to face independent defenders. In the case that $V < 2P$, this can be shown by comparing the challenger's expected profit against an independent defense, $(1 - \frac{V}{2P})^2 P$, and the challenger's expected profit against an alliance, $P(1 - \frac{V}{4P})^2$. The intuition is that under an alliance, the challenger faces a weaker defense and thus can obtain a high chance of winning at a relatively low cost. If $V \in [2P, 4P)$ then the challenger's expected profit is zero against independent defenders and is positive against alliance. Finally, when $V \geq 4P$ the challenger earns zero regardless of the defender's arrangement.

3 Experiment Design

Empirically, previous experiments have found that subjects in contests typically bid too aggressively. If this behavioral pattern applies similarly to both separate defenses and alliances, the comparative static predictions of the model should continue to hold. However, in non-contest settings, researchers have found strong evidence of altruism among in-group members. If the decision to form an alliance fosters this type of response, then alliances may bid even more

⁹Modifying the problem so that a successful challenge to an alliance results in an ambiguous preference ordering by the defenders, which depends on the relative sizes of V and P .

aggressively, which could reduce or eliminate the cost reduction associated with alliance membership.

To empirically test the predictions of the model, we conducted controlled laboratory experiments. To avoid influencing behavior, the experiments involved neutral language. No mention was made of challengers, defending, alliances, winning, etc. Instead, the task was framed as subjects bidding to claim two colored items. Defenders were identified as either Yellow or Blue and valued the item of the corresponding color at 256 (and valued the other item at 0). Challengers were identified as Green and valued both the yellow and the green item at 256.

Three experimental treatments were implemented: *Independent*, *Alliance*, and *Endogenous*. In the *Independent* treatment, Yellow and Blue (defenders) moved first and independently submitted bids for their respective items. In all cases bids were required to be non-negative and weakly less than the bidder's value of 256. Once the bids were submitted, Green (the challenger) observed the bids, chose an item on which to bid, and then placed a bid for the selected item. The item upon which Green did not bid was awarded to the defender who valued it. The allocation of the item upon which Green did bid was resolved via a proportional contest success function as described in Section 2, Case 1 with the winner receiving her value for the item. The results were revealed to all three participants and each person's profits were reduced by the amount of her bid and increased by the value of any item she claimed.

In the *Alliance* treatment, Yellow and Blue simultaneously submitted bids, knowing those bids would be combined into a single bid against Green. Green observed the bid by Yellow and Blue and then submitted her own bid. The outcome was determined using a proportional contest success function as described in Section 2, Case 2 above. If Green won the contest, Green was randomly assigned one of the items and the other was awarded to the defender who valued it.¹⁰ If Yellow and Blue won the contest, then both claimed their respective items. Regardless of the outcome, each participant had her bid deducted from her earnings and had the value of any claimed item added to her earnings.

The *Endogenous* treatment first presented Yellow and Blue with a binary choice to bid separately or to combine their bids. If both defenders opted to combine their bids then the experiment proceeded as in the *Alliance* treatment. Otherwise, the experiment proceeded as in the *Independent* treatment. Green knew that Yellow and Blue faced this choice and learned of the outcome before placing her bid. Because defender profits are higher under an *Alliance*, it

¹⁰It is possible that allowing the challenger to choose which target is claimed could impact behavior. For example, the challenger might be able to discourage defense investment by claiming the preferred target of the defender investing the most.

is expected that defenders will opt into the alliance when given the chance in the *Endogenous* treatment and thus the expected outcomes are the same for these two treatments.¹¹ Table 2 gives expected bids and profits by treatment. This table highlights the two changes in defense behavior associated with being in an alliance: the structure of the equilibrium changes from a mixed strategy to a pure strategy and defense investment falls dramatically. Notice that despite the multiple equilibria that exist in the alliance, there remains clear separation in predicted defender behavior between treatments. Specifically, bids in the interval (16, 64) by defenders should never be observed.

Table 2: Parameter Values and Expected Bids and Profits by Treatment

Parameters	
$V = P = 256$	
<i>Independent Defenses</i>	
$[\underline{b}, \bar{b}]$	[64, 192]
Expected Defender Bid	159.80
Expected Defender Bid Being Higher Bidding Defender	174.73
Expected Defender Bid Being Lower Bidding Defender	144.86
Expected Bid by Challenger	46.73
Chance of a Successful Attack	$\approx 25\%$
Expected Defender Profit	64
Expected Challenger Profit	17.67
<i>Alliance & Endogenous Defense</i>	
Optimal Defender Bid	8
Optimal Challenger Bid in Response to Optimal Defender Bid	48
Chance of a Successful Attack	$\approx 75\%$
Expected Defender Profit	152
Expected Challenger Profit	144

In each experimental session, subjects participated in 30 contests, 10 in each treatment. In half of the six sessions the treatment order was *Independent*, then *Alliance* and then *Endogenous*. In the other three sessions, the order of *Independent* and *Alliance* was reversed to control for ordering effects, but *Endogenous* was always implemented after subjects had familiarity with both defense methods so that their choice was informed.

The directions and the experiment were computerized using z-Tree (Fischbacher 2007). Subjects read treatment specific directions and answered comprehension questions just prior to participating in each segment of the experiment and did not know what if any other treatments would be implemented later in the session. Copies of the directions and comprehension ques-

¹¹Technically, there are two Pareto ranked Nash equilibria for the alliance formation game. Both opting to form the alliance and both opting not to for the alliance are equilibria; however, the players have a weakly dominant strategy to indicate a willingness to form the alliance.

tions are available in the Appendix.

When arriving at the lab, the twelve subjects in the session were seated at separate workstations isolated by privacy dividers. Subjects were then randomly assigned a color role that was maintained throughout the entire experiment.¹² However, each period subjects were randomly and anonymously rematched with other participants. This procedure eliminates the ability of subjects to build a reputation or engage in other repeated play strategies that might cause behavior to differ from the one-shot model described in Section 2.

The 69 participants were undergraduate students at the University of Arkansas recruited from the Behavioral Business Research Laboratory's subject pool.¹³ While some of the subjects had participated in other studies, none had participated in any related experiments. Subjects were paid in cash at the end of the approximately one hour experiment based upon their cumulative earnings. All of the values and bids in the experiment were denoted in Lab Dollars which were converted to \$US at the rate 250 Lab Dollars = 1 \$US. Because it is possible for subjects to lose money and in fact one of the three participants must lose money if they each place a positive bid, defenders were given an endowment of 750 while challengers were given an endowment of 1250. Asymmetric endowments were used for two reasons. First, challengers are involved in every contest while defenders are not. Second, with identical values for success in a contest, expected profits are greater for defenders in equilibrium. None of the subjects went bankrupt during the experiment. The salient earnings averaged \$18.98. Subjects also received an additional \$5 for participating.

4 Experimental Results

The model presented in Section 2 makes explicit predictions about how behavior should differ between the situation where the targets are defended independently and the situation where the defenders are in an alliance. The experimental results based on 210 contests are presented as a series of findings comparing what is observed in the lab with what is predicted by the model. Data from the *Independent* and *Alliance* treatments are used to evaluate bidding behavior while data from the *Endogenous* treatment is only used to determine defender preferences for forming alliances.¹⁴ Overall, the qualitative predictions of the model hold even though the

¹²All defenders viewed themselves as being Yellow and viewed the other defender as being Blue. This increased the number of different pairings that could occur.

¹³In one session only 9 subjects were present.

¹⁴This avoids issues of endogeneity when analyzing bidding behavior. However, behavior of defenders and challengers in the *Independent* and *Alliance* treatments is similar to behavior in the *Endogenous* treatment conditional on the

explicit quantitative predictions do not. Across all contests in the *Independent* treatment, the average bid was 95.5 while the average bid by a member of an alliance in the *Alliance* treatment was 62.6.¹⁵ The average bid by a challenger facing an independent defender was 107.3 and it was 98.4 when facing an alliance in their respective treatments. Challengers were successful in 57.4 percent of contests against independent defenders and in 44.8 percent of contests against alliances.

We begin with the behavior of defenders. Figure 2 shows a time series of the average defender bid by session across the *Independent* and *Alliance* treatments. Note that these data are from the first 20 periods of the experiment where the type of defense is exogenously determined. In the figure, sessions in which the subjects first experienced the *Independent* treatment are shifted to the right so that behavior in the same section of the figure is from the same treatment: the center section of the figure shows behavior in the *Independent* treatment which was periods 1-10 for some subjects and periods 11-20 for others. The shaded regions in the figure show investment levels that could be observed in equilibrium. Regardless of treatment order, defender behavior in the *Alliance* treatment is characterized by declining investments over time (see first and third sections of Figure 2), while defender behavior in the *Independent* treatment is generally flat (see middle section of Figure 2). Consistent with the predictions of the model, defender investments in the *Alliance* treatment are substantially below the defender investments in the *Independent* treatment. This provides the basis for Finding 1.

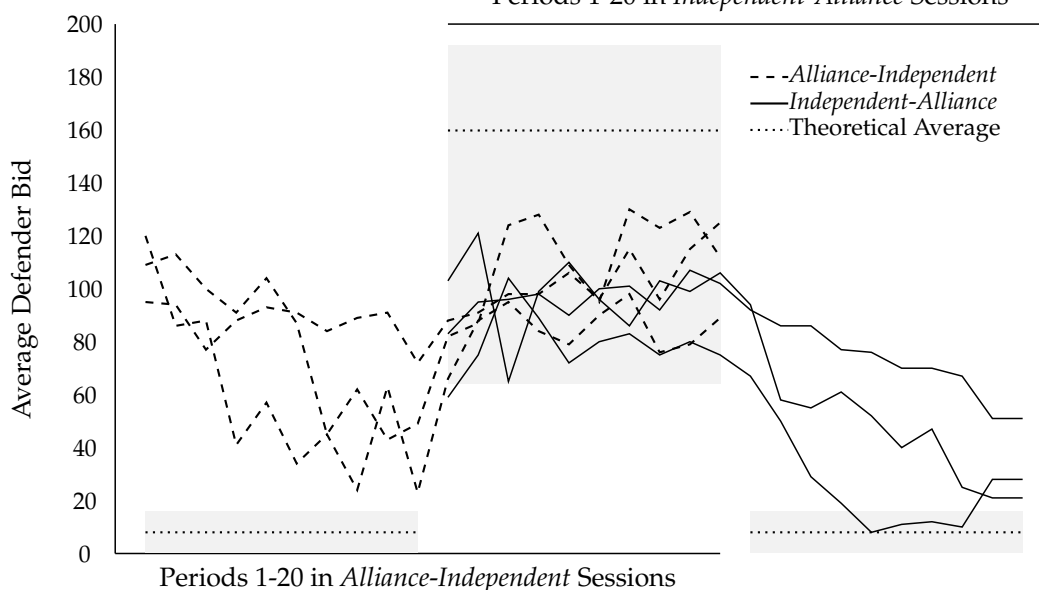
Finding 1: *Consistent with the theoretical predictions, defense investments are greater when the defenders are not in an alliance.*

This difference is statistically significant, as evidenced by the regression results presented in Table 3. For the first specification, the dependent variable is the investment by a defender. The explanatory variables are a constant, *AllianceDefense*, and *AllianceFirst*. *AllianceDefense* is a dummy variable that take a value of 1 if the observation was from a period in which the defender was in an alliance and a value of 0 otherwise. *AllianceFirst* is a dummy variable that take a value of 1 if the observation was from a session in which the defender experienced the *Alliance* treatment in periods 1 – 10 and a value of 0 otherwise. To handle the repeated measures in the data, standard errors are clustered at the session level. The treatment effect is captured

decision to defend independently or as an alliance, respectively.

¹⁵Within the *Endogenous* treatment, independent defenses had an average defender bid of 101.3 and allied defenses had an average defender bid of 45.5.

Figure 2: Average Defense Investment Each Period by Session, Separated by Treatment Order
 Periods 1-20 in *Independent-Alliance* Sessions



by the negative and significant value of *AllianceDefense*.¹⁶ The second specification in Table 3 is similar to the first except that, *Period*, a time trend variable is included. The interaction of *AllianceDefense* and *Period* allows for treatment specific trends. The results of this estimation suggest that while behavior does not differ between treatments initially, over time independent defenders are investing more while alliance members are investing less; that is, the predicted separation between treatments is becoming more pronounced over time.

Two additional features of Figure 2 and Table 3 are important. First, the distribution of investments in the *Alliance* treatment differs from the theoretically predicted (degenerate distribution at) 8. Instead, defenders invest an average of 62.6 in this treatment.¹⁷ This over-investment in defense is consistent with previous contest experiments. By contrast, in the *Independent* treatment, subjects are observed to under-invest. The observed average investment in this treatment was 95.5 while the predicted level was 159.8. These observations provide the basis of Findings 2 and 3.

¹⁶For robustness, we conducted the analysis clustering at the subject level. We also allowed for subject heterogeneity by regressing the treatment order on the difference in a subject's average bid in the two treatments. Finally, we constructed a between-subjects design by only considering data from the first treatment a subject experienced. All of these variations lead to the same conclusion - bids are significantly lower in the alliance treatment. We conducted similar analysis for challengers and find the results presented below for subjects in that role are also robust.

¹⁷The average behavior reported throughout the results comes directly from the raw data. But it can be calculated from the estimation in the tables up to rounding error. For instance, 62.6 is approximately $86.23 - 32.87 + 17.69/2$. The *AllianceFirst* coefficient is halved because *AllianceFirst* = 1 for half of the observations and 0 for the other half.

Table 3: OLS Estimate of Individual Defense Investment

Dependent Variable: Defender Bid		
	<i>Periods 1-20</i>	<i>Periods 1-20</i>
Constant	86.23*** (7.55)	77.58*** (8.16)
AllianceDefense	-32.87*** (6.07)	5.35 (9.82)
AllianceFirst	17.69 (11.60)	17.69 (11.61)
Period		1.57* (0.63)
Period × AllianceDefense		-6.95*** (1.23)
Observations	920	920
R ²	0.12	0.17
Subjects	46	46
SE Clusters	6	6

Notes: standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Finding 2: *When in an alliance, defenders over-invest.*

Finding 3: *When defending separately, defenders under-invest.*

The findings that investments are too high when defenders are in an alliance is supported statistically by testing $Constant + AllianceDefense = 8$ based on the first specification in Table 3, which can be rejected in favor of the alternative hypothesis that this sum is greater than 8 ($p\text{-value} < 0.001$). Finding 3 is supported statistically by testing $Constant = 159.8$, which can be rejected in favor of the alternative that it is less than 159.8 ($p\text{-value} < 0.001$).¹⁸

As a final point on defender behavior, we ask to what degree subjects in this role are responding to what they experienced in the previous period. Specifically, we consider *Other Defender's Bid*, the bid of the other defender from the previous period and *Challenger's Bid*, the bid of the challenger in the previous round. We also include dummy variables for each treatment. For the *Alliance* treatment, *Other Defender Lost* is a variable that takes the value one if the other defender did not claim their prize in the previous period's contest, and is zero otherwise. For

¹⁸This pattern could potentially be explained by subjects having an aversion to bidding near the extremes of the range of permissible bids.

the *Independent* treatment, *Attacked & Won* takes a value of one if the defender was challenged in the previous period and won, and is zero otherwise. The results are presented in Table 4. For independent defenders, larger bids from the other defender or the challenger lead to an increase in their bids next period. In the *Alliance* treatment, defenders only respond to the behavior of their alliance partner last period, which has a positive effect.

Table 4: OLS Estimate of Learning in Defense Investment

Dependent Variable: Defender Bid		
	<i>Alliance Treatment</i>	<i>Independent Treatment</i>
Constant	72.54** (23.30)	69.11*** (5.63)
AllianceFirst	-10.53 (18.24)	6.83 (7.19)
Period	-2.82** (1.08)	0.06 (0.59)
<i>Lagged 1 Period:</i>		
Other Defender's Bid	0.29*** (0.05)	0.35*** (0.05)
Challenger's Bid	0.05 (0.03)	0.08*** (0.02)
Other Defender Lost (1=Lost)	-6.98 (4.54)	
Attacked & Won (1=Won)		-5.11 (9.57)
Observations	414	414
R ²	0.23	0.15
Subjects	46	46
SE Clusters	6	6
Periods	20	20

Notes: standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

We now turn to the behavior of challengers. Challenger behavior is predicated on the actions of the defenders. When defense is done individually, challengers are expected to pursue the weaker defender's target. Indeed, this is the pattern that is observed as challengers opt to compete with weaker defender in 95.2% of the contests in *Independent* and 95.7% of the relevant contests in *Endogenous*. This is the evidence supporting Finding 4.

Finding 4: *When facing two independently defended targets, challengers overwhelmingly attack the weaker one, consistent with the theoretical predictions.*

The optimal response for challengers facing independent defenses is given by equation (2). The observed responses of challengers in the *Independent* treatment are given in panel (a) of Figure 3. The size of the markers in this figure denotes the relative frequency of the observation. As evidenced by panel (a) of Figure 3, challengers over-invest conditional on the level of defense, that is most of the observed responses lie above the best response curve. Further, challenger's tend to invest more in absolute terms than the defender as most responses actually lie above the 45 degree line.

For *Alliance*, the optimal challenger response is based on the total defense investment as shown in equation (3). Panel (b) of Figure 3 shows the observed challenger behavior for this treatment. As in *Independent*, challengers in *Alliance* tend to over-invest. However, in comparison to *Independent*, here challengers are more likely to have less than a 50% chance of a successful attack as they frequently bid less than the total level of defense. Of course, part of the explanation for the apparent difference in challenger behavior between the two treatments is the fact that the level of defense faced by the challenger was more likely to be large (above say 150) under *Alliance* and it is in this region where challengers are likely to have less than a 50% chance of success. Also evident from panel (b) of Figure 3 is that challengers do not give up when they should (i.e., when facing a defense that equals or exceeds 256), a result similar to Deck and Sheremeta (2012). These patterns provide the basis for finding 5.

Finding 5: *Consistent with previous experimental results, challengers over-invest regardless of the treatment.*

For econometric support of Finding 5 we offer Table 5, which is similar to the regression results presented above except that the dependent variable is the difference between the observed bid of the challenger and the optimal bid that the challenger should have made given the level of defense, $b_C - b_C^*$. In the first specification, over-investment is captured in the constant term, which is positive and significant. The lack of significance for *AllianceDefense* in the first specifications indicates that the level of overbidding does not differ by treatment in aggregate. However, when a time trend is added in specification 2, the results indicate that challenger overbidding is initially more severe in *AllianceDefense* but that this difference is diminished with experience.

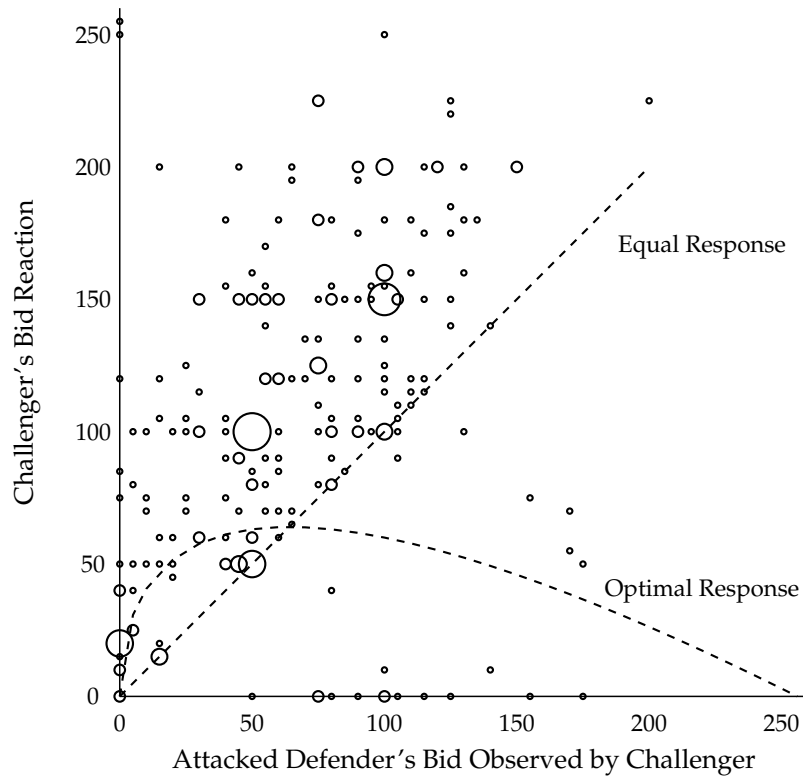
Given that challengers tend to over-invest and that the difference in defense investments, while significant, are not as dramatic as predicted, it remains to be determined if alliances are

Table 5: OLS Estimate of Challenger's Deviation from Optimal Investment

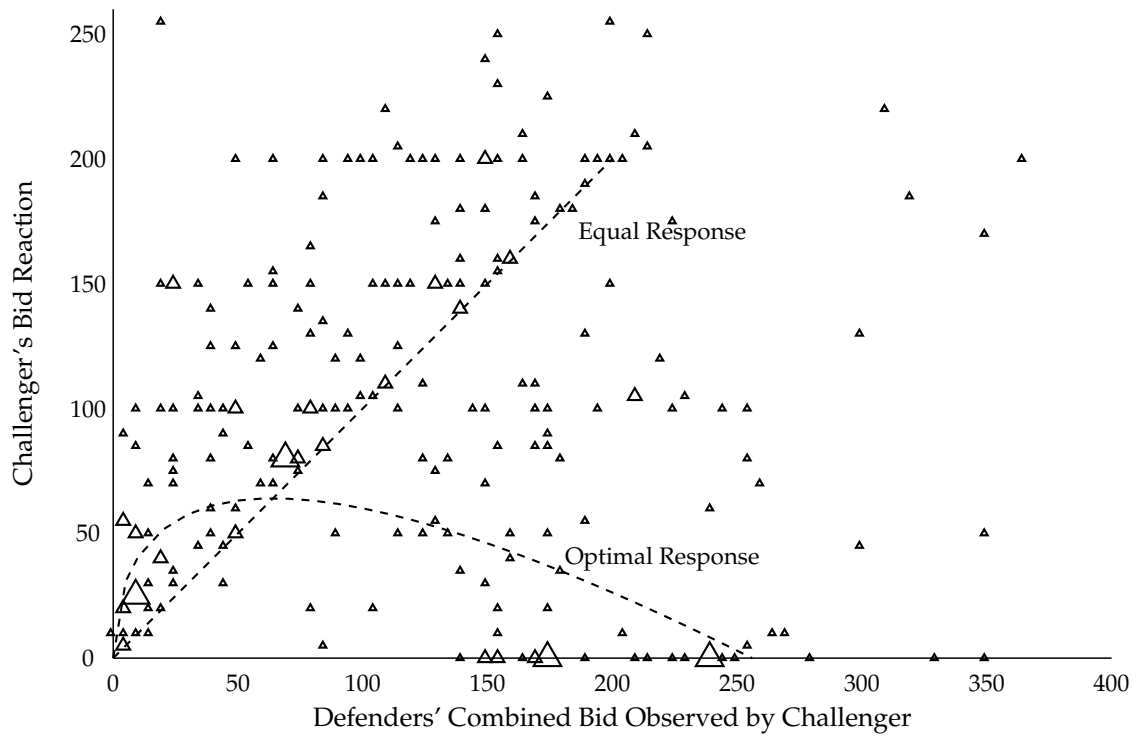
Dependent Variable: $b_C - b_C^*$		
	<i>Periods 1-20</i>	<i>Periods 1-20</i>
Constant	49.22** (16.30)	43.18** (16.01)
AllianceDefense	2.25 (7.27)	26.78* (12.54)
AllianceFirst	6.79 (16.25)	6.79 (16.28)
Period		1.10 (1.75)
Period \times AllianceDefense		-4.46* (2.06)
Observations	460	460
R ²	0.003	0.02
Subjects	23	23
SE Clusters	6	6

Notes: standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Figure 3: Challenger Responses Conditional on Defense Investment



(a) Independent Defense



(b) Alliance Defense

more profitable for and thus preferred by defenders.

The average defender profit in *Independent* was 87.08 while the average defender profit in *Alliance* was 136.08. This difference is significant, as supported by the regression results reported in the first two columns of Table 6. This estimation is similar to that reported above except that the dependent variable is defender profit. In the *Endogenous* treatment, 71.5% of the time when given a choice, defenders opted to form an alliance. Further, 30% of subjects attempted to join an alliance in every period, while only 17% preferred the independent defense a majority of the time. The combination of higher profits and expressed preference to form an alliance leads to Finding 6.

Table 6: OLS Estimate of Individual Profit

Dependent Variable:	Defender Profit		Challenger Profit	
	<i>Periods 1-20</i>	<i>Periods 1-20</i>	<i>Periods 1-20</i>	<i>Periods 1-20</i>
Constant	91.30*** (7.95)	93.64*** (5.72)	52.54** (16.92)	67.17** (17.72)
AllianceDefense	49.01*** (7.31)	44.74*** (9.71)	-23.40* (10.21)	-101.93*** (13.19)
AllianceFirst	-8.09 (8.26)	-8.09 (8.26)	-24.74 (24.31)	-24.74 (24.36)
Period		-0.43 (0.94)		-2.66 (1.68)
Period × AllianceDefense		0.77 (1.06)		14.28*** (1.99)
Observations	920	920	460	460
R ²	0.05	0.05	0.02	0.05
Subjects	46	46	23	23
SE Clusters	6	6	6	6

Notes: standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Finding 6: *Consistent with the theoretical model, defenders prefer to form alliances, which results in higher profits to defenders.*

To explore why some defenders sought to form an alliance while others did not, we conducted additional analysis where the dependent variable was the percentage of times a defender attempted to form an alliance. The results, omitted for brevity, indicate that realized profits from the first part of the experiment did not affect this choice, but bidding behavior did. In particular, those defenders that tended to bid more in *Alliance* were less likely to join an al-

liance in *Endogenous*. How one bid in *Independent* did not affect the chance that a defender tried to join an alliance.

Like defenders, challengers should fare better under alliances, but the average observed challenger profit was 39.63 in *Independent* and 16.23 in *Alliance*. This difference is significant as reported in columns 4 and 5 of Table 6 where the dependent variable is challenger profit. However, the specification in column 5 of Table 6 indicates that this difference is most pronounced in the early periods when defenders are still investing heavily. Also, challengers are predicted to be more successful when facing an alliance, but challengers were successful in 57.4% of attacks in *Independent* and in only 44.8% of attacks in *Alliance*. This difference is significant as evidenced by probit estimations with standard errors clustered at the session level (see Table 7) but again this result is being driven by the initial periods when defenders are just beginning to lower their investments in alliances. Perhaps with more experience, the predicted pattern would hold. These results are the basis of our next finding.

Finding 7: *Counter to the theoretical predictions, challengers are less successful in terms of expected profit and the likelihood of winning the contest when facing an alliance; however, this pattern diminishes with experience.*

Table 7: Probit Estimate of Attack Success

Dependent Variable: Challenger Success (1=Win)				
	Coefficients		Marginal Effects	
	Periods 1-20	Periods 1-20	Periods 1-20	Periods 1-20
Constant	0.29*** (0.09)	0.41*** (0.07)		
AllianceDefense	-0.32** (0.14)	-1.01*** (0.18)	-0.13** (0.06)	-0.39*** (0.06)
AllianceFirst	-0.19 (0.13)	-0.19 (0.13)	-0.07 (0.05)	-0.07 (0.05)
Period		-0.02 (0.02)		-0.01 (0.01)
Period × AllianceDefense		0.13*** (0.01)		0.04*** (0.004)
Observations	460	460		
Pseudo R ²	0.02	0.04		
Subjects	23	23		
SE Clusters	6	6		

Notes: standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Up to this point, we have focused on comparing behavior to the theoretical predictions developed in Section 2. We now ask what modifications to that theory can explain the observed pattern of behavior. Behavior consistent with risk aversion and loss aversion is commonly observed in a wide variety of lab experiments. The effect of risk aversion in contests is ambiguous unless one assumes a specific functional form (see Konrad and Schlesinger 1997). However, if subjects exhibit both risk aversion and prudence, a higher order risk property associated with the third derivative of the utility function, then challenger bids should be lower than the risk neutral prediction (see Treich 2009).¹⁹ Previous experimental studies have found that people identified as being more risk averse do tend to bid less in these types of contests (see Millner and Pratt, 1991). Other models of behavior under uncertainty such as prospect theory can lead challengers to bid more aggressively than predicted when facing a low level of defense; however, the level of defense required to encourage challengers with such preferences to drop out is lower as well, see Figure 4.²⁰ Thus, it does not appear that relaxing the risk neutrality assumption satisfactorily explains observed behavior.

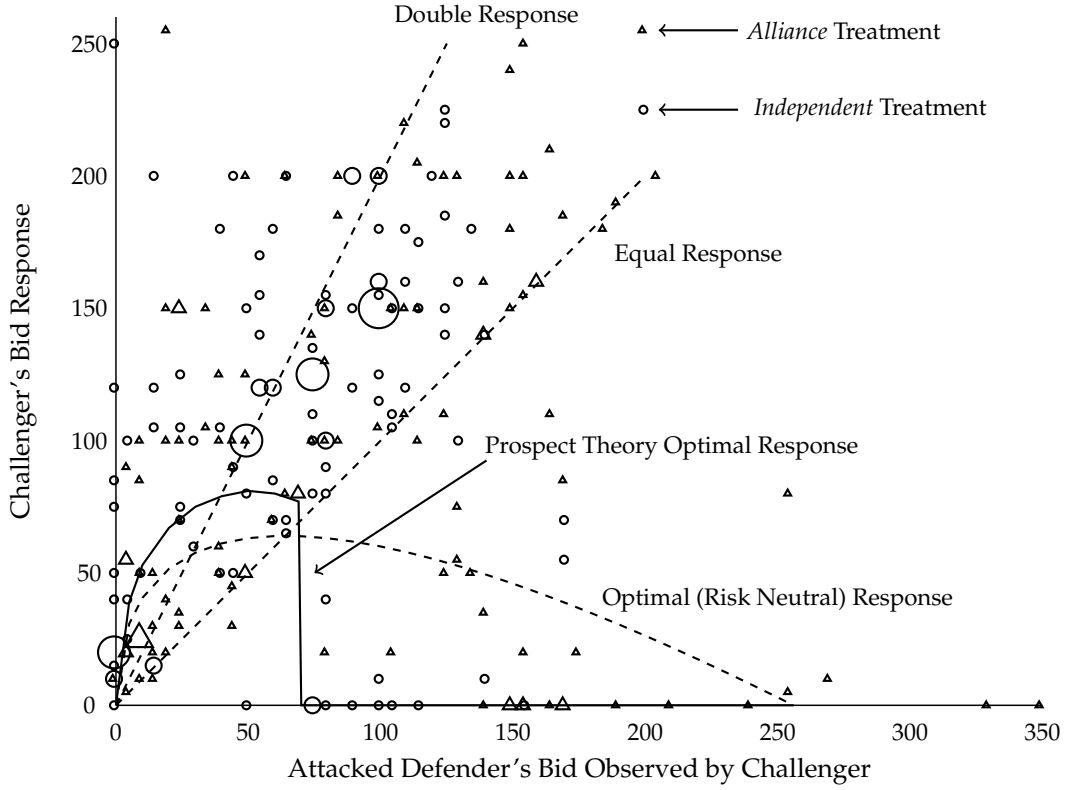
One explanation for overbidding that has received some success is that contestants receive a non-pecuniary, psychological benefit from winning the contest. Sheremeta (2010) reports that participants are willing to bid substantial amounts to anonymously and privately be declared the winner of a \$0 prize. If contestants have a utility of winning, then our analysis under-identifies the prize. An additive utility of winning denoted by U should lead to higher challenger bids as the response function in (2) would shift upwards to $\sqrt{(P+U)b_i} - b_i$. This means that each challenger's behavior should still fall along a curve similar to, but higher than, the ones shown in Figure 3, which does appear to be somewhat consistent with observed behavior. However, defender behavior in the independent defense offers evidence *against* a utility of winning explanation. With a utility of winning, the two defender profit functions shown in Figure 1 would shift out as shown in Figure 5. The result is that both \underline{b} and \bar{b} would shift to the right as $\frac{\partial b}{\partial V} = \frac{V}{2P} > 0$ and $\frac{\partial \bar{b}}{\partial V} = 1 - \frac{V}{2P} > 0$ when $V < 2P$. The new mixing distribution would first order stochastically dominate the original one. This is the opposite of the pattern we observe.

While risk and utility of winning do not appear to explain overbidding, inspection of Figure 4 suggests that at least for lower levels of defense, challenger behavior is approximated by a proportional rule to bid twice the level of defense. If defenders anticipate challengers will use

¹⁹The commonly assumed constant absolute risk aversion and constant relative risk aversion exhibit prudence.

²⁰The optimal response curve plotted to represent the prediction of prospect theory was determined by $\sum_{x \in X} p(x)|x|^\lambda(1 - \gamma\phi(x))$, where $X = \{v - b_C, -b_C\}$ and $\phi(x)$ is an indicator variable that equals 1 if $x \leq 0$ and 0 otherwise. To reflect the common result that losses are twice as costly as gains, we set $\gamma = -3$, and to depict moderate risk aversion we set $\lambda = 0.7$.

Figure 4: Challenger Responses in Last Five Periods of *Alliance* and *Independent* Treatments Conditioned on Defense Investment



a proportional rule, how should defenders behave? In the independent defense equation (4) becomes $\Pi_w = \frac{V}{(1+\beta)} - b_w$ under the assumptions that $b_C = \beta b_w$ and $b_w > 0$. The weak defender pushes her bid down as low as possible with the resulting security profit approaching $\frac{V}{1+\beta}$. Plugging this security profit in for the defender's profit in (5) along with $b_C = \beta b_w$ and solving yields the mixing distribution $F(b_i) = \frac{b_i(1+\beta)}{\beta V}$ over the interval $(\underline{b}, \bar{b}) = (0, \frac{\beta V}{1+\beta})$. Figure 6 plots this cumulative distribution function when challengers bid twice the defense level ($\beta = 2$) along with the observed cumulative distribution function from the experiment and the one defined in equation (7). The similarity between the observed behavior and that predicted for defenders who believe $b_C = 2b_w$ is striking. In an alliance, a proportional reaction by challengers would lead defenders to push their total bids toward zero, which is consistent with the pattern shown in Figure 2.

Figure 5: Effect of Utility of Winning for Independent Defense Behavior

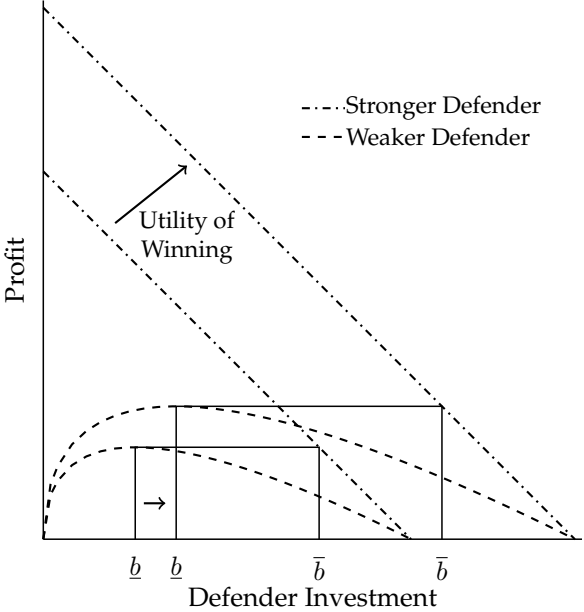
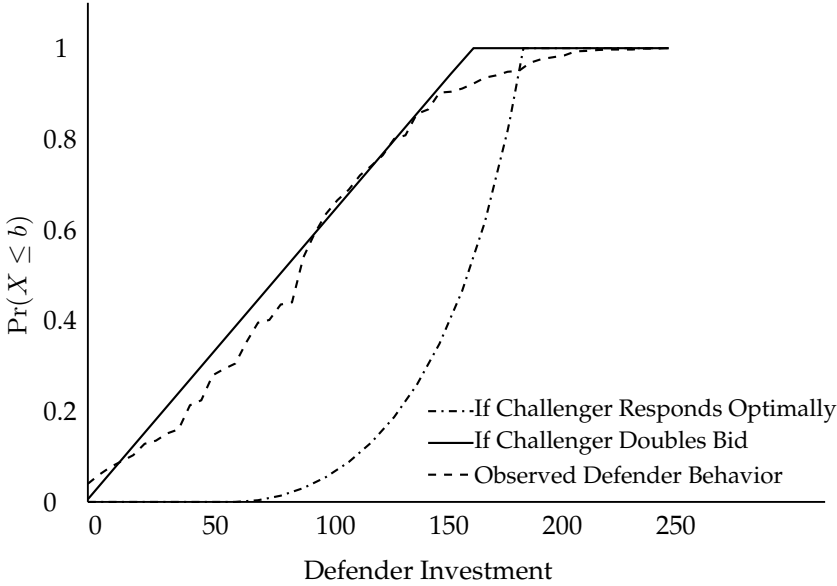


Figure 6: Cumulative Distribution of Independent Defender Investment



5 Conclusion

When defenders independently protect their own targets, the challenger's desire to focus on the weaker defender leads to an all-pay auction for the defenders. This leads defenders to invest heavily in their own defense. Unlike most of the all-pay auctions that have been studied previously, in this all-pay auction, the high bidder's profit net of the bid is fixed, but the low bidder's profit is a non-linear function of the bid. An alliance eliminates the need to outdo one's rival by avoiding the all-pay auction incentives and greatly reducing the average defense investment. Despite the fact that the challenger is theoretically more likely to be successful when facing an alliance, this loss is offset (in expectation) by the reduced defensive investment.

A series of controlled laboratory experiments largely confirms the qualitative predictions of the model: when selecting between two separately defended targets, challengers contest the weaker one; and defenders invest less and earn more in an alliance. Consistent with behavior in previous simultaneous contest experiments, in these sequential contests second mover challengers are observed to over-invest. In our sequential setting, challengers also fail to give up when it is optimal to do so. Over-investment relative to the theoretical prediction is also observed for defenders in alliances; however, that behavior diminishes with experience. In contrast, separate defenders under-invest relative to the theoretical prediction. This behavior is inconsistent with the utility of winning hypothesis that is commonly used to explain over-bidding. However, the investment pattern exhibited by independent defenders is optimal if defenders believe challengers will bid twice the observed level of defense, a heuristic that is not too far off of what challengers typically do.

While behavior does not match the point estimates predicted by the model, the general consistency we find between the theoretical and behavioral comparative static effects of changing from independent defenses to alliances is encouraging. It suggests that this framework is reasonable for exploring more complicated and realistic scenarios such as multiple challengers who coordinate their actions, larger numbers of defenders who can invest in defending both the alliance and their own target, or allowing defenders to choose whether or not the level of defense is observable.²¹ At the same time our results demonstrate two important behavioral properties. First, there is behavioral evidence that overbidding is driven by something other than a utility of winning - an explanation that warrants further investigation as over-bidding is

²¹If defense levels are not observable then the game is effectively simultaneous. Assuming that the challenger selects which contest to enter randomly then the case of independent defense collapses to a standard two player Tullock contest where the defender's prize is $\frac{V}{2}$ and the challenger's prize is P . An alliance would collapse to a standard three player Tullock contest with these same prize values.

one of the best documented aspects of contest behavior in the laboratory. Second, contestants appear to use the heuristic that others will use a proportional bidding rule, which has broad implications for contest design and implementation in a wide variety of settings.

References

- [1] Abbink, Klaus, Jordi Brandts, and Benedikt Herrmann et al. (2010): "Intergroup Conflict and Intra-Group Punishment in an Experimental Contest Game," *American Economic Review*, 100(1): 420-447.
- [2] Azaiez, Mohamed. and Vicki Bier (2007): "Optimal Resource Allocation for Security in Reliability Systems," *European Journal of Operational Research*, 181(2): 773-786.
- [3] Baik, Kyung Hwan (2008): "Contest with Group-Specific Public-Good Prizes," *Social Choice and Welfare*, 30(1): 103-117.
- [4] Baumeister, Roy and Mark Leary (1995): "The Need to Belong: Desire for Interpersonal Attachments as a Fundamental Human Motivation," *Psychological Bulletin*, 117(3): 497-529.
- [5] Baye, Michael, Dan Kovenock, and Casper de Vries (1996): "The All-Pay Auction with Complete Information," *Economic Theory*, 8(2): 291-305.
- [6] Baye, Michael, Dan Kovenock, and Casper de Vries (2012): "Contests with Rank-Order Spillovers," *Economic Theory*, 51(2): 315-350.
- [7] Bier, Vicki and Vinod Abhichandani (2002): "Optimal Allocation of Resources for Defense of Simple Series and Parallel Systems from Determined Adversaries," In: Proceedings of the Engineering Foundation Conference on Risk-Based Decision Making in Water Resources X. American Society of Civil Engineers, Santa Barbara, CA.
- [8] Bier, Vicki, Aniruddha Nagaraj and Vinod Abhichandani (2005): "Protection of Simple Series and Parallel Systems with Components of Different Values," *Reliability Engineering and System Safety*, 87: 315-323.
- [9] Clark, Derek and Christian Riis (1998): "Contest Success Functions: An Extension" *Economic Theory*, 11(1):201-204.
- [10] Davis, Douglas and Robert Reilly (1998): "Do Too Many Cooks Always Spoil the Stew? An Experimental Analysis of Rent-seeking and the Role of a Strategic Buyer," *Public Choice*, 95 (1-2): 89-115.
- [11] Deck, Cary and Roman Sheremeta (2012): "Fight or Flight? Defending Against Sequential Attacks in the Game of Siege," *Journal of Conflict Resolution*, 56 (6): 1069-1088.

- [12] Dighe, Nikhil, Jun Zhuang and Vicki Bier (2009): "Secrecy in Defensive Allocations as a Strategy for Achieving More Cost-effective Attacker Deterrence," *International Journal of Performability Engineering*, 5(1): 31-43.
- [13] Esteban, Joan and Jozsef Sakovics (2003): "Olson vs. Coase: Coalitional Worth in Conflict," *Theory and Decision*, 55(4): 339-357.
- [14] Fischbacher, Urs (2007): "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10(2): 171-178.
- [15] Garfinkel, Michelle R. (2004): "Stable Alliance Formation in Distributional Conflict," *European Journal of Political Economy*, 20(4): 829-852.
- [16] Gneezy, Uri and Rann Smorodinsky (2006): "All-pay Auction - An Experimental Study," *Journal of Economic Behavior and Organization*, 61(2): 255-275.
- [17] Hausken, Kjell (2008): "Strategic Defense and Attack for Series and Parallel Reliability Systems," *European Journal of Operational Research*, 186(2): 856-881.
- [18] Hausken, Kjell and Vicki Bier (2011): "Defending Against Multiple Different Attackers," *European Journal of Operational Research*, 211(2): 370-384.
- [19] Katz, Eliakim and Julia Tokatlidu (1996): "Group Competitions for Rents," *European Journal of Political Economy*, 12(4): 599-607.
- [20] Ke, Changxia, Kai Konrad and Florian Morath (2010): "Brothers in Arms - An Experiment on the Alliance Puzzle," CESifo Working Paper Series No. 3302. Available at SSRN: <http://ssrn.com/abstract=1742762>
- [21] Konrad, Kai (2004): "Bidding in Hierarchies," *European Economic Review*, 48(6): 1301-1308.
- [22] Konrad, Kai and Harris Schlesinger (1997): "Risk Aversion in Rent-seeking and Rent-augmenting games," *Economic Journal*, 107: 1671-1683.
- [23] Kovenock, Dan and Brian Roberson (2012a): "Strategic Defense and Attack for Series and Parallel Reliability Systems: Comment," *Defense and Peace Economics*, 23(5): 507-515.
- [24] Kovenock, Dan and Brian Roberson (2012b): "Strategic Defense and Attack for Series and Parallel Reliability Systems: Rejoinder," *Defense and Peace Economics*, 23(5): 521-524.

- [25] Kunreuther, Howard and Geoffrey Heal (2003): "Interdependent Security," *The Journal of Risk and Uncertainty*, 26(2/3): 231–249.
- [26] Lee, Dongryul (2012): "Weakest-link contests with group-specific public good prizes," *European Journal of Political Economy*, 28(2): 238-248.
- [27] Levitin, Gregory (2003a): "Optimal Multilevel Protection in Series-Parallel Systems," *Reliability Engineering and System Safety*, 81: 93–102.
- [28] Levitin, Gregory (2003b): "Optimal Allocation of Multi-State Elements in Linear Consecutively Connected Systems with Vulnerable Nodes," *European Journal of Operational Research*, 150: 406–419.
- [29] Lugovskyy, Volodymyr and Daniela Puzzello (2008): "An Experimental Investigation of Overdissipation in the All Pay Auction," *European Economic Review*, 54(8): 974-997.
- [30] Major, John (2002): "Advanced Techniques for Modeling Terrorism Risk," *Journal of Risk Finance*, 4(1): 15–24.
- [31] Milner E.L. and M.D. Pratt (1989): "An Experimental Investigation of Efficient Rent-seeking," *Public Choice*, 62: 139-151.
- [32] Muenster, Johannes (2009): "Group Contest Success Functions," *Economic Theory*, 41(2): 345-357.
- [33] Muller, Holger and Karl Warneryd (2001): "Inside Versus Outside Ownership: A Political Theory of the Firm," *RAND Journal of Economics*, 32(3): 527-541.
- [34] Nitzan, Shmuel and Kaoru Ueda (2008): "Collective Contests for Commons and Club Goods," *Journal of Public Economics*, 93(1-2): 48-55.
- [35] Noussair, Charles and Jonathon Silver (2006): "Behavior in All-pay Auctions with Incomplete Information," *Games and Economic Behavior*, 55(1): 189-206.
- [36] Nti, Kofi (1998): "Effort and Performance in Group Contests," *European Journal of Political Economy*, 14(4): 769-781.
- [37] O'Hanlon, Michael, Peter Orszag, Ivo Daalder, I.M. Destler, David Gunter, James Lindsay, Robert Litan, and James Steinberg (2002): "Protecting the American Homeland," Brookings Institution, Washington, DC.

- [38] Olson, Mancur and Richard Zeckhauser (1966): "An Economic Theory of Alliances," *The Review of Economics and Statistics*, 48(3): 266-280.
- [39] Potters, Jan, Casper de Vries and Frans van Winden (1998): "An Experimental Examination of Rational Rent-seeking," *European Journal of Political Economy*, 14(4): 783-800.
- [40] Ryvkin, Dmitry (2010): "The Selection Efficiency of Tournaments," *European Journal of Operational Research*, 206(3): 667-675.
- [41] Ryvkin, Dmitry (2011): "The Optimal Sorting of Players in Contest between Groups," *Games and Economic Behavior*, 73(2): 564-572.
- [42] Sheremeta, Roman, Emmanuel Dechenaux and Dan Kovenock (2012): "A Survey of Experimental Research on Contests, All-Pay Auctions and Tournaments," Economic Science Institute Working Paper.
- [43] Sheremeta, Roman (2010): "Experimental Comparison of Multi-Stage and One-Stage Contests," *Games and Economic Behavior*, 68: 731-747.
- [44] Sheremeta, Roman and Jingjing Zhang (2010): "Can Groups Solve the Problem of Overbidding in Contests?" *Social Choice and Welfare*, 35: 175-197.
- [45] Skaperdas, Stergios (1996): "Contest Success Functions" *Economic Theory*, 7(2): 283-290.
- [46] Spolaore, Enrico (2010): "Federalism, Regional Redistribution and Country Stability," In Bosch, Nuria, Marta Espasa and Albert Sole Olla, *Political Economy of Inter-Regional Fiscal Flows: Measurement Determinants and Effects on Country Stability (Studies in Fiscal Federalism and State-local Finance)*. Edward Elgar Publishing.
- [47] Treich, Nicolas (2010): "Risk-aversion and Prudence in Rent-seeking Games," *Public Choice*, 145(3-4): 339-349.
- [48] Tullock, Gordon (1980): "Efficient Rent Seeking," In Buchanan, James, Robert Tollison and Gordon Tullock, *Toward a Theory of the Rent-seeking Society*. Texas A& M University Press, College Station, TX.
- [49] Warneryd, Karl (1998): "Distributional Conflict and Jurisdictional Organization," *Journal of Public Economics*, 69(3): 435-450.

[50] Woo, Gordon (2002): "Quantitative Terrorism Risk Assessment," *Journal of Risk Finance*, 4(1): 7-14.

[51] Woo, Gordon (2003): "Insuring Against Al-Qaeda," Insurance Project Workshop, National Bureau of Economic Research, Inc. Available at <http://www.nber.org/confer/2003/insurance03/woo.pdf>.

6 Appendix: Subject Instructions

(Page 1)

Introduction

This is an experiment in the economics of decision making. In addition to the \$5 dollars you will receive for participating today, you have the opportunity to earn additional money. No person in the experiment (besides you) will know the decisions you make, and you will not be told the decisions of any other specific individual.

The experiment consists of three parts. At the end of the experiment you will be paid privately in cash for your total earnings in the entire experiment. However, the decisions you make in one part of the experiment will not impact any other part of the experiment. All amounts of money in the experiment are in Lab Dollars. At the end of the experiment, your Lab Dollars will be converted into \$US at the rate 250 Lab Dollars = 1 \$US. You will begin the experiment with 1250 (750) Lab Dollars. Any losses you incur during the experiment will be deducted from your Lab Dollars.

We will now walk through the instructions for part 1. Because the amount of money you will receive will depend upon the decisions you make, it is important that you understand the instructions completely. If you have a question at any point, please raise your hand and an experimenter will come to answer it. Otherwise, you should not talk or communicate with anyone else during this experiment.

(Page 2)

Each part of the experiment involves a series of decision periods. Each period, you will be randomly shuffled into a group of 3 people. There will be 3 types of decision makers in each group: Yellow, Blue and Green. You have been randomly assigned the role of Green (Yellow) and will remain in that role for the entire experiment.

Each period there are two items available to be claimed: a yellow item and a blue item, shown as colored boxes on your screen. There is only one item of each color and a decision maker can claim at most one item in a period. Yellow decision makers value the yellow item at \$256 but have no value for the blue item. Blue decision makers value the blue item at \$256 but have no value for the yellow item. Green decision makers are indifferent between the two items and value each at \$256.

If you claim an item, your earnings will increase by your value for the item. Since there are three decision makers and only two colored items, this means someone will not get an item.

What changes in each part of the experiment is how you claim items.

(Page 3 or 7: Independent Treatment)

How do I claim an item in this part of the experiment?

In this part of the experiment, you can try to claim a blue or yellow (yellow) by bidding on it. The other decision makers can also try to claim items by bidding on them.

If only one decision maker bids on an item, that decision maker will claim the item. If two decision makers bid on the same item, then who claims the item will depend in part on how much each decision maker bid and in part on chance. The larger your bid, the more likely it is that you will claim an item. However, each bidder must pay whatever amount he or she bid regardless of whether or not he or she actually claims an item or was the only one bidding for it.

(Page 4 or 8: Independent Treatment)

So how does bidding work?

The bidding process for this part of the experiment is as follows. First, Yellow and Blue will privately choose an amount to bid for the item of their respective color. Because Yellow and Blue only value the item in their own color, these decision makers can only bid on that item. They will place bids by typing their bid amounts in their separate boxes on their respective screens and pressing the Bid button. These bids must be a number from 0 to 256, because no one should be willing to bid more than their value for the item.

After Yellow and Blue bid, Green will then observe how much Yellow and Blue actually bid for the two items. Green will then choose one (and only one) of the two items on which to bid. Buttons will appear beside the two items on Green's screen. Green will select which item to bid on by clicking the button beside the item he or she wishes to bid on. After selecting which item to bid on, Green will then choose a bid from 0 to 256 by entering this amount in his or her box and pressing the bid button.

Because Green can only bid for one of the two items, this means either Blue or Yellow will be the only one bidding on the item Green does not bid on. The item that Green does not bid on automatically goes to the one decision maker who did bid on it (keep in mind that this decision maker still has to pay the amount of his or her bid). This will be denoted on your screen with a black arrow from the item to the decision maker that claimed it.

Let's look at an example: Suppose Yellow bids 30 and Blue bids 60.

If Green chooses to bid on Yellow, then Blue automatically receives a payoff of $256 - 60 = \$196$. If Green chooses to bid on Blue, then Yellow automatically receives a payoff of $256 - 30 = \$226$.

(Page 5 or 9: Independent Treatment)

What happens to the item for which two decision makers bid?

Who claims the item is determined as follows. The chance that the bidder will claim the item is equal to $own\ bid / (own\ bid + other's\ bid)$. This means that the chance that the bidder will not claim the prize is equal to $other's\ bid / (own\ bid + other's\ bid)$. With this proportional formula, the more a decision maker bids the more likely that person is to claim an item.

Continuing the example from before, suppose Yellow bids 30, Blue bids 60, and Green bids 30. If Green chooses to bid for the yellow item then the chance that Green claims the item is $30 / (30 + 30) = 0.5$ or 50% and the chance that Yellow claims the item is $30 / (30 + 30) = 0.5$ or 50%.

In this case, Blue receives $256 - 60 = \$196$. If Green claim the yellow item: Green receives $256 - 30 = \$226$ and Yellow receives $\$-30$. If Yellow claims the yellow item: Green receives $\$-30$ and Yellow receives $256 - 30 = \$226$.

However, if Green chose to bid for the blue item then the chance that Green claims the item is $30 / (30 + 60) = 0.33$ or 33% and the chance that Blue claims the item is $60 / (30 + 60) = 0.67$ or 67%.

In this case, Yellow receives $256 - 30 = \$226$. If Green claims the blue item: Green receives $256 - 30 = \$226$ and Blue receives $\$-60$. If Blue claims the blue item: Green receives $\$-30$ and Blue receives $256 - 60 = \$196$.

(Page 6 or 10: Independent Treatment)

After everyone has bid, a bar will appear on the left side of your screen to show you the chance that each decision maker will not claim the item he or she bid on. The total height of the bar represents 100% and each color segment denotes the chance that decision maker will be the one who does not claim an item. Notice that either Yellow or Blue will not appear in this bar because one of them is guaranteed to claim the item Green does not bid on. The computer will randomly place an X somewhere on the bar to determine who does not get to claim an item. A black arrow on your screen will indicate who claims the second item.

We are now ready to begin this part of the experiment. Keep in mind that this part of the experiment will last several periods and that you will be randomly shuffled into a group of 3 people each period. If you have any questions you'd like to ask before the experiment starts, please ask them now. Otherwise, press the button below that says BEGIN.

(Page 3 or 7: Alliance Treatment)

How do I claim an item in this part of the experiment?

In this part of the experiment, you can try to claim a by bidding on it. The other decision makers can also try to claim items by bidding on them, but the bids of the Yellow and Blue decision makers will be combined.

Who claims the item will depend in part on how much each decision maker bid and in part on chance. The larger your bid, the more likely it is that you will claim an item. However, each bidder must pay whatever amount he or she bid regardless of whether or not he or she actually claims an item.

(Page 4 or 8: Alliance Treatment)

So how does bidding work?

The bidding process for this part of the experiment is as follow. First, Yellow and Blue will privately choose an amount to bid for the item of their respective color. They will place bids by typing their bid amounts in their joint box on their respective screens and pressing the Bid button. These bids must be a number from 0 to 256, because no one should be willing to bid more than their value for the item. The amount that Blue and Yellow bid will be added together to become the Combined Bid.

After Yellow and Blue bid, Green will then observe the Combined Bid. Green will then choose a bid from 0 to 256 by entering this amount in his or her box and pressing the bid button.

The chance that Green claims an item is $\text{Green's Bid}/(\text{Combined Bid} + \text{Green's Bid})$.

If Green claims an item, then Green will be randomly assigned to claim either the yellow item or the blue item because Green values them equally. Notice that Green can claim only one item leaving the other item to be automatically claimed by the other decision maker who values it. So if Green claims the blue item then Yellow would claim the yellow item.

If Green does not claim an item, then Yellow and Blue both claim the item they value. The chance that Green does not get to claim an item is $\text{Combined Bid}/(\text{Combined Bid} + \text{Green's Bid})$.

With this proportional formula, the more a decision maker bids the more likely that person is to claim an item.

(Page 5 or 9: Alliance Treatment)

Notice that there are two ways that Blue or Yellow can claim an item. One way is if Green does not get to claim an item and the other is if Green gets to claim an item, but claims the item that is not valued. This means that there is only one way that Blue does not get to claim an item, which is Green gets to claim an item and it happens to be the blue one. Therefore, the chance that Blue does not get to claim an item is $\frac{1}{2} \frac{\text{Green's Bid}}{\text{Combined Bid} + \text{Green's Bid}}$. Similarly, Yellow has the same

chance of not getting to claim an item as Blue has.

Let's look at an example: Suppose Yellow bids 30, Blue bids 60, and Green bids 30. The Combined Bid would be Yellow's bid plus Blue's bid, which is $30 + 60 = 90$. Therefore, the chance that Green does not get to claim an item (and thus that Yellow and Blue both get to claim an item) is $90/(90 + 30) = 0.75$ or 75%. The chance that Green does get to claim an item is $30/(90 + 30) = 0.25$ or 25%. This means that the chance that Blue does not get to claim an item is $\frac{1}{2} \cdot 25\% = 12.5\%$ and the chance that Yellow does not get to claim an item is also 12.5%.

If Green does not claim an item then: Green receives \$-30. Yellow $256 - 30 = \$226$. Blue receives $256 - 60 = \$196$.

If Green does get to claim an item then: Green receives $256 - 30 = \$226$. If Green claims the Blue item then Blue receives \$-60 and Yellow receives $256 - 30 = \$226$. But, if Green claims the Yellow item then Blue receives $256 - 60 = \$196$ and Yellow receives \$-30.

(Page 6 or 10: Alliance Treatment)

After everyone has bid, a bar will appear on the left side of your screen to show you the chance that each decision maker will not claim the item he or she values. The total height of the bar represents 100% and each color segment denotes the chance that decision maker will be the one who does not claim an item. The computer will randomly place an X somewhere on the bar to determine who does not get to claim an item. Black arrows on your screen will indicate who claims each of the items.

We are now ready to begin this part of the experiment. Keep in mind that this part of the experiment will last several periods and that you will be randomly shuffled into a group of 3 people each period. If you have any questions you'd like to ask before the experiment starts, please ask them now. Otherwise, press the button below that says BEGIN.

(Page 11: Endogenous Treatment)

How do I claim an item in this part of the experiment?

How items are claimed depends on Yellow and Blue. In the first two parts of the experiment bids by Yellow and Blue were required to be separate or required to be combined. In this part of the experiment, Blue and Yellow can choose to bid separately or have their bids combined.

Buttons will appear on the screens of Blue and Yellow decision makers asking which process they would like to use for claiming regions that period. If both Yellow and Blue opt to have their bids combined, then that process will be implemented. If either or both Yellow and Blue opt to bid separately, then that process will be implemented.

Once the bidding process is determined, the period will progress accordingly following the

same sequence as in the corresponding previous part of the experiment. We are now ready to begin this part of the experiment. Keep in mind that this part of the experiment will last several periods and that you will be randomly shuffled into a group of 3 people each period. If you have any questions you'd like to ask before the experiment starts, please ask them now. Otherwise, press the button below that says BEGIN.