Soc Choice Welf (2015) 45:625–651
DOI 10.1007/s00355-014-0807-x

The UN in the lab

**Malcolm Kass · Enrique Fatas ·
Catherine Eckel · Daniel Arce**

Received: 29 January 2014 / Accepted: 3 February 2014 / Published online: 23 February 2014

Abstract We consider two alternatives to inaction for governments combating terrorism, which we term Defense and Prevention. Defense consists of investing in resources that reduce the impact of an attack, and generates a negative externality to other governments, making their countries a more attractive objective for terrorists. In contrast, Prevention, which consists of investing in resources that reduce the ability of the terrorist organization to mount an attack, creates a positive externality by reducing the overall threat of terrorism for all. This interaction is captured using a simple 3×3 "Nested Prisoner's Dilemma" game, with a single Nash equilibrium where both countries choose Defense. Due to the structure of this interaction, countries can benefit from coordination of policy choices, and international institutions (such as the UN) can be utilized to facilitate coordination by implementing agreements to share the burden of Prevention. We introduce an institution that implements a burden-sharing policy for Prevention, and investigate experimentally whether subjects coordinate on a cooperative strategy more frequently under different levels of cost sharing. In all treatments, burden sharing leaves the Prisoner's Dilemma structure and Nash equilibrium of the game unchanged. We compare three levels of burden sharing to a baseline in

Electronic supplementary material The online version of this article (doi:10.1007/s00355-014-0807-x) contains supplementary material, which is available to authorized users.

M. Kass · D. Arce
University of Texas at Dallas, 800 W. Campbell Road, Richardson, TX 75080, USA
e-mail: malcolmkass@utdallas.edu

E. Fatas (✉)
School of Economics, University of East Anglia, Norwich NR4 7TJ, UK
e-mail: e.fatas@uea.ac.uk

C. Eckel
Department of Economics, Texas A&M University, College Station, TX 77843, USA

a between-subjects design, and find that burden sharing generates a non-linear effect on the choice of the efficient Prevention strategy and overall performance. Only an institution supporting a high level of mandatory burden sharing generates a significant improvement in the use of the Prevention strategy.

1 Motivation

In combating terrorism, a government often faces a choice between two basic strategies: to pursue actions that directly undermine the capabilities of the terrorist organization, or actions that defend against terrorist attacks by hardening targets, thereby reducing damage in the event of an attack. Preventative actions, such as destroying a terrorist training ground, benefit all governments that are potential terrorist targets, as the terrorist's power to inflict damage is reduced, while the costs are borne entirely by the government that initiates the action. Defensive actions such as increased screening at points of entry or reinforcing the perimeter of an embassy may reduce the attractiveness of the defending government, while making other countries relatively more attractive as targets.

These tradeoffs are explored theoretically in a series of articles that identify the fundamental externalities associated with preventative and defensive counterterror strategies (Sander 2003; Sandler 2005; Arce and Sandler 2005a,b; Sandler and Siqueira 2006). In a "Nested Prisoner's Dilemma" framework, Arce and Sandler (2005a) model this choice as a pair of Prisoner's Dilemma games, combined into a 3×3 matrix consisting of three policy alternatives: Defense, Prevention, and Take No Action. They use this framework to illustrate the difficulty faced by countries that desire to coordinate on a policy of preventative action. While Prevention is the most efficient choice in the game, each country is tempted to free ride on the others' investment. With defensive policies, countries capture the benefits of their investment, but impose "costs" (in the sense of increased likelihood of attack) on other countries. The Nash equilibrium of the game is for both countries to invest in defensive actions, an outcome that is doubly inefficient in that it is both Pareto dominated and minimizes the utilitarian sum of the countries involved. Thus countries face a difficult problem in coordinating on effective counterterrorist policies. The incentive structure of the interaction implies that countries acting in their individual self-interest will fail to reach the optimal policy choice, and instead each invests in sub-optimal defensive policies.

International organizations have arisen to combat myriad collective action problems. With respect to terrorism, peacekeeping is an activity that is often rationalized as a means of stabilizing areas that may otherwise turn out to be terrorist breeding grounds. The two institutions that carry out the majority of peacekeeping missions—the UN and NATO—are financed by voluntary contributions, making the question of burden sharing for preventative counterterror policies a central one. INTERPOL is another international organization that is voluntarily financed, and its MIND/FIND database for stolen and lost travel documents, for example, has been shown to be cost effective in combating terrorism (Sandler et al. 2011).

A promising avenue for testing the effect of alternative policies to enhance counterterror policy coordination is to simulate countries' choices in the lab. Testing such policies in the field is costly or impossible. The lab provides a kind of wind-tunnel for exploring policy options at a relatively low cost. The impact of the incentive structure of a given policy can be evaluated, and alternatives compared. This is the approach we use in the current paper.

The current paper utilizes the experimental environment developed in Arce et al. (2011a,b), which experimentally tests the theoretical predictions of the Nested PD game.[1] In their experiment, subjects were matched in stable pairs for 20 rounds. In the 21st round subjects were re-matched and the stakes increased by a factor of 10. As is common in experimental studies of repeated standard PD games (Sally 1995), non-equilibrium behavior was clearly evident in early rounds of the 20-round section of the experiment. Initially 55 % of subjects played Prevention, 30 % chose Take No Action (maintained the status quo) and 15 % acted non-cooperatively and selected the Defense strategy, consistent with Nash equilibrium play. By the 20th round, 60 % of the subjects behaved non-cooperatively. In the 21st round, 58 % behaved non-cooperatively and 26 % cooperatively (the remainder chose Take No Action). Hence, while the experimental results are not as dire as theory predicts, the collective action problem remains one that subjects do not easily overcome on their own; hence we consider institution building in the form of cost-sharing rules for preventative counterterror policy.

We expand this framework to test the ability of a stylized institution, cost sharing, to alleviate the collective action problem. In the experiment we implement three different levels of cost sharing, while holding the Nash equilibrium of the game constant. Our experimental conditions can be interpreted as different policies enforcing the sharing of the cost of investing in Prevention: this policy diminishes the risk that others will free ride on such a decision. Policies vary in the extent to which players share the cost of choosing Prevention. Cost sharing serves to partially internalize the positive externality associated with Prevention, while keeping the negative externality of choosing Defense constant. We systematically manipulate the intensity of cost sharing in a linear way. The cost of Prevention is not shared at all in the Baseline, one unit is shared in the Low Burden Sharing condition, two units in the Moderate, and three in the High Burden Sharing treatment.

We find that cost sharing can be a powerful policy choice, boosting the utilization of the optimal policy, but only when the level of cost sharing is high. Low levels of cost sharing are more symbolic than effective, and have little impact on strategy choices, with more than half of the pairs playing the defensive strategy. But a high level of cost sharing (as in our High Burden Sharing treatment) induces almost around 85 % of the pairs to cooperate,[2] choosing mutual Prevention.

[1] This paper also compares the deterministic-payment Nested PD presented in the theory (see Table 1) with a probabilistic treatment where the expected values of the individual elements equal the deterministic values, but successful attacks lead to payoff losses. The authors find that the preventative choice is selected with less frequency in the probabilistic treatment.

[2] Aggregating rounds 1–21.

The experimental analysis of cost sharing in Prisoner Dilemma and related games is not new, and can be traced back to Smith (1980) analysis of contract mechanisms. It is far from the goal of this paper to make a complete survey of this particular branch of the extensive literature on cost-sharing mechanisms.[3] However, the natural reference for our paper is Andreoni and Varian (1999), who analyze a two-stage variant of a PD game in which players are allowed to make side payments. In this setting, the unique sub-game perfect equilibrium is efficient and supported by their experimental test. Charness et al. (2007) extend their analysis to show that side payments can be efficient in PD games when players have the chance of simultaneously choose a binding transfer in a preliminary stage. The compensation mechanism in Charness et al. (2007) transforms the PD into a coordination game in which mutual defection can still be an equilibrium in some treatments. However, as in Andreoni and Varian (1999), most transfers observed in the experiment were consistent with a SPE involving mutual cooperation, suggesting that cooperation gains are explained by equilibrium play.[4]

Our experimental test of different cost sharing mechanism differs from previous papers in three main respects. First, the degree of cost sharing is not selected by subjects, but rather exogenously imposed. We consider this to be the right approach when trying to identify the behavioral implications of different institutions. Second, all treatments share the same inefficient Nash equilibrium prediction: mutual Defense. Rational, selfish players should not change their behavior relative to our control treatment with no cost sharing. We believe this feature provides a strong test for the emergence of cooperation under different cost sharing rules. Our last contribution has to do with the double-externality game and the asymmetric nature of our cost sharing mechanism. While Prevention generates a positive externality, Defense generates a negative one. Cost sharing is implemented only for the Prevention strategy. This asymmetry is of particular interest in the domain of counterterrorism policies and has never been explored in the laboratory.

This paper contributes to the literature on conflict resolution. Our results strongly suggest that the effectiveness of different policies is heavily mediated by behavioral biases, hard to anticipate in a theoretical model, but easy to estimate in the lab. Our between-subjects experimental design tests different cost sharing policies in a simple decision environment. Subjects do react to the introduction of a cost sharing mechanism, even when all games share the same theoretical prediction. However, they respond only when the level of cost sharing is high. The significant increase in the utilization of the optimal strategy is stable across the 20 repetitions of the original game, and observed again in one last round in which the stakes are much higher and participants are randomly re-matched. In this sense, this paper extends cost sharing to a new environment, and shows the necessity of a better understanding of how different institutions solve conflict by inducing more out-of-equilibrium cooperative play.

3 Fong and Surti (2009) and Matsushima (2012) survey the recent theoretical literature.

4 Compensation mechanisms have also been analyzed in the parallel domain of public goods games. Not surprisingly, results are similar to the ones described for PD games. Chen (2003) and Razzolini et al. (2007) find that subjects easily converge to the efficient Nash equilibrium under different compensation regimes. Bracht et al. (2008) get a similar result in a more sophisticated environment (the public good uses a quadratic production function with interior equilibrium).

Finally, it illustrates the power of the lab in exploring policy alternatives. Our work contributes to the growing experimental literature on terrorism and counterterrorism policy (e.g., Cadigan and Schmitt 2010; Colombier et al. 2011; Arce et al. 2011a provide a survey.) The lab is increasingly being utilized to examine questions in this area that are difficult to address using observational or survey data, or theory alone. We return to this discussion in the conclusion.

The rest of the paper is organized as follows: in the next section we present the Nested PD game used in our experiment. Section 3 presents the experimental design, Sect. 4 presents our hypothesis, and Sect. 5 procedures. Section 6 analyzes the results and the last section concludes.

2 The nested PD game

The primary vehicle for our experiments is the 3×3 global security game developed in Arce and Sandler (2005a). The game combines elements of a voluntary-contribution public goods game, which corresponds to the externality generated by preventative counterterror actions, with those of a commons dilemma in the sense of Harden (1968), corresponding to the externality associated with defensive counterterror actions. In what follows, payoffs that are capitalized represent public consequences of strategic actions and lower-case variables measure private (individual) consequences. Specifically, consider a collective action environment with a public good having benefit $\boldsymbol{B}$ (i.e. to both countries) at individual private cost of provision $\boldsymbol{c}$. Similarly, an action that creates private benefit $\boldsymbol{b}$ does so by simultaneously creating a public 'bad' denoted by cost $\boldsymbol{C}$.

Arce and Sandler (2005a) note that preventative/counterterrorism actions such as compromising a cell or destroying a training ground have classic public goods properties. For example, the action Prevention keeps a terrorist from attacking any of its targets, regardless of nationality, thereby creating a public benefit $\boldsymbol{B}$ for all targeted nations at a private cost to the acting nation of $\boldsymbol{c}$. Such a 2-player situation is denoted in the northwest 2×2 collection of cells in the game in Table 1, where a value of 0 denotes the status quo payoff. Under the standard assumptions for the voluntary provision of public good, $\mathbf{2}\boldsymbol{B} > \boldsymbol{c} > \boldsymbol{B} > \mathbf{0}$, no action is taken—yielding Olson (1965) classic free rider problem.

By contrast, when a targeted nation engages in defensive counterterror activities such as body scans at airports or hardening the defenses around embassies, public costs, $\boldsymbol{C}$, are created in terms of lost civil liberties and/or the targeting of another nation or the acting nation's citizens in a foreign country. For example, the vast majority of US

Table 1 Nested prisoner's dilemma

	Prevention	Take no action	Defense
Prevention	$2\boldsymbol{B}-\boldsymbol{c}, 2\boldsymbol{B}-\boldsymbol{c}$	$\boldsymbol{B}-\boldsymbol{c}, B$	$\boldsymbol{B}-\boldsymbol{c}-\boldsymbol{C}, \boldsymbol{b}+\boldsymbol{B}-\boldsymbol{C}$
Take no action	$\boldsymbol{B}, \boldsymbol{B}-\boldsymbol{c}$	0,0	$-\boldsymbol{C}, \boldsymbol{b}-\boldsymbol{C}$
Defense	$\boldsymbol{b}+\boldsymbol{B}-\boldsymbol{C}, \boldsymbol{B}-\boldsymbol{c}-\boldsymbol{C}$	$\boldsymbol{b}-\boldsymbol{C}, -\boldsymbol{C}$	$\boldsymbol{b}-\mathbf{2}\boldsymbol{C}, \boldsymbol{b}-\mathbf{2}\boldsymbol{C}$

victims of terror are attacked outside of the US and the emphasis on defensive policy in the Western World subsequent to the attacks of 9/11, 3/11 and 7/7 in the US, Madrid and London, respectively, have caused fundamentalist terror activities to relocate to Southeast Asia. Defensive actions create private benefit, **b**, and public costs, **C**, where it is typically assumed that $\mathbf{2C > b > C > 0}$. The resulting 2-player game is given in the 2×2 collection of cells in the southeast of Table 1. Under standard assumptions this yields Harden (1968) Tragedy of the Commons, because Defense is dominant, but the (Defense, Defense) outcome is Pareto-dominated by mutual inaction.

This game combines a 2×2 Prisoner's Dilemma in the northwest four cells of Table 1 with another 2×2 Prisoner's Dilemma in the southeast four cells. For this reason the resulting 3×3 global security game is called a Nested PD game. Under the assumptions given above, Defense is the dominant strategy in the global security game, leading to an outcome that not only is Pareto-dominated by mutual Prevention and mutual inaction (Take No Action), but also represents the outcome with the lowest total payoffs. This leads to the novel observation that inefficiency increases when the public goods and commons versions of the Prisoner's Dilemma are combined. Under this scenario, we infer that mutual cooperation would be highly improbable.

Nonetheless, there are policies that can be used to alleviate the potential problems associated with choosing the socially-optimal Prevention strategy. For instance, "burden-sharing" institutions can be created among member states to share the costs of Prevention activities. For example, Sandler et al. (2011) have shown that INTERPOL's MIND/FIND database for stolen and lost travel documents is a particularly effective form of defensive counterterror policy, and member nations voluntarily contribute to INTERPOL's budget.

NATO and UN are examples of organizations that both coordinate activities that are akin to preventative counterterror measures (e.g. peacekeeping missions) and rely on voluntary contributions. Whether such institution-building is a worthwhile activity and what form such an institution should take can be best answered within an experimental framework.

In particular, within the structure of the Nested PD game, whenever a country chooses the Prevention action, part of the private cost **c** will be transferred to the other nation, as is consistent with pre-determined burden-sharing arrangements in many international organizations such as the UN, NATO, etc. Let $\boldsymbol{f}$ be the fraction of cost that is shared in the burden sharing agreement. In order to keep the current structure of the game, the following constraint on the burden share fraction $\boldsymbol{f}$ must be met:

$$f < 2/3$$

Above that critical value the fundamental structure of the game is altered and Defense is no longer a strictly dominant strategy. Note that, while Defense remains a dominant strategy for burden-sharing less than 2/3, lower values can alter the equilibrium of the Prevention subgame (in the upper left four cells of the matrix). In particular for $\boldsymbol{f}$ above 1/3, mutual Prevention would be an equilibrium of the stand-alone Prevention game. This is discussed in more detail in the experimental design below.

Table 2 demonstrates the changes that occur to Table 1 when a burden sharing mechanism is implemented for the cost of the Prevention strategy. As an example, if

Table 2 Changes in payoff due to burden sharing

	Prevention	Take no action	Defense
Prevention	0,0	$+f(c)$, $-f(c)$	$+f(c)$, $-f(c)$
Take no action	$-f(c)$, $+f(c)$	0,0	0,0
Defense	$-f(c)$, $+f(c)$	0,0	0,0

Table 3 Baseline nested prisoner's dilemma

	Prevention	Take no action	Defense
Prevention	9,9	5,11	1,13
Take no action	11,5	7,7	3,9
Defense	13,1	9,3	5,5

the row player selects Take No Action while the column player chooses Prevention, the payoff to the row player is $\boldsymbol{B}-\boldsymbol{fc}$ and the payoff to the column player is $\boldsymbol{B}-\boldsymbol{c}+\boldsymbol{fc}$. Within the range for f given above burden sharing has no theoretical implications for the global security game, but as will be seen below burden-sharing has significant normative implications for coordinating counterterror policy.

3 Experimental design

This experiment involves a baseline and three treatments testing different levels of cost sharing for Prevention. The baseline is a direct test of the theory presented in Arce and Sandler (2005a) and has been tested previously in Arce et al. (2011a,b). It is important to distinguish the current contribution from that paper. Both use the same baseline protocol, but Arce et al. compare the baseline game in Table 3 with a probabilistic version with the same expected payoffs. The current paper utilizes the same baseline, but compares the baseline to three different levels of burden-sharing. It does not include a probabilistic version of the game.

The Baseline (in both papers) consists of 20 rounds with payoffs shown in Table 3. This game is played for 20 rounds with stable partners (randomly assigned), with a 21st round where subjects are randomly rematched and payoffs increased by a factor of 10. The 21st round is included to provide a direct test of the theory. By the 21st round subjects are familiar with the game, and have some experience with one partner. The 21st round increases incentives sharply, and, in theory, random rematching should make the game play independent of prior rounds.

To prevent a potential focal point bias around the payoff of zero, we set 7 tokens as the status quo amount per round. Therefore, the Baseline payoffs are as follows, where $\mathbf{B}=4$, $\boldsymbol{b}=6$, $\boldsymbol{C}=4$, $\boldsymbol{c}=6$, and $\mathbf{f}=0$ (indicates no burden sharing):

The goal of this experiment is to determine if enforced sharing of the cost of choosing Prevention will lead to more cooperative play between two players, even if Defense is still the overall dominant strategy. Therefore, we designed 3 treatments with different levels of burden sharing where we adjust the proportion of Prevention costs that are shared. We call these treatments High, Moderate, and Low Burden Sharing.

In the Low Burden Sharing treatment, 1/6 of the cost of Prevention is transferred to the other player. Therefore, in Table 2, $\boldsymbol{f}=\frac{1}{6}$, and since the private cost of selecting

Table 4 Nested PD games with burden sharing

	Prevention	Take no action	Defense
A, Low burden sharing			
Prevention	9,9	6,10	2,12
Take no action	10,6	7,7	3,9
Defense	12,2	9,3	5,5
B, Moderate burden sharing			
Prevention	9,9	7,9	3,11
Take no action	9,7	7,7	3,9
Defense	11,3	9,3	5,5
C, High burden sharing			
Prevention	9,9	8,8	4,10
Take no action	8,8	7,7	3,9
Defense	10,4	9,3	5,5

Prevention is 6 units in our experiment, the total amount of the cost of Prevention that is transferred to the other player is equal to 1. Therefore, since the cost of selecting a strategy of Prevention is six, one unit of the cost of Prevention is transferred to the other player. Table 4A below shows the normal form payoff matrix used in the Low Burden Sharing treatment.

In the Moderate Burden Sharing treatment, $f = \frac{1}{3}$, and in the High Burden Sharing treatment, $f = \frac{1}{2}$, and since the private cost of selecting Prevention is constant at 6 units, the two units of cost are transferred to the other player in the Moderate Burden Sharing treatment, and three units in the High Burden Sharing treatment. Tables 4B, C show the normal form payoff matrices used in the Moderate and the High Burden Sharing treatments.

The burden sharing proportions where selected based on two criteria. First, we selected proportions that corresponded to integer shares of costs of Prevention to keep the decision environment as simple as possible for subjects. Second and more importantly, because the burden-sharing fraction is always below 2/3, the Defense is the strictly dominant strategy for both players.

While the burden sharing treatments do not alter the Nash equilibrium of the Nested PD, they do have an effect on the equilibrium of the stand-alone Prevention game (top left four cells). In the Baseline and Low Burden Sharing games, mutual Take No Action is the single Nash equilibrium of the stand-alone game. However, with Moderate Burden Sharing, the Prevention game has two weak equilbria: both common Prevention and common Take No Action are equilibria, with Prevention Pareto-dominant. With High Burden Sharing, mutual Prevention becomes the only Nash equilibrium of the Prevention game. Note that in all cases Defense remains the dominant strategy for the full Nested PD game.

Note that following the classic "unraveling" argument given in Luce and Raiffa (1957), the sub-game perfect equilibrium for any of the three versions of the game tested is the same as for the baseline stage game: the players will always select Defense, which strictly dominates Prevention and Take No Action. This is because we keep the burden-sharing fraction, f, less than the critical value of 2/3, above which Defense is no longer a strictly dominant strategy. In addition, while Take No Action strongly dominates Prevention for the baseline and the Low Burden Sharing treatment, it only weakly dominates Prevention in the Moderate Burden Sharing and is dominated by Prevention in the High Burden Sharing treatment. However, as long as all players are payoff-maximizers, they have no reason to deviate from Defense in any treatment.[5]

4 Hypotheses

The idea that cooperation in repeated PD is sensitive to changes in incentives to the size of the incentives to deviate from mutual defection and mutual cooperation goes back to Rapoport and Chammah (1965) and Rapoport (1967). The intuition is simple: cooperation will be decreasing in the size of the incentives to unilaterally deviate from mutual cooperation (and the penalty paid when deviating from mutual defection). In our paper, we systematically manipulate the cardinality of these relationships by reducing the incentives to deviate from mutual cooperation, and at the same time decreasing the penalty for deviating from mutual defection. Moreover, our manipulation has two interesting features: it is symmetric and linear.

We can illustrate this idea with a simple intuition coming from Tables 3 and 4. In our baseline, the incentive to deviate from mutual Prevention is 4. Payoffs increase from 9 to a maximum of 13 when a player deviates from mutual Prevention to Defense. Symmetrically, the penalty for a unilateral deviation from mutual Defense is also 4. Payoffs decrease from 5 to a minimum of 1 when a player unilaterally deviates from Nash and chooses Prevention. By introducing cost sharing, we manipulate these incentives (and penalties). In the Low Burden Sharing condition, the incentives to deviate are 3, in the Moderate condition 2, and in the High sharing treatment is only 1. Note that penalties for deviations from Nash also decrease from 4 in the baseline to 1 in the High Burden Sharing condition.

In an interesting experiment, Ahn et al. (2001) check the effect of cardinal manipulations and learning spillovers in a sequence of one-shot PD games. Even though our experiment is very different from theirs, we follow closely their behavioral conjectures to formulate a simple behavioral conjecture for this game:[6]

Conjecture 1: The rate of cooperation will be decreasing in the incentives (penalties) to deviate from mutual Prevention (Defense).

This simple conjecture allows us to order our four experimental treatments by the expected frequencies of Prevention (Defense):

---

5 There are additional theoretical considerations for the repeated version of this game, illuminated by Kreps and Wilson (1982) and Axelrod (1984). However, we limit our theoretical analysis to the stage game. Of course there are behavioral considerations for the repeated version of the game as well, following Shafir and Tversky (1992). Fairness concerns may potentially come into play as in Charness and Rabin (2002), as burden sharing reduces differences in payoffs associated with unilateral Prevention.

6 Ahn et al. (2001) are mainly interested in whether playing different coordination games in the Phase 1 of the experiment had any effect in two sequences of one-shot coordination games and PDs. They manipulate the size of the incentives to deviate in the one shot PDs in a binary way, much coarser than ours (high or low), and also alter the symmetry of incentives and penalties (High–High, High–Low, Low–High, Low–Low). Their behavioral conjectures are presented in pages 140 and 141.

$$Prev^{High-BS} > Prev^{Mod-BS} Prev^{Low-BS} > Prev^{BSL}$$
$$(Def^{High-BS} < Def^{Mod-BS} < Def^{Low-BS} < Def^{BSL})$$

Note that an alternative hypothesis based on the assumption of fully rational and selfish players applying backward induction (and this being common knowledge) would predict no differences across treatments.

5 Procedure

All sessions were conducted at the Center for Behavioral and Economic Experimental Science at the University of Texas at Dallas. A total of 142 undergraduates from various academic disciplines (including economics) and with no previous history of playing the Nested PD game were recruited via the ORSEE system (Greiner 2004). There were 12 sessions, three sessions per treatment, and twelve subjects participated in each session (except for one baseline treatment session, where only ten subjects participated). No subject participated in more than one session. Once subjects were present in the lab common area, each was randomly assigned to a computer terminal where partitions ensured that subjects could not see each other's decisions. Before beginning the experimental instructions, the subjects were told not to communicate with each other, but to feel free to ask questions. Then subjects read the self-paced, computerized experimental instructions, which explained the baseline game and externalities associated with different strategies, and the level of burden sharing tied with the Prevention game. They were shown the complete matrix itself and were informed that earnings would be cumulative. The experiment was programmed using z-tree software (Fischbacher 2007); instructions are available in online Appendix B.

Given the complexity of the asymmetric externalities in the Nested PD, we opted for instructions that are framed and have naturalistic, somewhat loaded language. Therefore, the Nested PD game was presented to subjects as an interaction between countries willing to reduce the possibility of a terrorist attack. We believe that by using a naturalistic context we were able to explain in a simpler and more convincing way the nature of the game, as well as the different consequences of Prevention and Defense. Given that instructions were identically framed in each treatment, they cannot explain any treatment effect.[7] Moreover, we control for the emergence of "hot" behavioral reactions of some subjects by running an in-depth questionnaire at the end of the

[7] The instructions were prepared using a frame that was designed to make the game easier for subjects to understand. The loaded language also may have affected perceptions about the appropriateness of alternative strategies: since the theory we are testing is explicitly motivated by the need to coordinate anti-terrorism policy, greater external validity is likely when the game is also framed in this way. Previous research on the prisoner's dilemma game finds increased levels of cooperation when the game is framed as a cooperative situation: for example, Liberman et al. (2004) found that naming the PD game a "Community" versus "Wall Street" game generated differences in levels of cooperation. However, in contrast, Abbink and Hennig-Schmidt (2006) find no effect of loaded terms such as "bribery" versus abstract wording in a bribery experiment. There is also some debate about the external validity of neutral framing, that norm activation and other psychology factors should be accounted for (see Eckel and Grossman 1996; Loomes 1999).

experiment. As it will be clear in the next section, individual preferences played no role in explaining the different treatment effects.[8]

Once the subjects finished reading the instructions, they were given six examples and a quiz to ensure that they understood the game. Participants also were notified that for the first 20 rounds of the experiment they would be playing with a single partner, and that on the 21st round, they would be randomly re-matched with another partner. Once the experiment started, players were shown the results and earnings of their decisions each round. Subjects were not shown a history of previous play, but were given a record sheet to record game play and earnings in each round. After the first 20 rounds, subjects then read a short description of the 21st round, where they were reminded that they would be randomly re-matched, and informed that payoffs for that round were increased to 10 times the amounts in the previous rounds. Once the experiment was finished, the subjects then completed a short survey, at the conclusion of which they were paid in private the $5 show up fee and experimental earnings.

On average the experiment took approximately 50 min, and the survey and payouts an additional 35 min. Average experiment earnings for the baseline were $20.93 ($14.75 without the 21st round), for the Low Burden Sharing treatment was $20.44 ($14.78 without the 21st round), for the Moderate Burden Sharing was $22.46 ($15.90 without the 21st round), and for the High Burden Sharing was $25.64 ($17.94 without the 21st round).[9]

6 Results

The results section is organized as follows. First we present the aggregate data, separating out rounds 1 and 21, because these are the only rounds where behavior is arguably independent of prior decisions. We then turn to a more in-depth discussion of individual behavior in rounds 1 and 21. This is followed by a panel data analysis of the repeated-game portion of the experiment: rounds 1–21. Finally, we discuss strategy coordination, asking how frequently pairs of subjects successfully coordinate on a common strategy choice.

6.1 Aggregate decisions

Table 5 contains information about the sample and the average subject play in the baseline and the three burden sharing treatments, for rounds 1–20 and round 21 separately.[10] From the table, we see that, as more of the cost of choosing Prevention is

[8] We present some additional analysis in a longer online version with an additional Appendix B (Kass et al. 2014). Table 1' in this Appendix B presents the information collected in the questionnaire, including political orientation, views on the US role in the world, or government spending. We also collected information about subjects' demographics in thirteen different dimensions.

[9] Table 1' in Kass et al. (2014), see Appendix B (Electronic Supplementary Material), contains comparison tests of the subject characteristics and beliefs between the treatments. Only the proportions of married subjects and full time students are different between the treatments below a 0.1 significance threshold, and neither drives results in our regression models.

[10] Baseline data are from Arce et al. (2011a,b).

Table 5 Descriptive statistics, rounds 1, 1–20, and 21

	Baseline	Low burden sharing	Moderate burden sharing	High burden sharing
Subjects	34	36	36	36
Pairs	17	18	18	18
Observations	714	756	756	756
Round 1				
Prevention	55.88 %	63.89 %	69.44 %	88.89 %
Take no action	29.41 %	11.11 %	8.33 %	2.78 %
Defense	14.71 %	25.00 %	22.22 %	8.33 %
Average profit	7.82 (3.31)	7.78 (3.35)	7.94 (2.64)	8.61 (1.46)
Rounds 1–20				
Prevention	55.59 %	54.17 %	72.08 %	89.31 %
Take no action	7.65 %	11.11 %	3.33 %	1.25 %
Defense	36.76 %	34.72 %	24.58 %	9.44 %
Average total profit	147.53 (31.98)	147.78 (30.68)	159.00 (26.31)	171.94 (16.53)
Average round profit	7.38 (2.78)	7.39 (2.52)	7.95 (2.14)	8.60 (1.33)
Round 21				
Prevention	26.47 %	11.11 %	38.89 %	69.44 %
Take no action	5.88 %	11.11 %	0 %	0 %
Defense	67.65 %	77.78 %	61.11 %	30.56 %
Average profit	61.76 (46.48)	56.67 (29.08)	65.56 (32.82)	81.11 (22.65)
Average total profit	209.29 (52.17)	204.44 (39.23)	224.56 (47.76)	256.39 (23.66)

Standard deviations are in parenthesis

shared with the counterpart, players choose Prevention more frequently. Except for round 21 where Defense selections are more common and constitute the most common strategy for all but the High Burden Sharing treatment. In the baseline, 55.59 % of subjects choose Prevention. This fraction is slightly lower for the Low Burden Sharing treatment, but increases steadily thereafter, reaching 89.31 % with High Burden Sharing. The last line of the table is average total profit for each subject (up to round 20), and it is increasing with the level of burden sharing as well. In the aggregate, there is evidence of social gains associated with burden sharing agreements, even though burden sharing itself does not alter the aggregate benefit-cost structure (payoffs are always identical across treatments, for each common pair of decisions). For round 21, in the baseline, 26.5 % of subjects choose Prevention. This fraction is considerably lower for the Low Burden Sharing treatment (11.1 %), but increases steadily thereafter, reaching 69.4 % with High Burden Sharing. Round 21 profits echo the results in rounds 1–20. In combining rounds 1–20 and round 21, only the Moderate and High Burden Sharing treatments generate social gains.

Figure 1 illustrates the fraction of Prevention decisions by round, and suggests that play in round 1 is a relatively good predictor of decisions in rounds 1–19, as average Prevention selections per round never deviate much, in any treatment. From

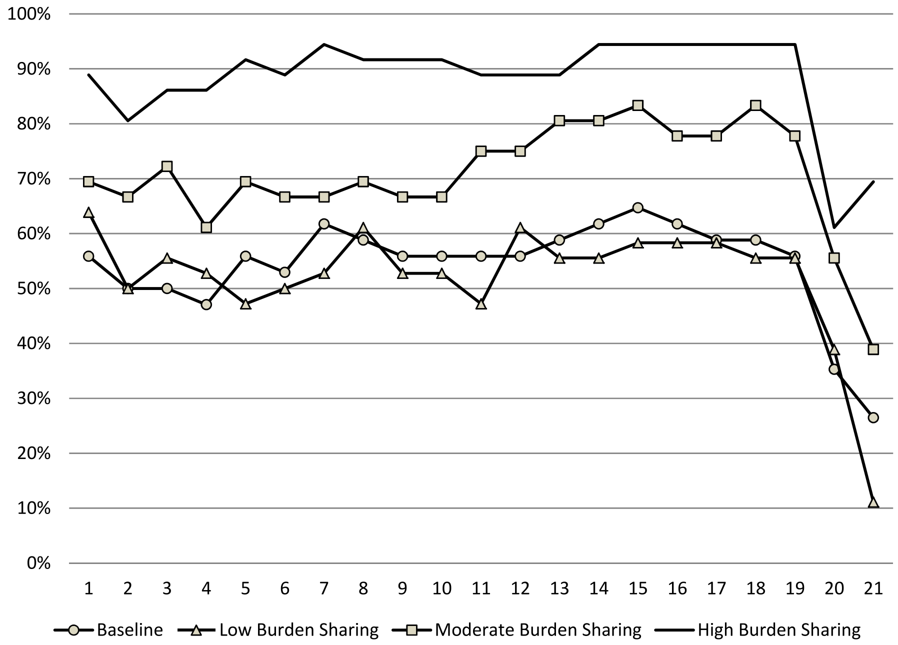

Fig. 1 Percent of subjects choosing Prevention in each round

Fig. 1, we see that different burden sharing institutions generate very similar Prevention dynamics, and the proportion of Prevention decisions remain relatively constant, with a strong endgame effect in round 20 and a further decline in the re-matched round 21. The exception is the High Burden Sharing treatment, in which Prevention selections increase slightly.[11]

There is no apparent reaction to Low Burden Sharing, and only a modest change when burden sharing reaches a Moderate level. This modest effect occurs as subjects abandon both the Take No Action and the Defense options. However, the dramatic shift observed in the High Burden Sharing treatment is not explained by a drop in the Take No Action selections (a difference of 2.08 % points), but rather by a reduction in Defense. We analyze these different patterns more formally below.

6.2 Round 1 and round 21

We next turn our attention to an investigation of rounds 1 and 21 play for two reasons. First, we consider this to be a natural way to analyze the data, as in round 1 there has been no interaction and data are free of any influence of prior play, and in round 21 subjects are experienced, but are matched with new partners for higher stakes, arguably making them independent from prior rounds.

[11] Figure 3 in Kass et al. (2014), see Appendix B (Electronic Supplementary Material), shows the distribution of Prevention decisions and earnings by subject. The histograms suggest that subjects are not reacting to changes in the burden sharing parameter in a linear, or even strictly monotonic manner.

Table 6 Nonparametric tests, round 1

	Prevention	Take no action	Defense
A. Actions selected in round 1			
No burden sharing	19	10	5
Low burden sharing	23	4	9
Moderate burden sharing	25	3	8
High burden sharing	32	1	3
	No burden sharing	Low burden sharing	Moderate burden sharing
B. Fisher exact test (prevention, take no action, defense), p values shown			
Low burden sharing	0.133	–	–
Moderate burden sharing	0.073	0.867	–
High burden sharing	0.004	0.043	0.127
C. Fisher exact test (prevention vs. non-prevention), p values shown			
Low burden sharing	0.626	–	–
Moderate burden sharing	0.323	0.803	–
High burden sharing	0.003	0.025	0.079

Table 6A presents the aggregate actions taken by subjects in round 1 and Table 6B, C contain p values from the Fisher Exact test. In Table 6B, we see evidence that the Moderate Burden Sharing and the High Burden Sharing treatments are significantly different from the Baseline. Table 6C presents a slightly different approach, and tests for differences in Prevention as opposed to the aggregate choice of other strategies (Take No Action and Defense decisions are aggregated into one "Non-Prevention" variable.) This more clearly illustrates that in round 1, subject decision-making only leads to more socially beneficial outcomes when subjects are interacting in the High Burden Sharing environment. We conclude that even in first round play, High Burden Sharing has a substantial impact on the play of Prevention as a counterterrorism strategy.

Table 7 presents the results of a binary Logit model using round 1 data where the dependent variable is Prevention and we control for treatment, demographics and personal preferences.[12] Table 7 strongly suggests that only the High Burden Sharing

[12] Demographics include binary variables gender (0 = male, 1 = female, white (Yes = 1), Married (Yes = 1), children (Yes = 1), live with parents (Yes = 1), grad student (Yes = 1), economics student (Yes = 1), full time student (Yes = 1), Christian (Yes = 1), US Citizen (Yes = 1), republican (Yes = 1). Ordinal demographic variables are father's education level and religiosity. Interval demographic variables are age, family income, and number of bathrooms in the subject's household. Preferences include binary variables such as if the person views him/herself as politically conservative, if the subject believes the US should have a strong leadership role, believes that she/he is politically astute, believes that minorities face discrimination, believes that the amount of immigrants in the US should decrease, believes that the US should be

Table 7 Strategy choice, round 1

Dependent variable: prevention	Model 1	Model 2[c]	
Dependent variable is prevention. Logit regression models, marginal effects shown			
Low burden sharing	0.066	0.151	
	(0.095)	(0.096)	
Moderate burden sharing	0.115	0.017*	
	(0.096)	(0.097)	
High burden sharing	0.362***	0.403***	
	(0.113)	(0.110)	
Demographics[a]	N	Y	
Preferences[b]	N	Y	
Log Likelihood	−81.59	−73.61	
McFadden's R^2	0.063	0.155	
McFadden's R^2 adjusted	0.017	−0.098	
AIC	171.19	191.23	
BIC	183.01	256.26	
Observations	142	142	
	Low burden sharing = moderate burden sharing	Moderate burden sharing = high burden sharing	Low burden sharing = high burden sharing
Wald test p values for differences between treatment coefficients			
Model 1	0.6174	0.0500	0.0173
Model 2	0.8136	0.0652	0.0390

Standard errors in parentheses; $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

[a] Demographics include binary variables gender (0 = male, 1 = female), white (Yes = 1), Married (Yes = 1), children (Yes = 1), live with parents (Yes = 1), grad student (Yes = 1), economics student (Yes = 1), full time student (Yes = 1), Christian (Yes = 1), US Citizen (Yes = 1), and republican (Yes = 1). Ordinal demographic variables are father's education level and religiosity. Interval demographic variables are age, family income, and number of bathrooms in the subject's household

[b] Preferences include binary variables such as if the person views him/herself as politically conservative, if the subject believes the US should have a strong leadership role, believes that she/he is politically astute, believes that minorities face discrimination, believes that the amount of immigrants in the US should decrease, believes that the US should be a more isolationist country, and if the subject believes that government spending should increase. Preference variables are based on 10 point Likert scales and include: if the subject views her/himself as a risk taker rather than avoiding risks, as impatient and satisfying immediate rather than long term concerns, and if the subject focuses on taking care of him/herself rather than others

[c] In Model 2, the only demographic and preference independent variables that are statistically significant are "white" at the 0.1 level and "impatience" at the 0.05 level

Footnote 12 continued
a more isolationist country, and if the subject believes that government spending should increase. Preferences variables that include 10-point Likert scales are measures of risk-tolerance, patience and altruism. See Table 2' in Kass et al. (2014), Appendix B (Electronic Supplementary Material), for coefficients on these demographic and personal preference effects.

Table 8 Round 21 (rematched, 10× stakes)

	Prevention	Take no action	Defense
A. Actions selected in Round 21			
No burden sharing	9	2	23
Low burden sharing	4	4	28
Moderate burden sharing	14	0	22
High burden sharing	25	0	11
	No burden sharing	Low burden sharing	Moderate burden sharing
B. Fisher exact test (prevention, take no action, defense), *p* values shown			
Low burden sharing	0.254	–	–
Moderate burden sharing	0.248	0.005	–
High burden sharing	0.001	<0.001	0.017
C. Fisher exact test (prevention vs. non-prevention), *p* values shown			
Low burden sharing	0.129	–	–
Moderate burden sharing	0.315	0.013	–
High burden sharing	<0.001	<0.001	0.017

treatment has a significant effect on the choice of Prevention, as the high burden-sharing coefficient is always positive and significant in the models.[13] The marginal effect suggests that subjects are 36–40 % points more likely to select Prevention in round 1 of the High Burden Sharing treatment than in the baseline, while the other two burden-sharing institutions (Low and Moderate) play at most a weak role in enhancing Prevention.

We now turn to a discussion of the round 21 data. The best way to examine the robustness of the results found in round 1 is to look at the decisions made in the final round, where subjects are re-matched, and stakes are multiplied by ten. Before round 21, they are reminded that they are to be re-matched, and are told that the stakes will be ten times the previous round. This allows us to test the theory on its own domain. The model is one-shot, and round 21 is a one-shot game with experienced subjects. Since they are knowledgeable about the game, their response to the new partner and high stakes gives the theory its best shot.

Tables 8A–C, and 9 presents information parallel to the tables for round 1 data. Table 8A presents the aggregate actions taken by subjects in round 21 and Tables 8B, C contains the nonparametric results from the Fisher Exact test. The results more strongly support the notion that the High Burden Sharing treatment produces differences in behavior. Akin to Table 6C, Table 8C combines the Take No Action and Defense decisions in a general "Non-Prevention" variable, demonstrating that the High Burden Sharing treatment encourages more Prevention choices as compared

13 Multinomial logit regressions, modeling Prevention and Defense choices, showing similar results can be found in the companion paper Kass et al. (2014), Table 3' in Appendix B (Electronic Supplementary Material).

Table 9 Strategy choice, round 21 (rematched, 10× stakes)

Dependent variable: prevention	Model 1	Model 2[b]	Model 3[b]	Model 4[b]
Dependent variable is prevention. Logit regression models, marginal effects shown				
Low burden sharing	−0.193*	−0.163	−0.168	−0.163
	(0.110)	(0.108)	(0.114)	(0.114)
Moderate burden sharing	0.106	0.109	0.057	0.062
	(0.095)	(0.094)	(0.093)	(0.092)
High burden sharing	0.400***	0.455***	0.309***	0.300***
	(0.110)	(0.093)	(0.089)	(0.088)
Prev. history (Rounds 1–20)[a]			0.0162***	
			(0.005)	
Coop history (Rounds 1–20)[a]				0.014***
				(0.004)
Demographics	N	Y	Y	Y
Preferences	N	Y	Y	Y
Log Likelihood	−78.42	−64.32	−56.60	−59.06
McFadden's R^2	0.159	0.310	0.393	0.367
McFadden's R^2 adjusted	0.116	0.075	0.147	0.120
AIC	164.84	172.64	159.20	164.13
BIC	176.67	237.67	227.19	232.11
Observations	142	142	142	142
	Low burden sharing = moderate burden sharing	Moderate burden sharing = high burden sharing	Low burden sharing = high burden sharing	
Wald test p values for differences between treatment coefficients				
Model 1	0.0099	0.106	<0.001	
Model 2	0.0203	0.003	<0.001	
Model 3	0.0321	0.0005	0.0796	
Model 4	0.0256	0.0003	0.0679	
Model 5	0.0548	0.0167	<0.001	
Model 6	0.0455	0.0156	<0.001	

Standard errors in parentheses, panel logit, 21st round data only $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

[a] Regressors as are follows: "Prev. history" is the count of Prevention selections by the subject from Rounds 1–20. "Coop history" is the aggregate number of instances where both subjects selected Prevention in a given round

[b] In Models 2, 5, and 6, the demographic and preference independent variables that are statistically significant (for all three models) include: economics students at the 0.1 level and subject preferences on government spending and attitudes toward risk, both significant at the 0.01 level

to the other treatments. For both experienced and inexperienced subjects, the High Burden Sharing environment appears unique in facilitating socially optimal play.

To account for history and subject characteristics, Table 9 contains several Logit regression models. Prevention is the binary dependent variable, and we control for demographics and personal preferences. As in round 1, we see that only when burden-sharing reaches its maximum level does it significantly increase the probability of Prevention, and subjects are 40–46 % points more likely to select Prevention than in the baseline game (shown in Models 1 and 2, which parallel Table 7). Indeed, Model 1 implies that Low Burden Sharing may be worse than no sharing at all; however, the difference between this treatment and the Baseline becomes insignificant when demographics and preferences are included (Model 2).

To account for experience, Models 3–4 incorporate two alternative ways of including subjects' previous history as an explanatory variable. Model 3 contains "Prev. history", which is defined as the sum of all Prevention choices by a subject from rounds 1–20. Model 4 incorporates the variable "Coop history," defined as the sum of all instances from rounds 1–20 where paired subjects both selected Prevention in a given round. Both models show that game history has a powerful effect on subjects' choices, even though they are matched with a new partner. Thus while higher incentives and rematching reduce the extent to which subjects select Prevention, still they are strongly affected by their previous experience in the game. Including these measures reduces the magnitude of the coefficients on the treatment variables somewhat, but maintains significance. Wald tests are reported at the bottom of Table 9.[14]

The next step in this analysis is to determine whether these treatments affect aggregate earnings. While Prevention is selected more frequently, these gains could be nullified by an increase in the play of Defense, leaving the cost-sharing treatments unable to generate a more positive aggregate social outcome. An earnings regression on round 21 data shows a positive, significant effect of the High Burden Sharing treatment, with subjects earning \$17-\$22 more in that treatment as compared with the baseline.[15]

6.3 Game dynamics

In this section we report panel data analysis in order to better understand the dynamics of play in the repeated Nested PD game. Table 10 shows the marginal effects estimated from a random effects logit model.[16] These models examine the binary decision of whether the subject selected Prevention or not. Model 1 includes only the treatment variables and period dummy variables to capture any time-dependent variation. In this analysis, there is evidence that the Moderate and the High Burden Sharing treatments increase the probability of choosing Prevention. As the Wald tests show, differences between treatments are consistently significant. In the Moderate Burden Sharing treatment subjects are about 20 % points more likely to choose Prevention, and in the High Burden Sharing treatment, 47 % points more likely (relative to the baseline). No sig-

14 Note that removing demographics and preferences from these regressions leaves the results qualitatively unchanged, except that the coefficient on "Low Burden Sharing" become marginally significant, as in Model 1. Table 4' in the Appendix (Electronic Supplementary Material) contains the coefficients on the demographic and personal preference effects from Table 9.

15 See Table 5' in Kass et al. (2014), appendix B (Electronic Supplementary Material) for more details.

16 Marginal effects are calculated in Stata using the "margins" command. Tables 6' and 7' in the companion paper Kass et al. (2014), Appendix B (Electronic Supplementary Material), control for demographics and show that results are robust to different specifications of the models.

Table 10 Strategy choice, rounds 1–20

Dependent variable: prevention	Model 1	Model 2	Model 3	Model 4
Dependent variable is prevention. Random effects logit regression models, marginal effects shown				
Low burden sharing	−0.008	0.0616	0.085	0.058*
	(0.1052)	(0.0775)	(0.0700)	(0.0323)
Moderate burden sharing	0.202*	0.150*	0.202***	0.081**
	(0.1112)	(0.0772)	(0.0638)	(0.0350)
High burden sharing	0.472***	0.250***	0.292***	0.106**
	(0.1429)	(0.0922)	(0.0786)	(0.0449)
Prevention lag (lagPrev)		0.341***		0.223***
		(0.0386)		(0.0229)
LowBS*lagPrev		−0.127**		−0.059*
		(0.0531)		(0.0307)
ModBS*lagPrev		−0.031		−0.027
		(0.0531)		(0.0310)
HighBS*lagPrev		0.040		−0.011
		(0.0735)		(0.0399)
Counterpart prevention lag (lagCP)			0.409***	0.247***
			(0.0406)	(0.0254)
LowBS*lagCP			−0.159***	−0.059*
			(0.0562)	(0.0327)
ModBS*lagCP			−0.150***	−0.066**
			(0.0528)	(0.0334)
HighBS*lagCP			−0.089	−0.069*
			(0.0704)	(0.0418)
Round dummies	Yes	Yes	Yes	Yes
Log likelihood	−863.067	−670.104	−635.218	−555.558
AIC	1,774.133	1,394.209	1,360.436	1,173.116
BIC	1,916.971	1,553.516	1,519.744	1,356.024
McFadden R^2	0.0769	0.2833	0.3206	0.4058
Adj McFadden R^2	0.0523	0.2555	0.2928	0.3738
Observations	2,840	2,698	2,698	2,698
Number of groups	71	71	71	71

nificant differences are found between the Low Burden Sharing treatment and the baseline.[17]

[17] When investigating the dynamics of this finitely repeated Nested PD game, persistence may also be a driver of behavior. As a test for persistence in our data generating processes, Figure 1 in Appendix B (Electronic Supplementary Material) of Kass et al. (2014) contains autocorrelation and partial autocorrelation functions (ACF and PACF) of subjects' Prevention choices indicating that the first lag (time t − 1) is a driver of the current Prevention choice at time t. In each treatment ACFs show an autoregressive process and in 3

Table 10 continued

	Low = moderate	Moderate = high	Low = high
Wald test *p* values for differences between treatment coefficients			
Model 1	0.043	0.010	<0.001
Model 2	0.184	0.022	<0.001
Model 3	0.090	0.290	0.014
Model 4	0.441	0.531	0.224

Standard errors in parentheses $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

[a] Regressors as are follows: "Prevention lag" is a variable whether the subject chose Prevention in the previous round. "LowBS*lagPrev" is an interaction term of the Low Burden Sharing treatment variable with the Prevention lag variable. "ModBS*lagPrev" is an interaction term of the Moderate Burden Sharing treatment variable with the Prevention lag variable. "HighBS*lagPrev" is an interaction term of the High Burden Sharing treatment variable with the Prevention lag variable. "Counterpart Prevention Lag" is a variable whether the subject's counterpart chose Prevention in the previous round. "LowBS*lagCP" is an interaction term of the Low Burden Sharing treatment variable with the Counterpart Prevention Lag variable. "ModBS*lagCP" is an interaction term of the Moderate Burden Sharing treatment variable with the Counterpart Prevention Lag variable. "HighBS*lagCP" is an interaction term of the High Burden Sharing treatment variable with the Counterpart Prevention Lag variable. "Round Dummies" are dummy variables capturing time specific effect. In our setting, the round dummy variables are capturing the drop in Take No Action selections as the subject move through the game

[b] In Model 5, the demographic and preference independent variable to come up as statistically significant is the "racism" variable, a variable measuring subject beliefs about if minorities experience racism. It is statistically significant at the 0.1 level

Models 2–4 incorporate both the subjects' and their counterparts' previous decisions and treatment interaction effects. Model 2 includes own choices and treatment interactions, Model 3 incorporates counterpart choices and interactions, and Model 4 both. We include the choices in the previous round to capture reference point/persistence effects and to account for adjustments to the counterpart's actions. Not surprisingly, there is evidence that subjects prefer to select Prevention when they and/or their counterparts selected Prevention in the previous period; failure to select Prevention in a round reduces the likelihood it will be subsequently selected. Additionally, lags subsume much of the treatment effects, dropping the direct effect of the High Burden Sharing treatment to 10–29 % points. However, even accounting for the lags, the Moderate and High treatment effects are robust. Among the lags, we find the individual and counterpart decision lags enter consistently in the different models, but the lags interacted with the treatments do not always show consistent patterns.[18]

In sum, Table 10 suggests a stronger effect for Moderate Burden Sharing than in the round 1 or 21 data alone. There is still evidence of a nonlinear impact of burden sharing, but not as dramatically as in previous analysis. From the Wald test results,

Footnote 17 continued
out of the 4 treatment PACFs only show a significant PACF at lag 1, indicating that the process has a lag at the previous period, but there is no evidence of a correlation with longer lags.

[18] The only interaction effect that is consistently significant is the counterpart lag interacted with the High Burden Sharing treatment. A potential explanation is that in the High Burden Sharing environment, some subjects may defect to capture additional earnings at time t when the counterpart shows willingness to play optimally at time $t - 1$.

there is less evidence of a clear-cut difference between treatment marginal effects from models that account for choice lags and interaction covariates.

The impact on earnings echoes this analysis. In an ordinary least squares model on total earnings as of round 20, we also find that the Moderate Burden Sharing, along with High Burden Sharing, produce increases in aggregate round 20 earnings in this repeated play environment.[19] These results show a consistent positive effect of the Moderate and High Burden Sharing treatments on earnings. Using Model 1, subjects earn about \$11 more per decision in the Moderate Burden Sharing treatment and \$24 more in the High Burden Sharing treatment than in the other treatments. In general, the direction and significance of results is robust to inclusion of demographics and preferences.

6.4 Cooperation as mutual prevention

While the individual decision data provide a measure of individual choices, it seems natural to analyze how different pairs of participants coordinate across the different symmetric solutions. We use the term coordination, even when mutual Prevention is never an equilibrium of the one-shot stage game. Exploring cooperation as mutual Prevention has an advantage, as the data are now cross-sectionally independent. Figure 2 presents the general picture. Aggregate mutual Prevention in rounds 1–20 increases with burden sharing. However, as we observed with individual observations on Prevention, coordination on a mutual strategy in round 21 is more difficult. Pairs are likely to coordinate on mutual Prevention only when the burden sharing is High, where almost 45 % of the pairs reached mutual cooperation.[20] Mutual Prevention is never observed in the baseline or the Low Burden Sharing treatment (therefore metric comparisons between baseline and treatments are not possible) and it rarely happens with Moderate Burden Sharing.

Table 11 contains random-effects regression models, and incorporates the previous ability of the pair to cooperate (as a test of persistence, Model 2) and the interaction terms of previous cooperation with the treatment dummies (Model 3).[21] As with the individual Prevention decisions, treatment effects are weakened by including the lag. Still, the High Burden Sharing marginal effect is significant and positive in each model, but the Moderate burden sharing treatment is only significant in Models 1 and 3. The Low burden sharing treatment is never significantly different from the baseline, supporting again the non-linear effect of burden sharing.

[19] See Table 8' in Appendix B (Electronic Supplementary Material), Kass et al. (2014).

[20] Since close to 70 % of subjects in round 21 chose Prevention in the High Burden Sharing treatment, the predicted value of two subjects forming a cooperative pair is around 49 %. Therefore 45 % of successful coordination on Prevention is very sensible given the predictions.

[21] Figure 2 in Appendix B (Electronic Supplementary Material) of Kass et al. (2014) shows autocorrelation and partial autocorrelation functions (ACF and PACF) of coordination. Except for the Low Burden Sharing treatment, ACFs show a clear indication of an autoregressive process and these PACFs only show a significant PACF at lag 1,indicating that the process has a lag at time $t - 1$.

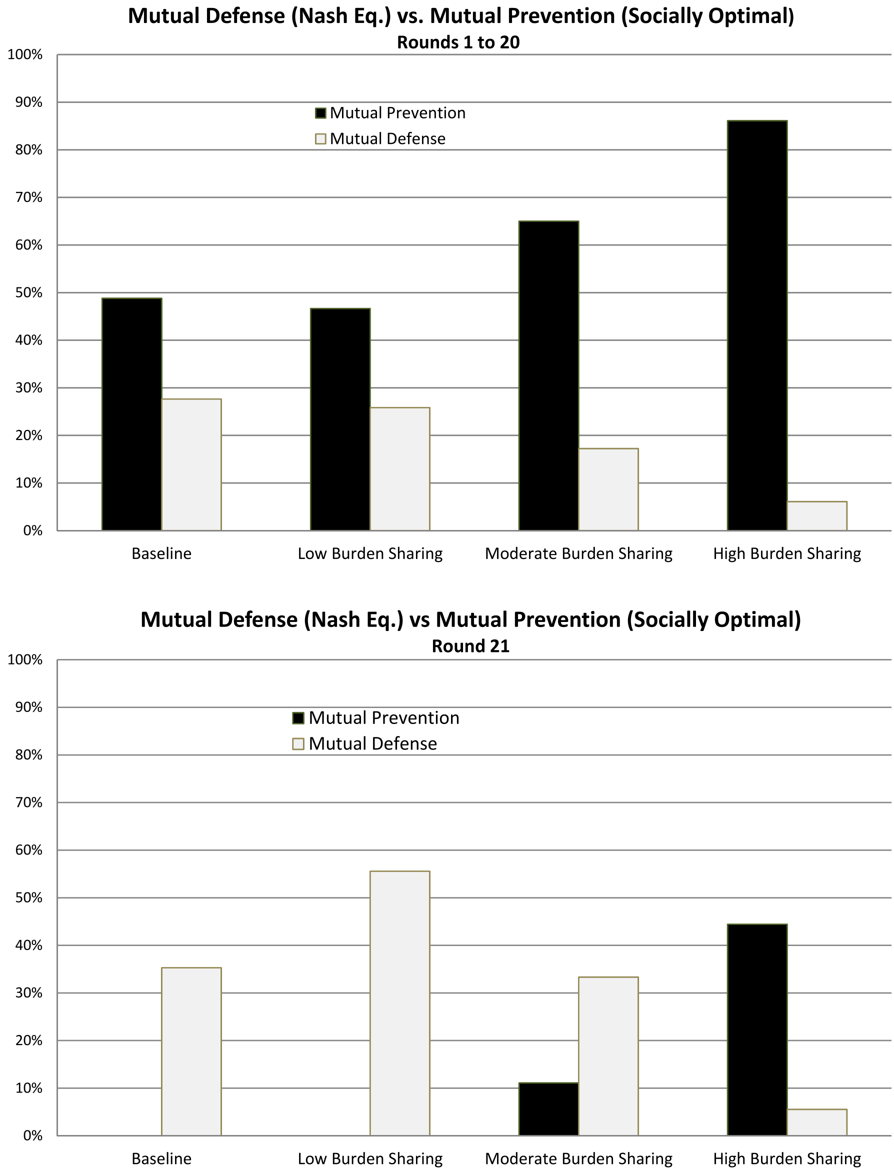

Fig. 2 symmetric choices in rounds 1–20 and round 21

7 Conclusions

In this paper we test the effect of burden sharing on the ability of countries to successfully coordinate on efficient counterterrorism policies. We use a laboratory experiment to test, on its own domain, a simple theoretical model first described in Arce and Sandler (2005a,b). In that model the policy alternatives faced by countries are combined

Table 11 Strategy coordination, rounds 1–20

Dependent variable: mutual prevention (cooperation)	(1)	(2)	(3)
Dependent variable is mutual prevention. Random effects logit regression models, at the pair level, marginal effects shown			
Low burden sharing	0.032	−0.003	0.060
	(0.1889)	(0.0459)	(0.0522)
Moderate burden sharing	0.351**	0.062	0.146***
	(0.1704)	(0.0506)	(0.0570)
High burden sharing	0.780***	0.115*	0.134**
	(0.1869)	(0.0600)	(0.0699)
Lagged cooperation		0.324***	0.429***
		(0.0322)	(0.0635)
LowBS*lagCoop			−0.135**
			(0.0687)
ModBS*lagCoop			−0.190***
			(0.0656)
HighBS*lagCoop			−0.094
			(0.0736)
Round dummies	Yes	Yes	Yes
Log likelihood	−303.598	−207.951	−202.419
AIC	655.195	463.903	458.837
BIC	781.397	588.874	599.430
McFadden R^2	0.2817	0.5080	0.5211
Adj McFadden R^2	0.2272	0.4535	0.4595
Observations	1,420	1,349	1,349
Number of group	71	71	71
	Low = moderate	Moderate = high	Low = high
Wald test p values for differences between treatment coefficients			
Model 1	0.142	0.061	<0.001
Model 2	0.156	0.270	0.024
Model 3	0.057	0.829	0.219

Standard errors in parentheses $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Regressors are as follows: “Cooperation lag” is a variable whether paired subjects mutually select Prevention in the previous round. “LowBS*lagCoop” is an interaction term of the Low Burden Sharing treatment variable with the Cooperation lag variable. “ModBS*lagCoop” is an interaction term of the Moderate Burden Sharing treatment variable with the Cooperation lag variable. “HighBS*lagCoop” is an interaction term of the High Burden Sharing treatment variable with the Cooperation lag variable. The positive log likelihoods are from narrow dispersion parameters by incorporating the lagCoop variable

into a single 3×3 matrix game, which embodies both the Prisoner’s Dilemma structure of Prevention policies and the Commons structure of Defensive policies.

Interestingly, this simple game captures well the behavioral tension between choosing Defense, at the collective price of increasing the cost of other players, and choosing

Prevention, at the cost of playing a dominated strategy, and risking free-riding by others. Prior experimental research using this game has shown that, while subjects are more successful than theory predicts at coordinating on the non-equilibrium, welfare-efficient strategy of Prevention, nevertheless policy coordination frequently fails (Arce et al. 2011a,b).

In this paper we test the effectiveness of one policy for enhancing the ability of players to coordinate on mutual Prevention: burden sharing of the cost of Prevention. The behavioral logic is simple. By implementing this policy the cost of choosing Prevention is shared with those players who benefit from it. As it is hard to imagine running an empirical test of this particular variation of the model with field data, we think the lab provides an ideal venue for wind-tunnel testing of the strategic response of subjects to enforced burden sharing.

The experiment tests three different levels of cost sharing: Low, Moderate, and High, corresponding to a shifting of up to half of the cost of Prevention to the counterpart. The cost of choosing Prevention decreases with the level of burden sharing, as Low Burden Sharing shifts one unit of the six-unit cost of Prevention onto the counterpart, Moderate Burden Sharing shifts two units, and High shifts three, or half, of the cost of Prevention.

An important feature of our design is that the burden-sharing level always falls below the level that would alter the equilibrium structure of the game. Defense is a dominant strategy in each experimental treatment. By reducing the strategic risk of choosing Prevention by decreasing its cost, we measure the corresponding impact on players' performance. We conjectured that burden sharing would monotonically serve to endorse the Prevention strategy, especially with the high level of contextual framing in the instructions. However, we find that the effect is quite discontinuous. While Low Burden Sharing has a negligible or negative impact on the choice of effective policies, the largest amount of burden sharing has a robust, substantial and significant effect on the ability to coordinate effective counter-terrorism policy.

A potential cause for this non-linear effect is that while the different levels of burden sharing in our study do not affect the Nash equilibrium of the Nested PD game, they do alter the payoffs of the game when coordination fails. A subject who selects Prevention is partially protected from the lowest payoffs in the event that coordination is unsuccessful. Moreover, the equilibrium prediction for the underlying 2×2 PD Prevention game actually changes when burden sharing is introduced. When the level of burden sharing is moderate or high, not choosing Prevention is no longer a dominant strategy. For the High Burden Sharing treatment, mutual Prevention becomes the unique equilibrium of this 2×2 game.

It should be noted that this is a different "solution" to the one proposed by Arce and Sandler (2005a,b), which requires one of the nations to be an asymmetrically favored target, by the terrorists. Under asymmetry, only the targeted nation prevents. What our experiment derives is more of a "UN solution." This is relevant because in the post-9/11 world terrorists have shown themselves to be willing to substitute one nation for another as a target. We identify in this paper a policy modifying the Prevention PD game, which then spills over experimentally to the 3×3 game, even though the Nash equilibrium of the 3×3 game is never changed.

This paper contributes to a growing literature on the behavioral analysis of conflict and security issues. Abbink (2012) is a nice and recent survey of the experimental analysis of conflict. Relative to other recent papers, our research tests the effectiveness of cost sharing in a deterministic environment. In that sense, our experiment tests the effectiveness of a cost sharing policy when known behavioral biases, as overweighting small probabilities, play no role. Keser and Montmarquette (2008) find that when subjects have the chance of reducing the likelihood of a bad outcome for all players they overweight small probabilities. Similarly, Cadigan and Schmitt (2010) test an environment where players can reduce the likelihood of attack by investing in protection. They find lower than predicted spending when protection generates a negative spillover. Relative to Colombier et al. (2011), our cost sharing policy is implemented vertically (as a rule of the game, exogenously imposed by design) rather than horizontally (as their sanctions or rewards, implemented by individual players). Maybe not surprisingly, rewards and sanctions increase investments in the international collective actions in a similar way to our cost sharing "UN". In that sense, we think our approach complements well theirs.

We believe this paper makes a significant contribution to the behavioral analysis of conflict resolution. By testing in the laboratory the implications of different policies, we identify an interesting and unexpected non-linearity. Burden sharing per se does not have the capability to significantly enhance coordination on the socially efficient strategy of Prevention, but higher levels of burden sharing do. Our results strongly suggest that behavioral tests are useful to understand subtle changes in the institutional environment.

Acknowledgments Experiments were conducted in the Center for Behavioral and Experimental Economic Science at the University of Texas at Dallas. Thanks to Wendy Lee Zhengzheng Wang, and Ravi Hanumara for programming and research assistance. The research was supported by the National Science Foundation (BCS-0905060). Enrique Fatas acknowledges the support of the ESRC Network for Integrated Behavioral Sciences (NIBS). We are especially grateful to the participants in the Workshop 'Experiments on Conflict, Inequality and Incentives' held at Rennes May 23-25, 2012 for very helpful comments.

References

Abbink K (2012) Laboratory experiments on conflict. In: Garfinkel MR, Skaperdas S (eds) Oxford handbook of the economics of peace and conflict. Oxford University Press, Oxford, pp 532–556

Abbink K, Hennig-Schmidt H (2006) neutral versus loaded instructions in a bribery experiment. Exp Econ 9:103–121. doi:10.1007/s10683-006-5385-z

Ahn TK, Ostrom E, Schmidt D, Shupp R, Walker J (2001) Cooperation in PD games: fear, greed, and history of play. Public Choice 106:137–155. doi:10.1023/A:1005219123532

Andreoni J, Varian H (1999) Pre-play contracting in the Prisoners' Dilemma. Proc Natl Acad Sci USA 96:10933–10938. doi:10.1073/pnas.96.19.10933

Arce DG, Sandler T (2005a) Counterterrorism: a game-theoretic analysis. J Confl resolut 49:183–200. doi:10.1177/0022002704272863

Arce DG, Sandler T (2005b) The dilemma of the prisoners' dilemmas. Kyklos 58:3–24. doi:10.1111/j.0023-5962.2005.00275.x

Arce DG, Croson RTA, Eckel CE (2011a) Terrorism experiments. J Peace Res 48:373–382. doi:10.1177/0022343310391502

Arce DG, Bakshi S, Croson RTA, Eckel C, Fatas E, Kass M (2011b) Counterterror strategies in a lab. Public Choice 149:465–478. doi:10.1007/s11127-011-9884-x
Axelrod R (1984) The evolution of cooperation. Basic Books, New York
Bracht J, Figuieres C, Ratto M (2008) Relative performance of two simple incentive mechanisms in a public goods experiment. J Public Econ 92:54–90. doi:10.1016/j.jpubeco.2007.04.005
Cadigan J, Schmitt PM (2010) Strategic entry deterrence and terrorism: Theory and experimental evidence. Public Choice 143:3–22. doi:10.1007/s11127-009-9488-x
Charness G, Rabin M (2002) Understanding social preferences with simple tests. Q J Econ 117:817–869. doi:10.1162/003355302760193904
Charness G, Frechette GR, Qin CZ (2007) endogenous transfers in the prisoner's dilemma game: an experimental test of cooperation and coordination. Games Econ Behav 60:287–306. doi:10.1016/j.geb.2006.10.010
Chen Y (2003) An experimental study of serial and average cost pricing mechanisms. J Public Econ 87:2305–2335. doi:10.1016/S0047-2727(02)00017-8
Colombier N, Masclet D, Mirza D, Montmarquette C (2011) Global security policies against terrorism and the free riding problem: an experimental approach. J Public Econ 13:755–790. doi:10.1111/j.1467-9779.2011.01519.x
Eckel CE, Grossman PJ (1996) Altruism in anonymous dictator games. Games Econ Behav 16:181–191
Fischbacher U (2007) z-Tree: Zurich toolbox for ready-made economic experiments. Exp Econ 10:171–178. doi:10.1007/s10683-006-9159-4
Fong YF, Surti J (2009) The optimal degree of cooperation in the repeated prisoners' dilemma with side payments. Games Econ Behav 67:277–291. doi:10.1016/j.geb.2008.11.004
Greiner B (2004) An online recruitment system from economic experiments. Forschung und Wissenschaftliches Rechnen 63:79–93
Harden G (1968) The tragedy of the commons. Science 162:1243–1248
Kass M, Fatas E, Eckel C, Arce DG (2014) The Un in the lab. CBESS working paper, University of East Anglia
Keser C, Montmarquette C (2008) Voluntary contributions to reduce expected public losses. J Econ Behav Organ 66:477–491. doi:10.1016/j.jebo.2006.06.007
Kreps DM, Wilson R (1982) Sequential equilibria. Econometrica 50:863–894
Liberman V, Samuels SM, Ross L (2004) The name of the game: predictive power of reputations versus situational labels in determining prisoner's dilemma game moves. Pers Soc Psychol Bull 30:1175–1185. doi:10.1177/0146167204264004
Loomes G (1999) Some lessons from past experiments and some challenges for the future. Econ J 109:35–45. doi:10.1111/1468-0297.00401
Luce RD, Raiffa H (1957) Games and decisions: introduction and critical survey. Courier Dover Publications, New York
Matsushima H (2012) Finitely repeated prisoner's dilemma with small fines: the penance contract. Japn Econ Rev 63:333–347. doi:10.1111/j.1468-5876.2011.00544.x
Olson M (1965) The logic of collective action: public goods and the theory of groups. Harvard University Press, Cambridge
Rapoport A, Chammah AM (1965) Prisioner's dilema: a study in conflict and cooperation. University of Michigan Press, Ann Arbor
Rapoport A (1967) A note on the" index of cooperation" for prisoner's dilemma. J Confl Resolut 11:100–103
Razzolini L, Reksulak M, Dorsey R (2007) An experimental evaluation of the serial cost sharing rule. Theor Decis 63:283–314. doi:10.1007/s11238-007-9038-2
Sally D (1995) conversation and cooperation in social dilemmas: a meta analysis of experiments from 1958 to 1992. Ration Soc 7:58–92. doi:10.1177/1043463195007001004
Sander T (2003) Collective action and transnational terrorism. World Econ 26:779–802. doi:10.1111/1467-9701.00548
Sandler T (2005) Collective versus unilateral responses to terrorism. Public Choice 124:75–93. doi:10.1007/s11127-005-4747-y
Sandler T, Arce DG, Enders W (2011) An evaluation of INTERPOL's cooperative-based counterterrorism linkages. J Law Econ 54(1):79–110. doi:10.1086/652422
Sandler T, Siqueira K (2006) Global terrorism: deterrence versus pre-emption. Can J Econ 39:1370–1387. doi:10.1111/j.1540-5982.2006.00393.x

Shafir E, Tversky A (1992) Thinking through uncertainty: nonconsequential reasoning and choice. Cogn Psychol 24:449–474. doi:10.1016/0010-0285(92)90015-T

Smith VL (1980) Experiments with a decentralized mechanism for public good decisions. Am Econ Rev 70:584–599