Trust on the Brink of Armageddon: The First-Strike Game

KLAUS ABBINK, THOMAS DE HAAN[1]

We introduce the experimental *first-strike* game. Two players accumulate earnings over many rounds. In any round a player can deactivate the opponent. This means that the opponent loses all earnings and faces severely reduced future payoffs. It also means that the opponent can no longer deactivate the aggressor in later rounds. Thus deactivation is a hostile act, but can serve as a self-protection device. We run four treatments to disentangle the effect of spite, fear of spite, and trust. We find the absence of pure nastiness, but nevertheless a strong fear of it. Mutual trust can considerably reduce hostility.

"If you say why not bomb them tomorrow, I say why not today? If you say why not today five o'clock, I say why not one o'clock?"

JOHN VON NEUMANN (quoted in Life Magazine (1957)).

1. Introduction

In his classic essay Schelling (1960) uses the example of a homeowner confronting a burglar, both carrying a gun, to illustrate how uncertainty about the opposition's motives can trigger unwanted attacks. If the homeowner assigns some probability to the burglar planning to shoot, then it might become the best response to strike first with a surprise attack. Suspecting this, the burglar might himself be tempted to attack. So a situation can arise in which both parties attack although neither has an actual desire to kill. Mutual fear of (possibly even irrational) aggression is sufficient to prompt a lethal exchange of hostilities.

The nuclear standoff of the two superpowers during the Cold War is an ultra-high-stakes version of Schelling's example. In the early stages of the confrontation a preemptive nuclear strike against the Soviet Union was a frighteningly popular idea within the US military (Gilpin (1963)), despite the millions of deaths such a strike would have caused. The introductory quote by a champion of strategic rationality shows that such contemplations were not considered irrational at the time. Only later, when the Eastern bloc had caught up in building up its nuclear arsenal, the idea of a first strike lost most of its persuasive force. It became clear that it would be unlikely to score a hit so decisive that the Soviet Union could not land a devastating second strike in retaliation (Wohlstetter (1958)).

These scenarios show that fear itself can be extremely dangerous. Situations of bilateral threat are characterized by layers of fear, which even reinforces the potential for damage. There is the direct fear that the opposition has the desire to hurt. This could immediately trigger a preemptive strike. We call this *first-order fear*. But even if I do not fear that the opponent is malevolent, I may fear that he fears that I am, and thus strikes against me. The opponent, however, may fear that I fear that he does that, and therefore strike first. This fear-of-fear-of-fear reasoning can be extended ad infinitum, we shall accumulate these levels of fear under the term *higher-order fear*.

---

[1] Abbink: CBESS, University of East Anglia, Norwich NR4 7TJ, United Kingdom. k.abbink@uea.ac.uk. De Haan: CREED, Faculty of Economics and Business, University of Amsterdam, Roeterstraat 11, 1018 WB Amsterdam, The Netherlands, t.dehaan@uva.nl. Financial support from the University of Amsterdam through the *Speerpunt Behavioural Economic*, and the University of East Anglia through the CBESS Research Fund is gratefully acknowledged.

Fortunately for us, the Cold War ended without the obliteration of entire nations. One reason could be that the two opponents were tied in a long-term relationship. Despite being bitter rivals some form of trust could emerge between them, or at least some notion of predictability from the common history. This is an important difference to Schelling's burglary example in which the homeowner and the burglar are strangers encountering each other for the first time. The latter have no common history in which they could have built up trust, and they have no common future for which to signal trust and hope for reciprocity.

In this paper we study the role of spite, fear and trust in a simple experimental game. We devise a new experimental paradigm we call the *first-strike game*. This game can be seen as a dynamic multi-stage variation of Schelling's idea, or an idealized toy version of the Cold War standoff. Two players each perform individual tasks over several rounds to accumulate earnings. In any round a player can decide to attack ("deactivate") the other player. Deactivation has four consequences. First, almost all earnings the opponent has accumulated are destroyed. Second, the opponent's future earnings capabilities are reduced to a small fraction. Third, the deactivator bears a cost, making it materially unattractive to attack. Fourth, and most importantly, the target cannot deactivate the attacker anymore in later rounds. Although there is no direct material benefit from deactivation, it can serve as a self-protection instrument, since an attacker who strikes first is safe from future attacks. In four treatments, described in the following section, we separate motives of fear and trust from one another.

Our experimental game is inspired by the Cold War standoff, but it is not meant to mimic it exactly. Obviously, the two operate on very different scales. Our main goal is to introduce a framework in which spite, fear and trust can be studied in situations of conflict. Though the present study is, to the best of our knowledge, the first to tackle this research agenda, it does not stand in isolation. Spite, or nastiness, which we define as a pleasure derived from lowering the well-being of others, is only sparsely studied in the experimental literature, in contrast to its nice counterpart, altruism. However, there is evidence that people are willing to pay money to reduce other people's income (Zizzo and Oswald (2001), Zizzo (2004)), even in absence of conventional motives like inequity aversion or reciprocity (Abbink and Sadrieh (2009), Abbink and Herrmann (2010)). Trust has been studied extensively in the context of gift exchange and investment games (Fehr et al (1993), Berg et al (1995)). In these games people frequently trust others to return favors even though it is not in the reciprocator's material self-interest. This notion of trust and reciprocity is of course fundamentally different from the one in the first strike standoff. In trust games the trustor has to rely on the reciprocator's willingness to reward kindness, in the first strike game trust is about the absence of malicious intent and fear of such.

Our game also bears some resemblance to other settings studied in the experimental literature. The orders of fear are reminiscent of the levels of reasoning as studied in beauty contest games (Nagel (1995, 1998), Großkopf and Nagel (2008)), auctions (Crawford and Iriberri (2007a)), and hide-and-seek games (Crawford and Iriberri (2007b)). The so-called level-k reasoning assumes some random behavior as level zero, and as level k the best response to level k-1. In contrast to the cascade of fear in the first-strike setting, this is a merely cognitive process. Similar processes of reasoning are at work in centipede games (McKelvey and Palfrey (1992)). In the centipede game players take turns to either pass or end the game. Passing increases both players' payoffs, but the player who ends earns more than the other player, and more than he would if the other player ended next turn. By this construction the subgame perfect equilibrium is for the first mover to end the game immediately, which is inefficient. In the first strike game, the universally best solution coincides with the subgame perfect equilibrium, which is not to deactivate each other.

With its war-like flavor, our experiment is also set in the recent tradition of conflict experiments (Durham et al. (1998), Duffy and Kim (2000), Abbink et al. (2010)). Their games typically model a conflict where individuals or groups engage in fighting to win an exogenous resource. In our first strike game economic gains as a motive to attack are deliberately ruled out to isolate the effects of fear and the use of trust to overcome it.

2. The experimental design

We have four treatments to separate the influences of nastiness, fear, and trust. The first treatment, *asym* for asymmetric, serves as a control to see whether (some) human players indeed have malicious intentions. In this treatment only one player can deactivate, the other has no such option. Thus there is no need for self-protection, so deactivations are motivated by pure nastiness.

In the second treatment (*prerec*), both players can deactivate one another. One player, however, does not make his deactivation decisions himself, but they are made for him by the computer who plays exactly like a randomly chosen active player from the asym treatment. So the player who makes decisions in this treatment responds to prerecorded choices from the asym treatment. Differences between the asym and the prerec treatment can be attributed to first-order fear of nastiness. To protect themselves from possible nasty behavior (participants do not know the number of nasty players in the asym treatment) they may choose deactivation.

Table 1. The lottery choices

	Lottery A				Lottery B			
	High Outcome		Low Outcome		High Outcome		Low Outcome	
No.	Ct	%	ct	%	Ct	%	ct	%
1	190	15%	150	85%	320	15%	10	85%
2	190	30%	150	70%	320	30%	10	70%
3	190	45%	150	55%	320	45%	10	55%
4	190	60%	150	40%	320	60%	10	40%
5	190	75%	150	25%	320	75%	10	25%
6	190	90%	150	10%	320	90%	10	10%
7	170	15%	170	85%	220	15%	80	85%
8	170	30%	170	70%	220	30%	80	70%
9	170	45%	170	55%	220	45%	80	55%
10	170	60%	170	40%	220	60%	80	40%
11	170	75%	170	25%	220	75%	80	25%
12	170	90%	170	10%	220	90%	80	10%
13	200	15%	120	0.85	260	15%	50	85%
14	200	30%	120	0.70	260	30%	50	70%
15	200	45%	120	0.55	260	45%	50	55%
16	200	60%	120	0.40	260	60%	50	40%
17	200	75%	120	0.25	260	75%	50	25%
18	200	90%	120	0.10	260	90%	50	10%

In the third treatment (*sym*) the complete symmetric game is played. Here higher-order fear can take effect, but, because of the long-term interaction between two humans, trust and reciprocity can potentially serve as means to overcome the threat of mutual destruction. If a player has not been deactivated in previous rounds, he can reciprocate by not deactivating the opponent in the current round. If players anticipate such behavior, they may refrain from attacking as a signal of their good intentions already in the first round. This effect would reduce hostility compared to the prerec treatment.

Hence, as we move from the prerec to

the sym treatment, two counteracting forces can drive behavior in opposite directions, potentially canceling out each other. To separate the two effects we conduct a fourth treatment called *hidden*. In this condition the full game is played, but subjects are not told during the game whether or not the opponent has deactivated them. Thus all levels of fear are still present, but we have removed the possibility to build up trust and reciprocity from the game. With this manipulation we can isolate the effect of higher-order fear.

Under common knowledge of own-payoff-maximizing rationality players would not deactivate in any of the treatments. In asym, deactivation is not done because it is costly. In the subgame perfect equilibrium of the full game no player would ever deactivate the opponent. Since it is costly and bears no benefit, regardless of the other player's choice, no player deactivates in the final round. Foreseeing this, no player has any reason to deactivate in the penultimate round, and this logic unravels backwards such that there are no deactivations in the first round. There are imperfect Nash equilibria in which one player deactivates in the first round and the other's strategy choice is to deactivate in some later round $t>1$ (which is then of course not carried out because at that point this player is already deactivated). In these equilibria the first-round deactivator plays one of the best responses to the later deactivator. The later player could not improve by revising her strategy either. She could choose to not deactivate or deactivate at some other round including the first, in which cases the outcome for her would be the same. Hence such a strategy combination is an equilibrium, albeit not a subgame perfect one. It hinges on the fact that the decision node at which the later player deactivates is not reached, since the player is already deactivated in that round. Should the node be reached, the player would revise her strategy and, in foresight of the following rational play, not deactivate.[2]

Table 2. Summary of the four treatments

Asym	Prerec	Full	Hidden
Only one player can deactivate	One player can deactivate, the other's deactivation decisions are simulated according to records from the asym treatment	Both players can deactivate	Like full, but players are not told whether they are deactivated
No fear, pure nastiness	First-order fear	First- and higher-order fear, trust	First- and higher-order fear, no trust

The experiment was conducted at the Centre for Research in Experimental Economics and Political Decision Making (CREED) at the University of Amsterdam. The experiment was computerized, with software developed using the *RatImage* programming package (Abbink and Sadrieh (1995)). Subjects were recruited from a database of students who had previously registered at CREED. They were mostly undergraduate students from a wide range of disciplines. Each subject was allowed to participate in only one session.

At the outset of each session subjects were handed out the instructions and given time to read them. When everybody indicated that they had fully read and understood them, participants

[2] Strictly speaking there are no subgames in the hidden treatment, but it can easily be seen that the non-subgame perfect equilibria in the full game are equilibria in weakly dominated strategies.

performed an on-screen quiz. A session only began when all participants had answered all questions correctly. In the prerec treatment the instructions of the asym treatment were handed out and quizzed first, followed by the actual instructions for the prerec treatment and a quiz about them. The quiz was rather difficult, to rule out miscomprehension as a possible driver of our results as far as possible. After all subjects had completed the quiz(zes) play started immediately. The 18 lottery choices were presented in random order. This minimizes the possibility that contamination across tasks could induce artificial effects. If such effects exist the individually randomized sequence would make them show up as noise rather than systematic effects.[3] A session lasted for about 60 minutes (much of the time spent to read the instructions). After the play subjects were paid anonymously in cash. Deactivated subjects earned on average 35-40 euros, those who could avoid that fate left with 8-9 euros. Hence the stakes were higher than in most economic experiments, to make sure that deactivation was highly consequential. At the time of the experiment, the exchange rate to other major currencies was approximately US-$1.30, £0.90, ¥130, and RMB9 for one euro.

We conducted three sessions in each treatment. A session comprised between 12 and 20 subjects, the variation is due to no-shows. In the sym and hidden treatments every subject could make deactivation decisions, in the asym and prerec treatment only the active players. We experiment had 27 active subjects in each of the asym and prerec treatments, 56 in the sym treatment, and 50 in the hidden treatment. In the first round, each active subject is an independent observation, for later periods a pair of subjects matched to one another is.

3. Results

Figure 1 shows the deactivation rates in the four treatments. We do not observe a single deactivation in the asym treatment, where only pure nastiness can motivate hostility. Nevertheless, players in the prerec treatment seem extremely fearful of meeting a nasty opponent. Deactivation rates shoot up to 77.8%. Hence participants feel a strong need for self-protection, despite the considerable damage this does to an innocent fellow participant.[4] Though the symmetric setup in the sym treatment potentially adds another layer of fear, we observe deactivation rates sharply drop to 33.9%. This is evidence for the strong effect that trust and reciprocity can exert to prevent hostility. As we remove this possibility in the hidden treatment, deactivation rates are back to levels very similar to the prerec figures (68.0%). The pairwise differences between asym, prerec, and sym are highly significant ($p < 0.001$ one-sided, Fisher's

[3] Traditionally such effects are avoided by not playing and paying every chosen lottery, but randomly drawing a single task to be played for real (Cubitt et al. (1998)). We could not use this method since we needed subjects to accumulate earnings over the rounds.

[4] Note that subjects do not need to be excessively paranoid to choose deactivation for self-protection. Since the costs of an attack are relatively low compared to the losses from being deactivated, a purely selfish player would already deactivate if he estimates the fraction of nasty people to be around 7% or higher. Though we do not observe any deactivation, the binomial test does not reject the hypothesis that the fraction in the population is indeed 7% (one-sided p=0.141). The critical value at which we can reject the null at 5% significance is a fraction of 10.5%. Note, further, that deactivation does not protect against first-round attacks. So the player must believe that there are 7% or more players who would deactivate later than in the first round. The nastiest action is last-round deactivation, as it leaves no opportunity to recover even a small part of the damage.

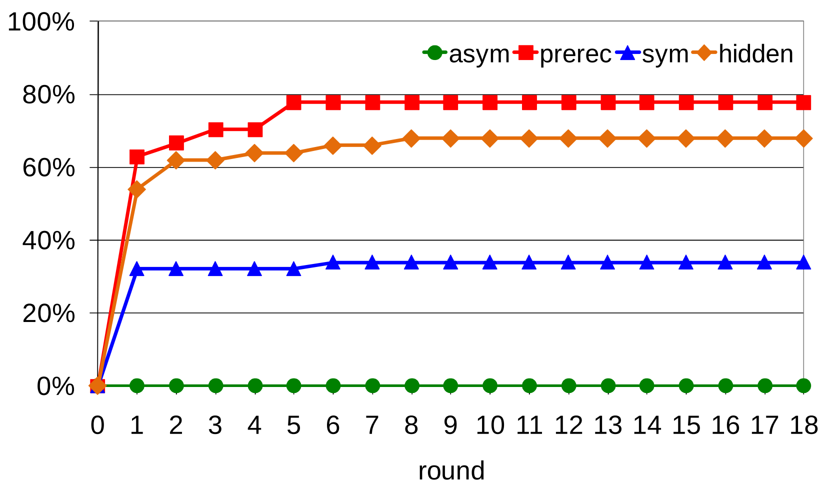

Figure 1. Deactivation rates in the four treatments over the 18 rounds of the experiment

exact test), sym and hidden are not statistically different (one-sided p = 0.145, all tests based on first-round decisions by individual participants[5]).

The figure also reveals that most of the deactivations take place in the very first round, especially in the sym treatment. This further indicates the role of reciprocity. Once the relationship has survived the initial stage without conflict, a history of non-aggression has been established and the urge to protect oneself against irrational attacks vanishes. Note that in later rounds, players can show actual reciprocity (a response to their counterpart's past behavior), while in the first round, where there is no history, such reciprocity must be anticipated. About one third of the subjects fail to do so, but the majority maintain non-aggressive relationships up to the end of the experiment. Note, further, that the only interaction between the two players is deactivation, so the message a player gets from not having been deactivated is unambiguous and no signals that may re-establish mistrust can be sent.

The question arises whether deactivation out of fear is driven by subjects' general risk-aversion. Our design with the lottery choices as money-earning task gives us the opportunity to check for such connections. Table 3 shows summary statistics on choice behavior. We divide the subjects into up to four categories depending on whether they chose to deactivate their opponent or not, and whether they were deactivated themselves. In addition, the asym and prerec treatments had two types of players, active and passive. Note that not in all treatments all four possible categories apply. In the asym treatment there were no deactivations, so we can only distinguish active and passive players. For this reason, no active player in the prerec treatment was deactivated. In the hidden treatment subjects did not know whether they were deactivated, so this distinction does not apply.

As a straightforward measure of risk-aversion we choose the number of times the subject has taken the safer choice, i.e. the one with the smaller difference between the two payoffs of a lottery. We observe very similar number of safe and risky choices between deactivators and non-deactivators, none of the differences is significant. Thus we do not find evidence that risk-aversion determines the propensity to deactivate. Subjects' attitudes towards strategic

[5] No result changes if we consider all rounds using pairs as independent observations. We focus on first-round choices to avoid minor comparability issues. In the sym treatment, a participant who is deactivated cannot deactivate anymore. Thus we do not know their strategy choices. In the other treatments this issue does not arise.

uncertainty (i.e. not knowing the behavior of others) does not seem to be related to their attitude towards risk (uncertainty about the outcome of a random draw with given probabilities).

Table 3. Summary of the decisions in the lottery choices

Sym	All subjects	Deactivators, non-deactivated	Deactivators, deactivated	Non-deactivators, non-deactivated	Non-deactivators, deactivated
#subjects	56	11	8	26	11
#choices	1008	198	144	468	198
#safe choices	584	107	88	298	100
%safe choices	57,94%	54,04%	61,11%	63,68%	50,51%

Asym	All subjects	Active	Passive
#subjects	52	26	26
#choices	936	468	468
#safe choices	601	301	300
%safe choices	64,21%	64,32%	64,10%

Hidden	All subjects	Deactivators	Non-deactivators
#subjects	50	34	16
#choices	900	612	288
#safe choices	538	360	178
%safe choices	59,78%	58,82%	61,81%

Prerec active	All subjects	Deactivators	Non-deactivators
#subjects	27	21	6
#choices	486	378	108
#safe choices	289	211	78
%safe choices	59,47	55,82%	72,22%

Prerec passive	All subjects	Deactivated	Non-deactivated
#subjects	27	21	6
#choices	486	378	108
#safe choices	246	180	66
%safe choices	50,62%	47,62%	61,11%

Only the sym treatment, in which sufficient data are available for all four categories (deactivators and non-deactivators, deactivated and non-deactivated), we can detect a statistically significant effect. Players who *were* deactivated show a significantly lower number of safe choices than those who were not (on average 9.3 versus 11.0, one-sided p=0.028, Fisher's two-sample permutation test). This is most likely due to the stakes in the lotteries, which were 90% lower after deactivation. Consistent with previous findings in the risky choice literature (Holt and Laury (2002)), subjects seem to be less risk-averse when the stakes are lower.

Table 4. Regression results

Choose deactivation	Coefficient	Std. Error	z	P>|z|
Age	-0.076	0.041	-1.860	0.063
Male	-0.166	0.249	-0.670	0.505
Economics	0.001	0.294	0.000	0.998
Religious	-0.151	0.260	-0.580	0.561
Foreign	-0.487	0.266	-1.830	0.067
Prerec	1.321	0.345	3.830	0.000
Hidden	1.014	0.268	3.780	0.000
Constant	1.711	1.026	1.670	0.095

N = 131
Pseudo r^2 = 0.171
Note: Dummies = 1 if as in label, = 0 otherwise

We handed out a post-experimental questionnaire to gather socio-demographic background data. We ran a probit regression with dummies for gender, religiousness (atheist/agnostic or others), subject of study (economics or others) and origin (Dutch or foreign). Dependent variable is the choice of deactivation. We restrict ourselves to the first round, where each individual constitutes an independent observation. The asym treatment is excluded from the analysis, as there were no deactivations there. Taking the full treatment as the baseline, we introduce treatment dummies for the prerec and hidden conditions. Table 4 shows the results of the regression. The treatment differences are significant, which replicates the findings from the non-parametric analysis. Individual characteristics, however, do not show up as significant determinants of behavior in the first-strike game.

4. Summary and conclusions

We introduce the experimental first-strike game to separate the effects of nastiness, fear and trust in a situation of bilateral threat. Though we do not find malicious behavior in our control treatment, fear of such is widespread and leads subjects to do considerable damage to each other. We show that this can entirely be attributed to first-order fear, i.e. the fear to meet a nasty person in the encounter. Higher orders of fear, the fear that the opponent may strike out of own fear and fear of fear, do not add anything to the propensity to strike. We even see slightly, but not significantly, higher deactivation rates in the treatment in which higher-order fear is removed.

These results show that fear is indeed a dangerous driver of hostile behavior. Note that we deliberately designed a game in which there is no economic benefit to be gained from an attack; it is only costly to do so. Other motives that could drive conflict in the real world do not apply either. Ideological dissent, ethnic divide, or personal antipathy are implausible motives for hostility, unlike in many real conflicts. Yet we observe that simple fear of nastiness is sufficient to make a large majority inflict considerable damage on another person.

Our data convey a more optimistic message as well. They show that humans can build up conflict-reducing trust in an ongoing relationship. The full game, in which all levels of fear as well as trust are effective, shows drastically reduced levels of hostility. This may be surprising under the adverse conditions we have created in our experiment. Once a player realizes that trust is not warranted, it is too late, since the player is already deactivated and thus robbed of the opportunity to respond in any way. Nevertheless, the possibility to signal trust at the start of the exchange reduces hostilities substantially, as the comparison between the full and the hidden treatment shows.

Our treatments allow the separation of spite, fear and trust in a very simple model of bilateral threat. We believe that with its simplicity the game can serve as a paradigm for studying situations with more complex threats. For example, future research can extend the game to look at scenarios where parties have a positive gain from attacking the opponent, for example to appropriate its resources. Further, we have looked at the simplest setup in which two individual players confront each other. Very often, however, conflicts emerge between groups or even nations. Though beyond the scope of this study, the inclusion of groups seems a promising area for future research.

References

Abbink, Klaus, Jordi Brandts, Benedikt Herrmann, and Henrik Orzen (2010): Inter-Group Conflict and Intra-Group Punishment in Experimental Contest Games. *American Economic Review*, 100, 420-447.

Abbink, Klaus, and Benedikt Herrmann (2010): The Moral Costs of Nastiness. *Economic Inquiry*, forthcoming.

Abbink, Klaus, and Abdolkarim Sadrieh (1995): RatImage – Research Assistance Toolbox for Computer-Aided Human Behaviour Experiments. SFB Discussion Paper B-325, University of Bonn.

Abbink, Klaus, and Abdolkarim Sadrieh (2009): The pleasure of being nasty. *Economics Letters*, 105, 306-308.

Berg, Joyce, John Dickhaut, and Kevin McCabe (1995): Trust, Reciprocity and Social History, *Games and Economic Behavior*, 10, 122-142.

Crawford, Vincent, and Nagore Iriberri (2007a): Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions? *Econometrica*, 75, 1721-1770.

Crawford, Vincent, and Nagore Iriberri (2007b): Fatal Attraction: Salience, Naivete, and Sophistication in Experimental "Hide-and-Seek" Games. *American Economic Review*, 97, 1731-1750.

Cubitt, Robin, Chris Starmer, and Robert Sugden (1998): On the Validity of the Random Lottery Incentive System. *Experimental Economics*, 1, 115-132.

Duffy, John, and Minseong Kim (2005): Anarchy in the Laboratory (and the Role of the State). *Journal of Economic Behavior and Organization*, 56, 297-326.

Durham, Yvonne, Jack Hirshleifer, and Vernon Smith (1998): Do the rich get richer and the poor poorer? Experimental tests of a model of power. *American Economic Review*, 88, 970-983.

Fehr Ernst, Georg Kirchsteiger, and Arno Riedl (1993): Does Fairness Prevent Market Clearing? An Experimental Investigation, *Quarterly Journal of Economics*, 108, 437-459.

Gilpin, Robert (1963): Two Prescriptions for a Sick World: Is the driving motive in international relations today reciprocal fear? *Science*, 139:3559, 1040-1043.

Großkopf, Brit, and Rosemarie Nagel (2008): The two-person beauty contest. *Games and Economic Behavior*, 62, 93-99.

Holt, Charles, and Susan Laury (2002): Risk Aversion and Incentive Effects. *American Economic Review*, 92, 1644-1655.

McKelvey, Richard, and Thomas Palfrey (1992): An experimental study of the centipede game. *Econometrica*, 60, 803-836.

Nagel, Rosemarie (1995): Unraveling in guessing games: An experimental study. *American Economic Review*, 85, 1313-1326.

Nagel, Rosemarie (1999): A Survey on Experimental Beauty Contest Games. In David Budescu and Ido Erev: Games and Human Behavior – Essays in Honor of Amnon Rapoport. Mahwah, NJ: Questia.

Schelling, Thomas (1960): The Strategy of Conflict. Part IV. Surprise Attack: A Study in Mutual Distrust. Harvard University Press.

Wohlstetter, Albert (1958): The Delicate Balance of Terror. RAND Corporation.

Zizzo, D (2004): Inequality and Procedural Fairness in a Money Burning and Stealing Experiment. In F.A. Cowell (ed.), Research on Economic Inequality, vol. 11. Elsevier.

Zizzo, D, and AJ Oswald, (2001): Are People Willing to Pay to Reduce Others' Incomes? *Annales d'Economie et de Statistique* 63-64, 39-62.