Tjalling C. Koopmans Research Institute

Discussion Paper Series nr: 12-24

Cooperation and the common enemy effect

Kris De Jaegher
Britta Hoyer

Universiteit Utrecht

Utrecht School of Economics

Tjalling C. Koopmans Research Institute
Utrecht School of Economics
Utrecht University

Kriekenpitplein 21-22
3584 EC Utrecht
The Netherlands
telephone +31 30 253 9800
fax +31 30 253 7373
website www.koopmansinstitute.uu.nl

The Tjalling C. Koopmans Institute is the research institute and research school of Utrecht School of Economics.
It was founded in 2003, and named after Professor Tjalling C. Koopmans, Dutch-born Nobel Prize laureate in economics of 1975.

In the discussion papers series the Koopmans Institute publishes results of ongoing research for early dissemination of research results, and to enhance discussion with colleagues.

Please send any comments and suggestions on the Koopmans institute, or this series to J.M.vanDort@uu.nl

ontwerp voorblad: WRIK Utrecht

How to reach the authors

Please direct all correspondence to the first author.

Kris De Jaegher
Utrecht University
Utrecht School of Economics
Kriekenpitplein 21-22
3584 TC Utrecht
The Netherlands.
E-mail: k.dejaegher@uu.nl

Britta Hoyer
Universität Paderborn
Fakultät für Wirtschaftswissenschaften
Lehrstuhl für Institutionenökonomik und Wirtschaftspolitik
Warburger Str. 100
33089 Paderborn
Germany
E-mail: britta.hoyer@wiwi.uni-paderborn.de

Utrecht School of Economics
Tjalling C. Koopmans Research Institute
Discussion Paper Series 12-24

Cooperation and the common enemy effect

Kris De Jaegher[a]
Britta Hoyer[b]

[a]Utrecht University School of Economics, Utrecht University

[b]University of Paderborn

December 2012

Abstract
This paper presents a game-theoretic rationale for the beneficial effect of a common enemy on cooperation. In a defence game against a common natural threat, the value of the public good of defence is equal to the sum of the players' defensive efforts. The game therefore takes the form of a prisoner's dilemma, leading to free-riding. When the same defence game is played against a common enemy, the value of the public good of defence is equal to the smallest defensive effort. The game now takes the form of a stag hunt, so that a cooperative equilibrium becomes possible. For this reason, an informed and benevolent government may not want to inform the public that it is facing a common natural threat rather than a common enemy. At the same time, the common enemy has an incentive to mimic nature, and perform only random rather than targeted attacks.

Keywords: Common Enemy Effect; Defence Games; Prisoner's Dilemma; Stag Hunt.

JEL classification: D74, H41, C72

1. Introduction

A diverse literature treats what may be termed the *common enemy effect*. In *sociology*, Simmel (1955) and Coser (1956), in a hypothesis known as the *in-group out-group hypothesis*, argue that the presence of a common enemy leads to more cohesion in groups (for a contemporary example, see McLauchlin and Pearlman, 2012). In *political science*, the common enemy effect has been used to explain international alliances between countries (Waltz, 1979). Further, following an argument already made by Coser (1956), it is argued in political science that governments may seek to exploit the common enemy effect by looking for a common enemy in order to increase cohesion among citizens. This has been referred to alternatively as the theory of *diversionary foreign policy* or as the *scapegoat hypothesis* (see Levy, 1989; for a recent treatment, see Sirin, 2011).

In *social psychology*, attempts have been made to provide micro-foundations for the common enemy effect put forward in sociology. For instance, in his *balancing theory*, Heider (1946) argues that when two enemies *A* and *B* have a common enemy *C*, this creates cognitive dissonance, as when *A* hurts *B*, this serves the interests of *A*'s enemy *C*. For this reason, *A* and *B* may become friends (for a contemporary and dynamic treatment, see Antal et al., 2006). In *evolutionary psychology*, a version of the common enemy effect is found in the argument that, while the evolution of altruism may not be possible in a single group, it may evolve when several groups are competing, in a group-selection model (Boyd and Richerson, 1990).

In *economics*, to our knowledge little attention has been paid to the common enemy effect. A first exception is a recent theoretical paper by Hoyer (2012), who considers a model of network formation where players bilaterally decide whether to form links for sharing information, given that they obtain information both from their direct and their indirect links (cf. Jackson and Wolinsky (1996)). Hoyer adds a common enemy to this game, who can disrupt links in the network in order to minimise information sharing, and shows that network formation may only take place conditional on the presence of a common enemy. The intuition is that in the presence of a common enemy, links not only ensure that information sharing takes place, but also that the network cannot be disrupted by the common enemy. This may cause an increase in the benefit of information sharing, which makes network formation more attractive. A second exception in economics is found in two laboratory experiments on the effect of intergroup competition on coordination in games with both cooperative and non-cooperative equilibria, but where the non-cooperative equilibria are more risky to players (Bornstein et al. (2002), Riechmann and Weimann (2008)). In the experiments, such games are played among groups of participants, where in the baseline treatment the participants' payoffs do not depend on the performance of other groups, and in the competitive treatments participants receive a higher payoff if they perform better than a competing group. The experiments show that participants are more likely to coordinate on cooperative equilibria in the competitive treatments, leading the authors to conclude that competition turns cooperative equilibria into focal points (Schelling, 1960). A third and final exception in economics is found in Kovenock and Roberson (2012). In this paper, two players *A* and *B* each play a separate, two-player Colonel Blotto game against the same player *C*. In each Colonel Blotto game, the two players allocate their respective endowments over n battlefields. In each battlefield, the player allocating more resources wins, and players are better off the more battlefields they win. If player *A* has more resources than player *B*, then the fact that C is a common enemy may create an equilibrium where player *A* transfers some of his endowment to player *B*, where both *A* and *B* are better off than if such transfers would not be possible.

Finally, in *evolutionary biology*, Mesterton-Gibbons and Dugatkin (1992) consider a standard two-player prisoner's dilemma, and consider both the benefit of acting rather than not acting when the other player acts (benefit 1), and the benefit of acting rather than not acting when the other player does not act (benefit 2), where benefit 1 is larger than benefit 2. In the prisoner's dilemma, both these benefits are negative. The authors argue that if the degree of environmental adversity increases, the first benefit to become positive is benefit 1, such that the game is turned into a stag hunt. As environmental adversity is further increased, both benefits become positive, and joint action is the only equilibrium. This argument is known in evolutionary biology as *byproduct mutualism* (see Mesterton-Gibbons, 2001, Section 5.8). While the authors refer to environmental adversity as a common enemy, no differentiation is made on whether this environmental adversity comes from a strategic player, or from the physical environment.[1]

With the exception of economics and evolutionary biology, the common theme in this literature is that when a group faces a common enemy, this changes the psychology or the preferences in this group. The purpose of this paper is to show that, if cooperation consists of defence against a common threat, a simple rational explanation for the common enemy effect exists. The difference between our argument and the argument of Bornstein et al. (2002) and Riechmann and Weimann (2008) is that the presence of a common enemy, rather than causing coordination on cooperative equilibria that are present with or without a common enemy, creates cooperative equilibria where none existed before. Further, our model differs from Mesterton-Gibbons and Dugatkin (1992) in that we predict cooperation only when the players face a strategic common enemy, and not when they face a common natural threat. Finally, contrary to Kovenock and Roberson (2012), in our model cooperation caused by the presence of a common enemy does not require that the cooperators are in asymmetric positions, and make transfers between each other.

Intuitively, consider a group of players who benefit from a common resource, but face possible destruction of this common resource when not taking defensive measures. In the basic variant of our model, attack against a *single* player destroys the *entire* common resource, unless the player takes a defensive measure. On the positive side, a single protected player is able to defend the entire common resource, if this player happens to be the one to be attacked. As examples of a common resource which is vulnerable even there is only a single player which fails to defend it, we may think of the following. An accident in a single nuclear plant in one country may damage the environment in several surrounding countries as well. In neighbouring countries sharing a coast line, flooding of one country may mean that several countries are flooded. In a population facing disease, one individual who does not take preventive measures against disease, such as vaccination, may lead the entire population to be infected. In a price cartel among competing firms, failure of one firm to take measures to keep the cartel secret may mean that the entire cartel is destroyed, if the competition authority happens to audit this firm. Finally, in a distributed network, such as a road network or computer network, failure of one link can make the entire network break down.

When the attack on the common resource has a natural cause, such as an accident, the target of the attack is random. Let it typically be the case that, if an attack takes place, it is targeted at only a single player (this occurs if the probability of an attack is small, so that the probability of more than one player being attacked is vanishing), and take as a starting point the situation where every player takes defensive measures. Then the individual player runs little risk from deviating and not taking defensive measures, as the probability that he will be

[1] Recently, Arenas et al. (2011) treat an evolutionary public good game, which differs only from a standard public good game in that next to the standard strategies of cooperating and defecting, there is a third strategy of inflicting damage on other players, which yields a zero payoff to the player inflicting the damage, but reduces the payoff of other players. The presence of such a strategy ensures the stability of a mixed equilibrium where players mix between cooperating, defecting, and inflicting damage. In this model, it is the potential cooperators themselves who may become common enemies. Intuitively, inflicting damage serves as a punishment strategy, which ensures that some cooperation is still possible.

targeted is small. For this reason, cooperation in taking defensive measures breaks down, leading to an inefficient equilibrium where no player takes defensive measures (Section 2). Consider now instead a situation where the players face an attack by an intelligent adversary, the common enemy. Then this common enemy will always attack a player who did not take defensive measures. Take as a starting point a situation where all players take defensive measures. The fact of facing an intelligent common enemy makes it very costly for a single player to deviate and no longer take defensive measures, as this leads to complete destruction of the common resource. Players may now coordinate on the efficient equilibrium. Compared to the situation where attack has a natural cause and is not caused by a strategic player, players thus benefit from the presence of a common enemy (Section 3).

We provide two applications of the difference between the situations where attack has a natural cause, and where attack comes from an intelligent common enemy. In Section 4, we first show that, if players cannot always identify whether they are playing a game against a common enemy, the possibility of facing a common enemy may induce cooperation in games where no such enemy is faced. Next, we extend this idea by considering a game between a benevolent government which cares about the overall wellbeing of a population, and the mentioned population which owns a common resource and does not know whether this common resource faces an attack that has a natural cause, or an attack by an intelligent common enemy. It is shown that as long as the players believe that the possibility of a common enemy is not too small, in equilibrium the benevolent government does not inform the population whether they are facing a common natural threat or an intelligent common enemy. The government may then be seen as maintaining the idea that there might be a common enemy in order to boost cooperation, yielding a rationale for the scapegoat hypothesis of political science. In Section 5, we show that, as a random attack leads to a better result for the attacker than a targeted attack, an intelligent attacker can prevent cooperation by credibly committing herself to only performing random attacks, where the attack remains random even if only a single player deviates, and where the common enemy simulates the behaviour of a common natural threat. Strategic delegation to a blind and random attacker may here serve to the common enemy as a commitment device. Section 6 contains robustness results in terms of some variants of the model, which leads to insights in the exact conditions when the common enemy effect applies. We end with a discussion in Section 7.

2. Model 1: Attack by Nature

Consider the following two-player game. At stage 1, two players, referred to as Defenders, simultaneously decide whether or not to take defensive measures against a potential attack by Nature. At stage 2, Nature randomly attacks one of the Defenders, where each Defender is attacked with equal probability (note that Nature is only referred to here as a player for expositional reasons). Finally, both Defenders obtain their payoffs. If Nature's attack was targeted at a Defender who took defensive measures, each Defender obtains payoff V, namely the value of the common resource. If Nature's attack was targeted at a Defender who did *not* take defensive measures, the common resource is destroyed, and both Defenders obtain payoff zero. Put otherwise, the defensive measures of a single Defender i protect both Defenders if Nature's attack happens to be targeted at Defender i; if Nature's attack is instead targeted at Defender j, and if this Defender did not take defensive measures, then the defensive measures of Defender i have no effect, and the common resource is destroyed. A Defender who takes defensive measures additionally incurs cost C. We assume that $C < V < 2C$, so that the game against Nature takes on the form of a standard prisoner's dilemma, with a unique Nash equilibrium where no Defender takes defensive measures, and where each

Defender obtains payoff zero – this in spite of the fact that both Defenders would be better off if they would both take defensive measures. This Prisoner's Dilemma is represented in Table 1.[2]

	Act	Don't
Act	$V - C, V - C$	$V/2 - C, V/2$
Don't	$V/2, V/2 - C$	0, 0

Table 1. Payoff matrix against Nature (Prisoner's Dilemma)

In an n-player extension of this game, it is again assumed that Nature can attack one random Defender, so that each Defender is attacked by Nature with probability $1/n$. If defensive measures are taken by a single Defender, this suffices for all n Defenders to be protected. It follows that if n_1 Defenders take defensive measures, the expected payoff of a Defender who takes such measures equals $V^*(n_1/n) - C$. It is optimal for this Defender to deviate if $V^*(n_1 - 1)/n > V^*(n_1/n) - C$ or $V < nC$. Thus, for $V < nC$, the n-player game against Nature is an n-player Prisoner's Dilemma with a unique Nash equilibrium where no player acts. The condition for having a unique non-cooperative Nash equilibrium is thus less tight the more players there are. This is intuitive, as the probability that the individual player who does not take defensive measures is targeted, is smaller the more players there are.

It should be noted that the game in Table 1 takes exactly the same form as a classical public-good game where the amount of the public good produced equals the sum of the players' contributions (in Hirshleifer's (1983) taxonomy of social composition functions, this is a *summation function*), so that the players' contributions are perfect substitutes in the production function of the public good.[3] To see this, let each player be able to contribute V/n to the production of a public good, where if n_1 players contribute, $V^*(n_1/n)$ of the public good is produced. If the cost of making a contribution with a value of V/n equals C, and if $C > V/n$, then in equilibrium zero units of the public good are produced. In our model, the public good consists of the defence of a common resource, where a contribution to the common good consists of a reduction of the probability that the common resource is destroyed by $1/n$.

3. Model 2: Attack by Informed Attacker

Next consider a variant of the game above where at stage 2, it is not Nature which moves to attack one of the players, but a third player, the Attacker, who can choose either to attack Defender 1 or Defender 2, after having observed for each Defender whether or not the Defender took defensive measures at Stage 1. The Attacker is assumed to obtain minus the sum of the payoffs of the Defenders, where additionally the Attacker incurs a cost K when attacking, with $2V > K$ (in which case, if at least one Defender does not take defensive measures, the Attacker prefers attacking and obtaining payoff $0 - K$ to not attacking and obtaining payoff $-2V$). Table 2 presents now, for the two Defenders, each strategy profile with the payoffs adjusted for the response that the Attacker adopts after observing the Defenders´ strategy profile. It is best response for the Attacker not to attack if both Defenders adopt

[2] As the dominant strategy is not to take defensive measures, it does not matter whether the Defenders adopt their strategies simultaneously, or sequentially after observing the choice made by the other player.

[3] While the Defenders' defensive production function can be seen as a weakest-link production function, as attack on the weakest link completely destroys the common resource, the fact that Nature only performs blind attacks turns the defensive production function into a summation function.

defensive measures. If only Defender *i* takes defensive measures, then it is a best response for the Attacker to attack Defender *j*. If none of the Defenders take defensive measures, then the Attacker may attack any of them.

	Act	Don't
Act	$V - C, V - C$	$-C, 0$
Don't	$0, -C$	0, 0

Table 2. Payoff matrix Game against Informed Attacker (Stag Hunt)

As can be seen from Table 2, the presence of a common enemy to the Defenders in the person of an Attacker, turns the game into a Stag Hunt rather than a Prisoner's Dilemma. As the Attacker can identify who does not take defensive measures, not taking such measures is a safe action, as the common resource is then always destroyed. Seeing the game as a Stag Hunt, not taking defensive measures is the equivalent of Hunting Hare. Taking defensive measures is a risky action, as it involves a cost, and only yields a high benefit if the other player also takes the risky action. In Stag Hunt terms, taking defensive measures is the equivalent of Hunting Stag. As any Stag Hunt, the game now has a payoff dominant but risk dominated equilibrium where players coordinate on taking the risky action, and a payoff dominated equilibrium but risk dominant equilibrium where players coordinate on taking the safe action.[4] Risk dominance suggests that players will coordinate on the payoff dominated equilibrium, i.e. will not take defensive measures. Yet, as pointed out by Devetag and Ortmann (2006, p.342) in their literature review of laboratory experiments on coordination games, "while coordination failures are common in the lab, they are by no means ubiquitous."

An *n*-player generalisation is easily obtained if we continue to assume that an attack on a single player who does not take defensive measures reduces the value of the common resource to zero. As soon as at least one Defender does not take defensive measures, the Attacker can attack such a Defender, so that all Defenders obtain value zero from the common resource. As taking defensive measures is costly, as soon as one Defender does not take defensive measures, it is a best response for all players not to take defensive measures. Moreover, if all players take defensive measures, a single player who deviates sees his value of the common resource reduce from V to 0. It follows that it continues to be the case that the game has a payoff dominant equilibrium where all players take defensive measures, and a payoff dominated equilibrium where no player takes defensive measures.

A remarkable fact is now the following. Compared to the game in the previous section, nothing has changed, except that the players are now facing an intelligent Attacker who strategically looks for the weakest point in the defence. At first sight, it would seem that facing an intelligent adversary instead of a common natural threat could not possibly make the Defenders better off. For any given strategy profile chosen by the Defenders at stage 1, the Attacker would seem to have a second-mover advantage when observing this strategy profile, as she can then attack any Defender who did not take defensive measures. Yet, the point is that the Defenders choose different strategies when they are facing an Attacker. Because of this, what appears at first sight a second-mover advantage for the Attacker, namely observing who took defensive measures, is instead turned into a second-mover disadvantage: if the Attacker had not been able to observe who took defensive measures, there would have been no cooperation and she would have been better off.

[4] The distinction between payoff dominance and risk dominance is due to Harsanyi and Selten (1988).

The reason for this result is the following. When facing blind Nature, failure of a single Defender to take defensive measures only leads to a bad outcome with small probability, given that Nature will only attack the deviating Defender's position with small probability. The punishment of deviating is therefore small. When facing an intelligent Attacker, however, failure of a deviating Defender to take defensive measures leads to an attack exactly on this Defender, and to a loss of the full resource. For this reason, the individual Defender does not have any incentive to deviate from an equilibrium situation where all Defenders take defensive measures. Inadvertently, by playing her best response and always targeting a weakest link among the Defenders, the Attacker is punishing defectors and thereby enforcing cooperation. In terms of public goods, by being intelligent and knowing which Defenders do not take defensive measures, the Attacker turns what was before a public good determined by the sum of the players' inputs, i.e. with perfect substitutability between defensive measures, into a public good produced by means of a production function that counts only the smallest input level, i.e. with perfect complementarity between defensive measures. In terms of Hirshleifer's (1983) taxonomy of public good games, the latter corresponds to *weakest link* social composition function (on specific applications of this taxonomy to defence games, see Sandler, 2006).

Noting that Nature can be seen as an Attacker who is not able to observe which Defenders took defensive measures (and who effectively plays a simultaneous move game against the Defenders), another way of stating our result is that the Attacker may be worse off when having more information. It is a standard result in game theory that with strategic interaction, the value of information may be negative (Hirshleifer, 1971). If, as we assume in this section, the Attacker is not able to commit herself to not knowing which Defenders took defensive measures, she cannot credibly commit herself to not attacking the first Defender who does not take defensive measures. For this reason, knowledge of the Defenders' actions is a poisoned gift that makes the Attacker the ideal enforcer of collective action, by punishing deviations from collective action. The Defenders are made better off by the presence of an intelligent common enemy. A common enemy effect is obtained here without any appeal to changes in the Defenders' preferences or psychology when an intelligent common enemy is present: they continue to act in an equally self-interested manner.

4. Spillover common enemy effects, and the scapegoat hypothesis

The common enemy effect of Sections 2 and 3 describes such an effect as arising in the specific context of defence games, where players only cooperate in defence. Yet, the literature reviewed in the Introduction suggests that the common enemy effect should have a much broader effect on cooperation, beyond joint defence against the common enemy. We here treat a simple variant of the model in Sections 2 and 3, where the players do not know whether or not they are playing a game against a common enemy. As recently argued by Mengel (forthcoming), players may find it costly to fine-tune their strategies to every different game that they may be playing. For this reason, they may partition the games which they play into broad categories, where for all games in the same category, players adopt one and the same strategy. In our case, if players consider the game in Table 1 and the game in Table 2 as belonging to one category of games, and if the players consider it sufficiently likely that they are facing a common enemy, the players may cooperate in both games.

Concretely, consider the following variant of the games in Tables 1 and 2. At Stage 0, a random event occurs, such that with probability $(1 - p)$, the players face the defence game against an intelligent Attacker (see Table 2), and with probability p, they face the game with the payoffs specified in Table 1. We here interpret the game in Table 1 in the broadest sense, where the cooperative action of each player contributes $V/2$ to a public good, and where the value of the public good is the sum of the contributions. Then the players' payoff matrix becomes as in Table 3. It is clear now that if p is such that $V - C > pV/2$, or $p < 2(V - C)/V$, meaning that the probability of facing the game in Table 1 is not too large, it continues to be the case that a cooperative equilibrium exists, even though there is no cooperative equilibrium in the game in Table 1. Simply, if a standard prisoner's dilemma is pooled with a game against a common enemy, then the common enemy effect which applies to the latter game may spill over into the prisoner's dilemma.[5]

	Act	Don't
Act	$V - C, V - C$	$pV/2 - C, pV/2$
Don't	$pV/2, pV/2 - C$	0, 0

Table 3. Payoff matrix Game with uncertainty about type of attack faced

In political science, the theory of diversionary foreign policy, also referred to as the scapegoat hypothesis (for an overview, see Levy, 1989) posits that governments may create a foreign enemy, or find a foreign scapegoat, in order to serve domestic purposes. Following Simmel (1955), a rationale that is often provided by political scientists for this government strategy is that it creates internal cohesion. A further variant of the model in Sections 2 and 3 leads to a game-theoretical rationale for the scapegoat hypothesis. Consider exactly the same setting as proposed above, where players do not know whether they are facing the game in Table 1 or 2, but know with which probabilities these games occur. Add now a Government as yet another player, where this government's payoff is equal to the sum of the payoffs of the players (excluding the Attacker). Let the Government find out which game the players play at stage 0. Also, at this stage, let the government be able to send a cheap signal in a commonly known language to the players, informing them about which game is played. Assume further that in line with the experimental literature, players of the game in Table 2 play the cooperative equilibrium with positive probability. Let the players believe that the Government is honest in the signals that it sends. Then it is a best response for the Government to always tell the players that they are playing the game against the common enemy. Given this fact, a separating equilibrium where the Government honestly reveal which game is played does not exist, and the only Nash equilibrium is a pooling equilibrium. If it is again the case that $p < 2(V - C)/V$, so that the probability that the game in Table 1 is played is not too high, in the pooling equilibrium it is possible that players play cooperatively. By not revealing which game is played, the Government then has actively promoted cooperation.

We end by noting that the results obtained in this section depend on the probability that the game in Table 1 is played is not too high. If this probability is high ($p > 2(V - C)/V$), the only equilibrium is one where players do not cooperate, and rather than a common enemy effect spilling over into the game in Table 1, it is the inability of players in the game in Table 1 to cooperate that spills over into the game in Table 2. In this case, with cheap talk, the Government again is not able to credibly inform the players about which game they play, but this now leads to a pooling equilibrium without cooperation. If in this case the Government can give credible evidence of what game the players are playing, it will give the players evidence that they are facing a common enemy, so that cooperation becomes possible at least in the game in Table 2.

[5] For a different way of modelling cooperative spillover effects, see Spagnolo (1999), who in a repeated setting shows how cooperation in one prisoner's dilemma (e.g. a social setting) may spill over into cooperation in another prisoner's dilemma (e.g. a professional setting).

5. A Rationale for Random Attacks

Terrorists seem often to resolve to random attacks. Enders and Sandler (2002) and Arce M. and Sandler (2005) argue that this is done in order to maximise terror, by making everyone feel at risk at every possible time and location. In our model, a rationale for random attacks can be provided without appealing to the psychology of fear.

A direct corollary of our analysis in Sections 2 and 3 is the following result. Consider an Attacker who observes the defensive measures of the Defenders, and is able to attack a single Defender at a cost of *C*. After having observed the Defenders' defensive measures, the Attacker will attack a Defender who has not taken defensive measures, leading to the destruction of the common resource. Because the Attacker makes it costly for a single Defender to deviate from a strategy profile where both Defenders take defensive measures, the Defenders may coordinate on an equilibrium where both take defensive measures. Yet, as follows from the analysis of the game against Nature, if the Attacker would be able to commit herself to mimicking the behaviour of Nature, and in any circumstance attack a random Defender, she would be better off, as the Defenders would never take defensive measures.

Unfortunately, if the Attacker observes the defensive measures of the Defenders, such a commitment is not credible. Whenever the Attacker observes at least one of the Defenders not taking any defensive measures, she has no incentive to randomise her attack. Moreover, even if she would be able to commit herself to a random attack, if she expects only one player to take defensive measures, a random attack may now be too costly for her. In Section 4, we assumed that $2V > K$, ensuring that an Attacker who has not committed herself to random attacks always finds it worth to attack a Defender who did not take defensive measures. However, if $2V > K > V$, the Attacker who has somehow committed herself to a random attack, finds it too costly to attack if she expects that a single player takes defensive measures, as it is then the case that $-½(2V) - K < -2V$.

Still, while the Attacker cannot credibly commit herself to performing a random attack, she could make use of *strategic delegation* (see Vickers (1985) for strategic delegation of profit-maximising firms to non profit-maximising managers). Let us introduce a third type of player into the game in the form of an Agent. The Agent can perform an attack on a single Defender at a cost *K*. Let the Attacker now offer the Agent a contract where the Agent gets a reimbursement *K* from the Attacker only when performing a fully random attack, and no reimbursement when not performing a random attack – where the assumption is that the Attacker can perfectly monitor the actions of the Attacker. The Agent is then just willing to perform the random attack. It the Attacker can now make visible to the Defenders that she has delegated the task of attacking to an Agent which she has given such incentives, the random attack is credible, and the Defenders do not take defensive measures. As an attack is now always performed, and as Defenders do not cooperate, the net benefit of the Attacker is $-K$. Without strategic delegation, the Attacker's payoff equals $-\beta(2V) - (1 - \beta)K$, where β denotes the probability that the Defenders coordinate on the cooperative equilibrium in the game in Table 2. Given our assumption that $2V > K$, as long as the Defenders cooperate with positive probability in the game in Table 2, the Attacker always prefers to delegate to an Agent performing random attacks.[6]

[6] Note that if the Attacker decides on how much to invest in the capacity to monitor the defensive measures of the Defenders *before* deciding on her attacking strategy, she will visibly not invest in this capacity, and she can credibly commit herself to a random attack without strategic delegation. If, however, such monitoring is costly but these costs are only incurred after the players took their defensive measures, a random attack is only credible with strategic delegation.

As a possible illustration, consider again the case where the Defenders are firms in a price cartel, where defensive measures take the form of precautions to keep the cartel secret.[7] A competition authority which can detect which firms are not taking such precautions may encourage cartel formation, by encouraging each firm to take precautions. If, however, the competition authority hires employees who are paid for performing random audits, the individual firm gets an incentive to free-ride on taking precautions, and the cartel breaks down.

6. Robustness

Further insight into the key features of our model leading to the common enemy effect may be obtained by considering some variants of our analysis where the common enemy effect is not at work.

6.1. Lower costs of defensive measures

As we now show, a requirement for our results is that the cost of defensive measures is relatively high. To see why, consider again our two-player game, but assume that the parameters are such that $V > 2C$, so that the costs of taking defensive measures are relatively low. Then in the game against Nature (Section 2), the only Nash equilibrium is one where both Defenders take defensive measures. Starting from a situation where no Defenders take defensive measures, it is a best response for the individual Defender to take defensive measures. This is because some benefits are still obtained for this Defender, as Nature may accidentally attack him rather than the other Defender, and because the cost of defending is relatively low. For the same reason, if one Defender takes defensive measures, then so does the other Defender.

In the game against the Attacker, however, both a Pareto dominant, risk dominated equilibrium exists where both players take defensive measures, and a Pareto dominated, risk dominant equilibrium where neither Defender takes the defensive measure. This is because, when a single Defender deviates from the latter equilibrium and takes defensive measures, even if taking defensive measures is cheap, they are to no avail as it is the other Defender who is attacked. Players may thus get stuck at a Pareto dominated, risk dominant equilibrium. As long as in the game against the Attacker this risk dominant equilibrium is played with positive probability (as is predicted by laboratory experiments, see Devetag and Ortmann (2007)), this means that under these alternative parameters our results are completely reversed (where it should be noted that this is not because of changes in the game against the Attacker, but because of changes in the game against Nature). The presence of a common enemy with positive probability makes the Defenders worse off. Applying the analysis in Section 4, rather than not informing the Defenders that they are facing an attack by Nature, the government should now not inform the Defenders that they are facing a common enemy. Finally, applying the analysis of Section 5, the common enemy does not want to perform random attacks, but targeted ones. The common enemy does not face commitment problems here, and need not strategically delegate.

If we attempt to interpret these results, as pointed out by Kriesberg (1973, p.249), it is conceivable that the presence of a common enemy decreases rather than increases cohesion. A recent overview of literature that argues that a common enemy may decrease cohesion can be

[7] Considering V as the cartel profit under price competition, the analysis is directly applicable.

found in McLauchlin and Pearlman (2012). As we show here, a game-theoretic rationale for such an opposite effect is the following. If the costs of defence are low enough such that each Defender is unilaterally willing to take defensive measures, then substituting an intelligent Attacker for Nature may have adverse effects, as with a common enemy players may get stuck at a non-cooperative equilibrium, where defensive measures by a single player do not make any difference, as the common enemy continues to target weak links anyway.[8, 9]

As an example, consider owners of adjacent houses in a neighbourhood who face the threat of a fire. As the houses are close to one another, any fire quickly spreads and damages all of the houses. If installing a fire alarm is cheap, and if fires occur accidentally, then it may be worth for the individual house owner to install a fire alarm even if none of the other house owners do so. If the fire now happens to start at the house of this individual owner, then at least the fire will be stopped. If, however, fires are caused by an intelligent arsonist who observes which house owners install alarms, the individual house owner will not be willing to install a fire alarm by himself. Even though such a fire alarm is relatively cheap, installing it does not make any sense, as the arsonist will always put fire to unprotected houses anyway.

6.2 Private nature of defended resource

Consider a variant of the analysis in Sections 2 and 3 where defensive measures defend not only a common resource of value V, but also a resource that is private to each Defender, and yields value G. Let each Defender attach weight α to the common resource, and weight $(1 - \alpha)$ to the private resource. As before, whenever a player is attacked who did not take defensive measures, this destroys the value of the common resource for *both* players. However, attack on a player who did not take defensive measures only destroys his own private resource, and not the private resource of the other player. Thus, a player who takes defensive measures always keeps the value of his private resource, and a player who does not take defensive measures obtains the private value of his resource exactly half of the time, no matter what the other player does. The payoff matrices of the games against Nature and against the Attacker are then modified as in Tables 4 and 5.

	Act	Don't
Act	$\alpha V + (1 - \alpha)G - C$, $\alpha V + (1 - \alpha)G - C$	$\alpha V/2 + (1 - \alpha)G - C$, $[\alpha V + (1 - \alpha)G]/2$
Don't	$[\alpha V + (1 - \alpha)G]/2$, $\alpha V/2 + (1 - \alpha)G - C$	$(1 - \alpha)G/2$, $(1 - \alpha)G/2$

Table 4. Payoff matrix with private resource: Game against Nature

[8] For sociological literature that suggests that natural disasters increase cohesion, see the papers cited in Carroll et al. (2005).

[9] In a similar manner, in Hoyer's (2012) model of network formation, the presence of a common enemy leads to increased efficiency in the case of low linking costs (empty, or redundant networks with a common enemy, rather than minimal networks without a common enemy). With high linking costs, the presence of a common enemy may lead to increased efficiency (network formation rather than no network formation).

	Act	Don't
Act	$\alpha V + (1 - \alpha)G - C$, $\alpha V + (1 - \alpha)G - C$	$(1 - \alpha)G - C$, 0
Don't	0, $(1 - \alpha)G - C$	$(1 - \alpha)G/2$, $(1 - \alpha)G/2$

Table 5. Payoff matrix with private resource: Game against Attacker

Assuming that $\alpha V + (1 - \alpha)G > C > [\alpha V + (1 - \alpha)G]/2$, it is clear that it is dominant in the game against Nature not to take defensive measures, whereas in the game against the Attacker, there is both a Nash equilibrium where both take defensive measures and where neither takes defensive measures. This would then seem a generalisation of the analysis in Sections 2 and 3: the common enemy effect continues to apply, in that in the game against the Attacker, coordination on defensive measures becomes possible. However, as soon the Defenders attach positive weight to the private resource ($\alpha < 1$), it need not be the case that the equilibrium where players both take defensive measures is Pareto-efficient. If $\alpha V + (1 - \alpha)(G/2) > C > [\alpha V + (1 - \alpha)G]/2$, the Defenders are better off because of the common enemy effect, as coordinating on defensive measures is Pareto-efficient. Yet, if $\alpha V + (1 - \alpha)G > C > \alpha V + (1 - \alpha)G/2$, in the game against Nature, it is both dominant not to act, and Pareto-efficient not to act; in the game against the Attacker, there is both an equilibrium where Defenders coordinate on defensive measures, and where Defenders do not coordinate on defensive measures, where only the latter is Pareto-efficient. Thus, in this case while a common enemy effect may promote coordination on an equilibrium where all take defensive measures, such an equilibrium is actually to the Defenders' detriment. Thus, assuming that $\alpha V + (1 - \alpha)G > C > [\alpha V + (1 - \alpha)G]/2$, the common enemy effect only has the beneficial effect obtained in Sections 2 and 3 if $C > [\alpha V + (1 - \alpha)G]/2$, or $\alpha > (C - G/2)/(V - G/2)$, i.e. if the weight attached to the common resource is sufficiently large.

The intuition for this result is the following. Let the Defenders' defensive measures also protect private resources owned by the Defenders. If the cost of defensive measures is high, it is efficient for players not to take defensive measures. However, let the Defenders face an intelligent Attacker. Take as a starting point a situation where both Defenders take defensive measures. Then, while defensive measures are costly, it may still not be worth for the individual Defender to deviate and not take defensive measures, as this Defender is then automatically attacked. Thus, when facing an intelligent Attacker, Defenders may get stuck in a Pareto-inefficient equilibrium where they take defensive measures.[10] By contrast, in the game against Nature, unilaterally deviating and not taking defensive measures is a best response, because one is not automatically attacked then. Applying the analysis in Sections 4 and 5, it continues to be the case that the Attacker is better off when able to commit to a random attack; a benevolent government, however, should not promote that Defenders take defensive measures, and therefore should not inform the Defenders that there is a common enemy.

As examples, we may think of the defence expenses of countries with a common enemy, but where attack of one country does not cause much damage to the other country, because the countries are not neighbours and/or do not have a lot of trade relationships. In this case, the presence of a common enemy can create an inefficient equilibrium where all countries overinvest in defence. Similarly, in a neighbourhood, while the probability of burglary in any given house may be too small to justify the costs of burglar alarms, in the presence of an

[10] This result is in line with Arce M. and Sandler (2005), who in their game where Defenders can both defend and attack, find a Pareto-inefficient defensive equilibrium, and a Pareto-efficient offensive outcome, where both Defenders cooperate in destroying the Attacker.

intelligent thief, if many other house owners have burglar alarms, the individual house owner is forced to install one as well, since otherwise he becomes a likely target of burglary. Finally, compare firms that collude into a cartel to firms that do not pay their taxes. Consider as defensive measures the measures taken by these firms not to be caught by the authorities. While the cartel profits have common-resource aspects (if one firm is caught, the cartel is destroyed), the profits from not paying taxes are a private resource. In both cases, the presence of an intelligent auditor who performs targeted audits on firms that do not take precautionary measures may cause a common enemy effect. Yet, in the case of cartels, this common enemy effect may be to the firms' advantage, whereas in the case of tax evasion, it may be to the firms' detriment.

6.3 Defence technology

In our analysis so far, we consider a particular defence technology. Consider a Defender that is not attacked, or is attacked but took defensive measures, as an input of "1" into the defensive production function; consider a Defender that is attacked and did not take defensive measures, as an input of "0" in the defensive production function. Thus, the inputs here are not only given by whether or not the Defenders take defensive measures, but also by whether or not they are attacked: a Defender who does not take defensive measures but is not attacked still delivers an input of 1. Then the production function that we have so far considered takes the form $\min(r, c)*V$, where r is the input of the Row Defender, and c is the input of the Column Defender. This shows that we have so far treated the extreme case where the Defenders' inputs are perfect complements, in the sense that when attacked by an Attacker, only the lowest input counts. More generally, following Cornes (1993), we may consider a CES production function $[0.5(r^{\pi} + c^{\pi})]^{1/\pi}*V$, where π is proportional to the elasticity of substitution between the Defenders' inputs. As π approaches minus infinity, we approach our production function $\min(r, c)*V$ ("weakest link" production function). As π approaches one, we approach a production function $[(r + c)/2]*V$, where the level of the common resource maintained equals the average of the Defenders' inputs ("average" production function). Finally, as π approaches plus infinity, we approach the production function $\max(r, c)*V$, where attack can only destroy the common resource if neither Defender takes defensive measures, and if at least one of them is attacked ("best shot" production function). Since for negative π, the CES production function is undefined for zero inputs, in this case we take instead of a zero input an input level arbitrarily close to zero.

For the CES production function, the game against Nature can generally be represented by Table 6, where $x(\pi)$ is a fraction of V, with $½ \leq x(\pi) \leq 1$. $x(\pi) = ½$ corresponds to the production function $\min(r, c)*V$, $x(\pi) = ¾$ corresponds to $[(r + c)/2]*V$, and $x(\pi) = 1$ corresponds to $\max(r, c)*V$. In the same way, the game against Nature can generally be represented by Table 7, where $y(\pi)$ is a fraction of V, with $0 \leq y(\pi) \leq 1$. $y(\pi) = 0$ corresponds to the production function $\min(r, c)*V$, $y(\pi) = ½$ corresponds to $[(r + c)/2]*V$, and $y(\pi) = 1$ corresponds to $\max(r, c)*V$. It is the case that $x(\pi) \geq y(\pi)$, where the difference $[x(\pi) – y(\pi)]$ decreases in π. $[x(\pi) – y(\pi)] = ½$ corresponds to the production function $\min(r, c)*V$, $[x(\pi) – y(\pi)]= ¼$ corresponds to $[(r + c)/2]*V$, and $[x(\pi) – y(\pi)] = 0$ corresponds to $\max(r, c)*V$, in which case the game against Nature and the game against the Attacker are identical.

	Act	Don't
Act	$V - C$, $V - C$	$x(\pi)V - C$, $x(\pi)V$
Don't	$x(\pi)V$, $x(\pi)V - C$	0, 0

Table 6. Payoff matrix with substitutability: game against Nature

	Act	Don't
Act	$V - C$, $V - C$	$y(\pi)V - C$, $y(\pi)V$
Don't	$y(\pi)V$, $y(\pi)V - C$	0, 0

Table 7. Payoff matrix with substitutability: game against Attacker

The set of equilibria in the games in Tables 6 and 7 is given by two factors, namely *first* whether or not the individual player prefers to act when the other player acts (i.e. whether or not $V - C > x(\pi)V$ or $[1 - x(\pi)]V > C$, and in the same way whether or not $[1 - y(\pi)]V > C$); and whether or not the individual player prefers to act when the other player does *not* act (i.e. whether or not $x(\pi)V - C > 0$ or $x(\pi)V > C$, and in the same way whether or not $x(\pi)V - C > 0$ or $x(\pi)V > C$). It follows that the set of equilibria is given in the first place by the relation between the four values $x(\pi)$, $y(\pi)$, $[1 - x(\pi)]$ and $[1 - y(\pi)]$ (which depends on π, and therefore on the degree of substitutability between the Defenders' inputs). The second factor determining the set of equilibria (holding V fixed) is the relation between C and these four values.

It is now easily checked that, as a function of the level of π, the following cases are possible. Define a level π^*, with $0 < \pi^* < 1$. Then in Case 1, we have $\pi \leq \pi^*$, i.e. a low degree of substitutability between the Defenders' inputs. In this case, we have that $[1 - y(\pi)] > x(\pi) > [1 - x(\pi)] > y(\pi)$. This means that the largest marginal benefit of acting across the two games is when facing an Attacker, and when the other player already acts. This is because not acting means that one is automatically attacked, whereas one is crucial in the defence of the common resource. For the same reason, the smallest marginal benefit of acting is when facing an Attacker, and when the other player does not act: the other player is targeted anyway. In between lie the marginal benefits of acting when facing Nature. The highest among these benefits is the one of acting when the other player is not acting. Half of the time, one is then attacked and this leads to full defence of the resource, half of the time the other player is attacked, but if defensive measures are not perfect complements, this does still lead to some degree of defence of the common resource. The marginal benefit of acting along with an acting player in the game against Nature is smaller, because half of the time, the other player is targeted, so that one's contribution does not make any difference.

Given this order of the marginal benefits of acting, as one starts from a high cost C, and gradually lower this cost, the first equilibrium encountered where at least one player acts, is joint defence against the Attacker. It follows that for high C, the common enemy effect as laid out in Sections 2 and 3 applies. For intermediate C, in the game against Nature the only equilibria are ones where a single Defender takes defensive measures (asymmetric equilibria), whereas in the game against the Attacker there continue to be two symmetric equilibria. In this case, which becomes possible as soon as the defensive measures are not perfect complements, whether or not the common enemy effect applies depends on which equilibrium is played in the game against the Attacker. Finally, for low C, in the game against Nature the only equilibrium is one where both players take defensive measures, whereas in the game against the Attacker both symmetric equilibria continue to exist; this is a generalisation of the case treated in Section 6.1, where the common enemy effect applies, but leads to inefficient cooperation on defensive measures.

In Case 2, we have $\pi^* < \pi \leq 1$, and there is an intermediate level of substitutability between the Defenders' inputs. It is now the case that $x(\pi) > [1 - y(\pi)] > y(\pi) > [1 - x(\pi)]$. The difference with Case 1 is that the highest marginal benefit of acting across the two games is now, in the game against Nature, the benefit of acting rather than not acting when the other player does not act. Given that the degree of substitutability of the Defenders' inputs is now higher, in the game against the Attacker the consequences of deviating from joint defence are less severe, whereas the benefits of acting alone in the game against Nature have increased.

Given this modified order of the marginal benefits of acting, as one starts from high action costs and gradually lowers these costs, the first equilibrium encountered where at least one of the players acts is an asymmetric equilibrium in the game against Nature. Thus, for high C, in the game against Nature the only equilibria are asymmetric ones where one Defender takes defensive measures; in the game against the Attacker, the only equilibrium is a symmetric one without defensive measures. In this case, the presence of an Attacker eliminates rather than stimulates defensive measures, and the common enemy effect does not apply. For intermediate C, asymmetric equilibria continue to exist in the game against Nature, while two symmetric equilibria, namely with and without cooperation, exist in the game against the Attacker. The validity of the common enemy effect here again depends on how often Defenders coordinate on the Pareto dominant equilibrium. For low C, asymmetric equilibria still exist in the game against Nature, but the only equilibrium in the game against the Attacker is one where both players take defensive measures, so that the common enemy effect applies here. We conclude that if we increase the degree of substitutability between the Defenders' inputs, the common enemy effect only applies for lower costs of taking defensive measures.

Finally, in Case 3 we have $\pi > 1$,[11] for which it is true that $x(\pi) > y(\pi) > [1 - y(\pi)] > [1 - x(\pi)]$. This case, in which the degree of substitutability between the Defenders' inputs is high, differs from Case 2 in that in the game against the Attacker, the marginal benefit of acting alone is larger than the marginal benefit of acting together with the other player. The reason is simply because action by one player almost completely defends the resource. In this case, for high C, there are two asymmetric equilibria in the game against Nature, whereas in the game against the Attacker, no defensive measures are taken in equilibrium. Thus, the common enemy effect does not apply here. For intermediate C, all equilibria are asymmetric in both games. For low C, finally, there are only asymmetric equilibria in the game against Nature, whereas in the game against the Attacker, joint defence always takes place. This again confirms that if the degree of substitutability between the defensive measures is relatively high, the common enemy effect only applies for lower action costs.

As an illustration, consider missile defence systems[12] set up by two countries, where the setting up of such a defence system in one country also protects to a large extent the other country. If such a system is set up against meteorites or crashing satellites, the only equilibrium may be where only one country sets up the system, thus protecting both countries. This is because the benefit to one country of installing such a system is large: there is full protection if the attack is in the domestic country and still some protection if the attack takes place in the foreign country. When the attack instead takes the form of a missile attack by a common enemy, the benefit of setting up a missile defence system by oneself is smaller, because the country that does not set up such a system will always be attacked. Therefore, if setting up missile defence systems is expensive, it will be the case that a single country sets

[11] Cornes (1993) refers to such a production function as a "better shot" production function, which is a generalisation of the "best shot" production function $\max(r, c)*V$.

[12] A similar example is given by Cornes (1993) for the "best shot" production function.

up such a system in the game against Nature, whereas none do in the game against the common enemy. If setting up the missile defence system is cheap, however, an equilibrium exists in the game against the common enemy where both players set up the system, so that the common enemy effect applies again.

7. Discussion

As shown in this paper, the threat of an intelligent common enemy performing targeted attacks, rather than an impersonal enemy in the form of nature performing random attacks, may encourage cooperation between players. By attacking weakest links, the common enemy inadvertently punishes the formation of weakest links, so that no player wants to be the weakest link. Given this fact, a benevolent government that seeks to promote cooperation has an incentive to make players believe that there is a common enemy, even if there is none. The common enemy on the other hand has an incentive to commit herself to replicating the behaviour of an impersonal common threat, and to perform random attacks.

The purpose of our model is not to diminish the literature that provides a psychological explanation for the common enemy effect. Indeed, in a series of laboratory experiments, Gary Bornstein and co-authors find common enemy effects for which a rational explanation is not ready at hand. In particular, Bornstein and Ben-Yossef (1994) show that participants are twice more likely to cooperate in a prisoner's dilemma involving group competition than in an equivalent prisoner's dilemma not involving such competition. Contrary to what is the case in Bornstein et al. (2002), where it is shown that participants are more likely to coordinate in stag hunts if they face group competition, cooperation in the prisoner's dilemma cannot be explained by means of focal points.

Rather, our purpose has been to show that there also exists a simple rational basis for the common enemy effect. Given the focus of a rational model on incentives, the attractive feature of our model is that it predicts the presence of a common enemy effect in particular contexts. The context which is the main focus of our paper is one where 1) defence concerns to a sufficient degree a common resource; 2) the complementarity between the players' defensive measures is relatively high; 3) the costs of defence are relatively large. The fact that defence concerns a common resource ensures that the common enemy effect, when it applies, also leads to efficient play; if defence concerns to a large degree private resources, defence in the presence of a common enemy may merely impose a negative externality on other players, as these are then more likely to be targeted. The fact that there is complementarity between the players' defensive measures ensures that, in the presence of a common enemy, it is expensive to free-ride, so that a cooperative equilibrium exists. The fact that the costs of defence are large ensures that there is no cooperative equilibrium in the game against Nature. If defence costs are low, a cooperative equilibrium exists in the game against Nature, whereas players of the game against the common enemy may get locked into inaction: unilateral defensive measures are to no avail, as it is the players who do not take such measures who will be targeted. Thus, our analysis reveals that there are also contexts where the presence of a common enemy may discourage rather than encourage cooperation.

In an extension of our model, we also checked whether the common enemy effect continues to apply if the players' defensive measures have a higher degree of substitutability. In this case, in the game against Nature, the individual player has a large incentive to take unilateral defensive measures: the blind attack by Nature may be targeted at this player, leading to full protection; even if the attack is not targeted at the player, because of substitutability the common resource still is somewhat protected. The incentive to take unilateral defensive measures in the game against the common enemy is smaller, as the player who takes defensive measures is never targeted. Therefore, for relatively high defence costs, it may be that in the game against Nature, only one player takes defensive measures, whereas in the game against the common enemy, no player takes defensive measures. In this case, the common enemy effect does not apply. The common enemy effect now only applies for lower defence costs, when a cooperative equilibrium exists in the game against the common enemy. This extension thus shows that the common enemy effect is fairly general, in that it may also apply if defensive measures are substitutes.

We conclude that game theory offers a rational explanation of the common enemy effect in specific and identifiable contexts, but in other contexts predicts what could be termed an anti common enemy effect, where the presence of a common enemy discourages rather than encourages cooperation. We hope that the concrete predictions of our model lead to further empirical research on the common enemy effect, to identify the importance of this effect, and to identify the extent to which it is due to rational and/or psychological factors.

References

Antal, T., Krapivsky, P.L. and Redner, S. (2006) Social balance on networks: the dynamics of friendship and enmity. *Physica D* 224, 130-136.

Arce M., D.G and Sandler, T. (2005) Counterterrorism: a game-theoretic analysis. *Journal of Conflict Resolution* 49, 183-200.

Arenas, A., Camacho, J., Cuesta, J.A. and Requejo, R.J. (2011) The joker effect: cooperation driven by destructive agents. *Journal of Theoretical Biology* 279, 113-119.

Arquilla, J. and Ronfeldt, D. (2001) *Networks and Netwars: The Future of Terror, Crime and Militancy.* Santa Monica, CA: Rand.

Bornstein, G. and Ben-Yossef, M. (1994) Cooperation in intergroup and single-group social dilemmas. *Journal of Experimental Social Psychology* 30, 52-67.

Bornstein, G., Gneezy, U. and Nagel, R. (2002) The effect of intergroup competition on group coordination: an experimental study. *Games and Economic Behavior* 41, 1-25.

Boyd, R. and Richerson, P.J. (1990) Group selection among alternative evolutionary stable strategies. *Journal of Theoretical Biology* 145, 331-342.

Carroll, M.S., Cohn, P.J., Seesholtz, D.N. and Higgins, L.L. (2005) Fire as a galvanizing and fragmenting influence on communities: the case of the Rodeo-Chediski fire. *Society and Natural Resources* 18, 301-320.

Cornes, R. (1993) Dyke maintenance and the other stories: some neglected types of public goods. *Quarterly Journal of Economics* 108, 259-271.

Coser, L. A. (1956) *The Functions of Social Conflict*. Glencoe, IL: Free Press.

Devetag, G. and Ortmann, A. (2007) When and why? A critical survey on coordination failure in the laboratory. *Experimental Economics* 10, 331-344.

Harsanyi, J.C. and Selten, R. (1988) *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge.

Heider, F. (1946) Attitudes and cognitive organization. *Journal of Psychology* 21, 107-112.

Hirshleifer, J. (1971) The private and social value of information and the reward to inventive activity. *American Economic Review* 61, 561–574.

Hirshleifer, J. (1983) From weakest-link to best-shot: the voluntary provision of public goods. *Public choice* 41, 371-386.

Hoyer, B. (2012) Network disruption and the common enemy effect. TKI Discussion Paper 12-06, Utrecht University School of Economics.

Jackson, M.O. and Wolinsky, A. (1995) A strategic model of social and economic networks. *Journal of Economic Theory* 71, 44-74.

Kovenock, D. and Roberson, B. (2012) Coalitional Colonel Blotto games with application to the economics of alliances. *Journal of Public Economic Theory* 14, 653-676.
Kriesberg, L. (1973) *The Sociology of Social Conflicts*. Englewood Cliffs, N.J.: Prentice-Hall.
Levy, Jack S. (1989). The diversionary theory of war: a critique. In M. M. Midlarsky (ed.) *Handbook of War Studies*, London: Unwin Hyman.
McLauchlin, T. and Pearlman, W. (2012) Out-group conflict, in group unity? Exploring the effect of repression on intramovement cooperation. *Journal of Conflict Resolution* 56, 41-66.
Mengel, F. (forthcoming) Learning across games. *Games and Economic Behavior*.
Mesterton-Gibbons, M. (2001) *An Introduction to Game-Theoretic Modelling* (2^{nd} Ed.) Student Mathematical Library, Volume 11.
Mesterton-Gibbons, M. and Dugatkin, L.A. (1992) Cooperation among unrelated individuals: evolutionary factors. *Quarterly Review of Biology* 67, 267-281.
Reichmann, T. and Weimann, J. (2008) Competition as a coordination device: experimental evidence for a minimum effort coordination game. *European Journal of Political Economy* 24, 437-454.
Sandler, T. (2006) Hishleifer's social composition function in defense economics. *Defence and Peace Economics* 17, 645-655.
Simmel, Georg. [1908] 1955. Conflict. Glencoe, IL: Free Press.
Sirin, C.V. (2011) Is it cohesion or diversion? Domestic instability and the use of force in international crises. *International Political Science Review* 32, 303-321.
Spagnolo, G. (1999) Social relations and cooperation in organization. *Journal of Economic Behavior and Organization* 38, 1-25.
Vickers, J. (1985) Delegation and the theory of the firm. *Economic Journal* 95, 138-147.
Waltz, K.N. (1979) *Theory of International Politics*. Reading, Addison-Wesley.