doi:10.1111/j.1558-5646.2010.01176.x

A STRATEGY TO INCREASE COOPERATION IN THE VOLUNTEER'S DILEMMA: REDUCING VIGILANCE IMPROVES ALARM CALLS

Marco Archetti[1,2,3]

[1]*Department of Organismic and Evolutionary Biology, Harvard University, 26 Oxford Street, Cambridge, Massachusetts 02138*

[2]*Faculty of Business and Economics, University of Basel, Peter Merian-Weg 6, 4002 Basel, Switzerland*

[3]*E-mail: archetti@fas.harvard.edu*

Received July 27, 2010
Accepted October 14, 2010

One of the most common examples of cooperation in animal societies is giving the alarm in the presence of a predator. A reduction in individual vigilance against predators when group size increases (the "group size effect") is one of the most frequently reported relationships in the study of animal behavior, and is thought to be due to relaxed selection, either because more individuals can detect the predator more easily (the "many eyes" effect) or because the risk of predator attack is diluted on more individuals (the "selfish herd" effect). I show that these hypotheses are not theoretically grounded: because everybody relies on someone else to raise the alarm, the probability that at least one raises the alarm declines with group size; therefore increasing group size does not lead to relaxed selection. Game theory shows, instead, that increasing the risk that the predator is not reported (by reducing vigilance) induces everybody to give the alarm more often. The group size effect, therefore, can be due to strategic behavior to improve the production of a public good. This shows how a selfish behavior can lead to a benefit for the group, and suggests a way to solve social dilemmas in the absence of relatedness and repeated interactions.

KEY WORDS: **Alarm calls, brinkmanship, cooperation, game theory, group size effect, public goods, social dilemma, strategy, vigilance.**

THE GROUP SIZE EFFECT: MANY EYES AND THE SELFISH HERD

Many animals living in groups rely on alarm calls as a defense against predators. Vigilance against predators is one of the most common and simple examples of cooperation in animal societies (Clutton-Brock, 1999; Searcy and Nowicki 2005), commonly mentioned when describing the benefits of living in a group: by taking advantage of the vigilance of many individuals a group can detect a predator more easily; moreover the risk of an attack is diluted among all group members.

A reduction in individual vigilance against predators when group size increases (the "group size effect") is one of the most frequently reported relationships in the study of animal behavior (Roberts 1996; Beauchamp 2003, 2008). The adaptive explanation of the group size effect remains unclear. The two main hypotheses are based on the assumption that selection is relaxed when group size increases: the "group vigilance" or "many eyes" hypothesis states that, by taking advantage of the vigilance of many group members, individuals are allowed a reduction in their own vigilance to obtain the same level of security (Pulliam 1973); the "dilution of risk" or "selfish herd" hypothesis, instead, states that increasing group size allows a reduction in vigilance because the risk of predation for each group member is diluted (Hamilton 1971); there seems to be no consensus yet on the relative importance of these two hypotheses (Roberts 1996; Beauchamp 2003, 2008).

The rationale of this article is that these two hypotheses are theoretically ungrounded and that there is a more likely explanation for the group size effect. First, I show that living in a group is not necessarily an advantage against predators: in fact, although the probability that someone detects the predator increases, the probability that someone raises the alarm *decreases* with group size; therefore selection is not necessarily relaxed when group size increases. Second, I suggest an alternative explanation for the group size effect: reducing individual vigilance can be a deliberate strategy to increase the probability that someone else gives the alarm by inducing a risk.

VIGILANCE GAMES

Game theory has been used to understand the group size effect by Pulliam et al. (1982) and their model has been modified by Parker and Hammerstein (1985) and extended by McNamara and Houston (1992). In these models an individual in a group can be vigilant or not against a predator attack; if nobody sees the predator the attack is always successful and the predator kills one of the group members; if one is vigilant he always sees the predator and always gives the alarm, which reduces the probability of a successful attack. Vigilance is costly, and the game between group members consists in deciding the optimal amount of time spent being vigilant. The result is that individual vigilance can decrease when group size increases, because of the "many eyes" effect and "selfish herd" effect.

Here I consider a different type of game. As in previous models an individual in the group can be vigilant or not; if one is vigilant, however, he can choose to raise the alarm or not (the alarm is not automatic); if at least one individual raises the alarm this reduces the probability of a successful attack. Vigilance itself is not costly, but raising the alarm is; The game therefore has two levels of decision: first, decide the fraction of time devoted to vigilance; second, provided one is vigilant, decide whether to raise the alarm. Why this is important can be understood looking at strategic behavior in a simple public goods game, the volunteer's dilemma.

THE VOLUNTEER'S DILEMMA

The volunteer's dilemma (Diekmann 1985; Archetti 2009a, 2009b) is an N-person game in which a public good is produced if and only if at least one player volunteers to pay a cost. The basic model is the following: N individuals witness a crime; the criminal could be arrested if at least one of the witnesses called the police, which costs c; if nobody calls the police the criminal remains at large, at a cost a ($>c$) for everybody. Clearly volunteering produces a public good, and a volunteer benefits from his action if nobody else volunteers, but the cost of volunteering is wasted if someone else volunteers; this is the dilemma. A similar problem arises in the case of alarm calls against predators: the alarm is a public good; one individual is enough to give the alarm, and giving the alarm may have nonnegligible costs, for example because it increases the risk of being attacked by the predator (Searcy and Nowicki 2005). What is the best strategy for an individual?

Let us assume, for the moment, that all individuals always detect the predator with certainty. An individual that detects a predator approaching can give the alarm (*Volunteer*) or not (*Ignore*). The game has N asymmetric pure-strategy equilibria in which only one player gives the alarm, but they require coordination and can only work if the players decide in advance who is going to volunteer. The game has also a symmetric mixed-strategy equilibrium, which does not require coordination, in which the two pure strategies (*Ignore* and *Report*) are played with a probability. The payoffs of the strategy *Volunteer* is $1 - c$ because the volunteer pays the cost and the public good is produced; the payoff of the strategy *Ignore* is $\gamma^{N-1}(1 - a) + (1 - \gamma^{N-1})(1)$, where γ is the probability of playing *Ignore*, because with probability γ^{N-1} nobody volunteers (therefore the public good is not produced at a cost a) and with probability $1 - \gamma^{N-1}$ at least someone volunteers (and the public good is exploited at no cost); the mixed ESS is $\gamma_{eq} = (c/a)^{1/(N-1)}$ (Archetti 2009a,b).

THE BYSTANDER EFFECT

The volunteer's dilemma has interesting and counterintuitive consequences. At equilibrium the probability γ_{eq} of ignoring the predator increases with N, which is intuitive because when there are more players each relies more on somebody else volunteering. What is perhaps less intuitive is that the probability that nobody volunteers, and that therefore the alarm is not given, also increases with N. Because $\gamma_{eq} = (c/a)^{1/(N-1)}$ then $\gamma_{eq}^{N-1} = c/a$, and $\gamma_{eq}^{N} = \gamma_{eq}c/a$, which is increasing in N, because γ_{eq} is increasing in N. Therefore the more individuals see the predator, the *less* likely it is that at least one will actually give the alarm. This is known in the social sciences as the "bystander effect"; it is the theoretical explanation for the commonly observed fact that the probability that someone reports a crime or an accident decreases with group size (Diekmann 1985). Increasing group size therefore leads to a reduction of the probability that the public good (the alarm) is produced.

Because so far I have assumed that everybody detects the predator with certainty, changing group size does not have any effect on the *probability* to detect the predator; moreover I have assumed that the cost of predation is not diluted among group members. To assess the "many eyes" hypothesis we must generalize the game to cases in which the predator is detected with a probability <1; to assess the "selfish herd" hypothesis, we must allow a dilution of risk. Increasing group size will be beneficial (because of relaxed selection) only if the benefit due to the increased probability to detect the predator (for the may eyes effect) or the dilution of risk (for the selfish herd effect) overcomes the cost due to the reduced probability that the alarm is raised (the bystander effect).

Model

Let us assume that each individual has a probability $\omega \leq 1$ to detect the predator; ω can be thought of as a measure of individual vigilance, and the game can be thought of as a distribution of subgames in which M of the N individuals of the group detect the predator; each of these subgames occurs with probability

$$p_M = \binom{N}{M} \omega^M (1-\omega)^{N-M}.$$

In a subgame with $M > 0$, we can write an individual's payoff as

$$\zeta_{(M)} = \gamma_{(M)}^{M}(1-a) + \left(1-\gamma_{(M)}^{M}\right)(1),$$

when he is nonvigilant, irrespective of his own strategy, because in this case the payoff is $(1-a)$ if all the M witnesses play *Ignore* (which happens with probability $\gamma_{(M)}^{M}$), and 1 if at least one of the witnesses plays *Volunteer* (which happens with probability $1-\gamma_{(M)}^{M}$); and

$$\xi_{(M)} = \gamma_{(M)}\left[\gamma_{(M)}^{M-1}(1-a) + \left(1-\gamma_{(M)}^{M-1}\right)(1)\right] + (1-\gamma_{(M)})[1-c]$$

when he is vigilant, because in this case the payoff is $(1-c)$ if he plays *Volunteer* (which happens with probability $1-\gamma_{(M)}$), whereas if he plays *Ignore* (which happens with probability $\gamma_{(M)}$) it depends on the strategies of the other $M-1$ witnesses: if they all play *Ignore* (which happens with probability ${\gamma_{(M)}}^{M-1}$), the payoff is $(1-a)$; if at least one plays *Volunteer* (which happens with probability $1-\gamma_{(M)}^{M-1}$) the payoff is 1.

In other words, if $M > 0$ we can write the payoffs of the two pure strategies *Volunteer* and *Ignore*, respectively, as

$$\begin{aligned} W_{V(M)} &= \omega_o[1-c] \\ &\quad + (1-\omega_o)\left[\gamma_{(M)}^{M}(1-a) + \left(1-\gamma_{(M)}^{M}\right)(1)\right] \\ W_{I(M)} &= \omega_o\left[\gamma_{(M)}^{M-1}(1-a) + \left(1-\gamma_{(M)}^{M-1}\right)(1)\right] \\ &\quad + (1-\omega_o)\left[\gamma_{(M)}^{M}(1-a) + \left(1-\gamma_{(M)}^{M}\right)(1)\right], \end{aligned}$$

where ω_o is the vigilance level of the focal individual.

The probability of playing *Ignore* at equilibrium in a game with $M > 1$ witnesses can be found by equating $W_{V(M)}$ and $W_{I(M)}$, which yields the familiar result

$$\gamma_{eq(M)} = (c/a)^{1/(M-1)}.$$

If $M = 1$ there is no mixed equilibrium and $\gamma_{eq(1)} = 0$; $W_{V(1)} = W_{I(1)} = 1-c$. If $M = 0$ clearly $W_{V(0)} = W_{I(0)} = 1-a$. Fitness therefore is

$$W = p_0[1-a] + \sum_{M=1}^{N} [\omega_o q_M \xi_{(M)} + (1-\omega_o) g_M \zeta_{(M)}],$$

where

$$g_M = \binom{N-1}{M} \omega^M (1-\omega)^{N-M-1}$$

is the probability that exactly M among the $N-1$ other players with vigilance ω observe the predator, and

$$q_M = \binom{N-1}{M-1} \omega^{M-1} (1-\omega)^{N-M}$$

is the probability that exactly $M-1$ among $N-1$ other players with vigilance ω observe the predator. The probability that the public good is produced is

$$G = p_0[0] + p_1[1] + \sum_{M=2}^{N} p_M\left[1-\gamma_{(M)}^{M}\right]$$

We still have not defined what $\gamma_{(M)}$ is, however, as this depends on the assumptions of the model. We have three different cases: (1) the number of actual witnesses (M) is common knowledge; (2) the number of actual witnesses is unknown but the average vigilance level (ω) is common knowledge; (3) neither the number of witnesses nor the average vigilance level is known.

If M (the number of witnesses in the current subgame) is common knowledge, a player can simply choose a probability of playing *Ignore* for each M, that is $\gamma_{(M)} = \gamma_{eq(M)}$. If ω is common knowledge but M is unknown, players must choose a single value of $\gamma_{(M)}$ for all subgames: $\gamma_{(M)} = \gamma^{\bullet}_{eq}$; this value $\gamma^{\bullet}_{eq}$ can be found numerically by setting $\gamma_{(M)} = \gamma^{\bullet}_{eq}$ in $W_{V(M)}$ and $W_{I(M)}$ and equating the payoffs of *Ignore* and *Volunteer* over all the possible subgames (M values):

$$p_0[1-a] + \sum_{M=1}^{N} p_M W_{V(M)} = p_0[1-a] + \sum_{M=1}^{N} p_M W_{I(M)},$$

An approximation of this method, which can be useful for computational purposes when N is large, is to take the weighted (proportional to p_M) average of the optimal values calculated for each M, that is

$$\gamma_A = p_1[0] + \sum_{M=2}^{N} p_M\ \gamma_{eq(M)}.$$

If neither ω nor M are known, clearly there is no calculable optimal strategy; one could assume, however, that individuals have a belief on the distribution of ω values and calculate their optimal response given that belief. I do not discuss this model (the results are not very different from the other models, unless there is very high uncertainty on the value of ω*;* for large N computation time increases rapidly with N).

Finally, the level of vigilance that maximizes fitness can be found by introducing mutants with vigilance $\omega^* = \omega - \varepsilon$ ($\varepsilon > 0$; I use $\varepsilon = 0.01$); when a mutant is introduced ($\omega_o = \omega^*$), a new stable value of volunteering γ_{eq} is calculated, and the public good (and therefore the payoffs for all individuals) updated, to check whether the mutation invades and replaces the resident vigilance level ω.

Results

The results for the case in which M is common knowledge are shown in Figure 1 and Table 1. The case in which the number of witnesses (M) is unknown (and only vigilance is common knowledge) leads to similar results, but fitness and the probability to produce the public good are slightly higher (Fig. 2, Table 2). The results of the approximate method are also similar. The rest of the article is based on the case in which M is common knowledge.

PERFECT VIGILANCE

Consider, first, the dilution of risk in the case of perfect vigilance ($\omega = 1$). The bystander effect is not affected by whether the risk of predation is constant or diluted among all group members ($a = 1/N$) because the dilution of predation reduces individual risk, but just because this risk is lower it also induces everybody to volunteer (to raise the alarm) less often: the two effects (the bystander effect and the dilution of risk) have opposing effects that exactly cancel each other; the strength of selection, therefore, does not change with group size. It is easy to see why if we look at the mixed equilibrium: because at equilibrium the payoffs of the two strategies *Volunteer* and *Ignore* must be equal, and because the fitness of *Ignore* is constant ($1 - c$), the fitness of the mixed strategy must also remain constant (and equal to $1 - c$), as long as the mixed equilibrium exists (the mixed equilibrium disappears if $a < c$, that is if the cost of predation is diluted, for example $a = 1/N$, and the size of the group increases to more than a/c; in this case there is no volunteering at all and no alarm calls).

Result 1: With perfect vigilance, as group size increases, the cost of the bystander effect (a higher probability that the alarm is *not* produced) exactly cancels out the benefit of the selfish herd effect (relaxed selection due to the dilution of risk). Increasing

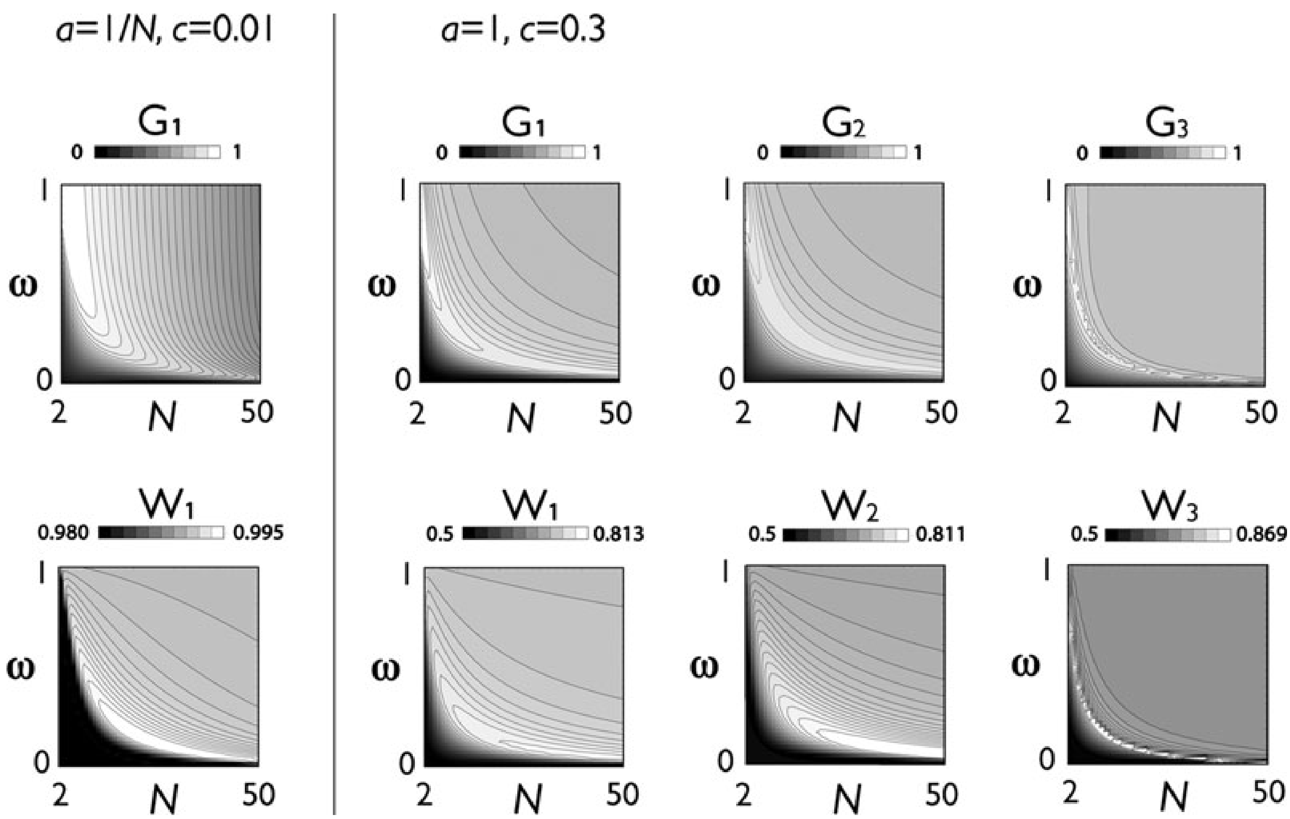

Figure 1. **The bystander effect with reduced vigilance. The probability that at least one raises the alarm (G_i) and fitness (W_i) at the mixed equilibrium as a function of group size (N) and vigilance (ω) for different values of the cost of raising the alarm (c) and of the cost paid if nobody raises the alarm (a), when the number of witnesses is common knowledge ($i = 1$); using the approximate method ($i = 2$); when only individual vigilance is common knowledge ($i = 3$).**

Table 1. Optimal vigilance, fitness, and public goods when the number of witnesses is common knowledge. For each value of *c* (the cost of volunteering) and *N* (group size) the table shows fitness (W_1) and the probability to produce the public good (G_1) with full vigilance ω=1; the stable level of vigilance (ω_{eq}), fitness (W_{eq}), and the probability to produce the public good (G_{eq}) with vigilance ω_{eq}. The number of players that are actually vigilant (*M*) is common knowledge; *a*=1.

	N					
	5	10	20	30	50	100
c=0.9						
W_1	0.1	0.1	0.1	0.1	0.1	0.1
W_{eq}	0.378	0.400	0.410	0.413	0.411	0.412
G_1	0.123	0.110	0.105	0.103	0.102	0.101
G_{eq}	0.462	0.440	0.429	0.426	0.418	0.416
ω_{eq}	0.24	0.12	0.06	0.04	0.02	0.01
c=0.5						
W_1	0.5	0.5	0.5	0.5	0.5	0.5
W_{eq}	0.632	0.655	0.666	0.671	0.671	0.676
G_1	0.580	0.537	0.518	0.512	0.507	0.503
G_{eq}	0.719	0.698	0.688	0.685	0.679	0.680
ω_{eq}	0.39	0.22	0.12	0.08	0.04	0.02
c=0.3						
W_1	0.7	0.7	0.7	0.7	0.7	0.7
W_{eq}	0.775	0.796	0.806	0.810	0.813	0.815
G_1	0.778	0.738	0.718	0.712	0.707	0.704
G_{eq}	0.851	0.833	0.825	0.822	0.820	0.819
ω_{eq}	0.49	0.28	0.15	0.10	0.06	0.03
c=0.01						
W_1	0.99	0.99	0.99	0.99	0.99	0.99
W_{eq}	0.991	0.993	0.994	0.994	0.994	0.994
G_1	0.997	0.994	0.992	0.991	0.991	0.990
G_{eq}	0.998	0.996	0.996	0.995	0.995	0.995
ω_{eq}	0.82	0.55	0.33	0.22	0.14	0.07

group size, therefore, does not lead to relaxed selection against predation with perfect vigilance.

REDUCED VIGILANCE

Even with reduced individual vigilance ($\omega < 1$) the probability that someone will give the alarm decreases with group size, unless *N* is very small, or ω is very low. This is true even when the cost of predator attack is diluted among group members (Fig. 1). Again, as in the case of perfect vigilance ($\omega = 1$), although it is true that the cost of predation is diluted among many individuals, increasing the number of witnesses also leads to the bystander effect: the alarm will be given less often. Now, however, the two effects do not cancel each other. With imperfect vigilance, the combined effects of the dilution of risk and of the bystander effect lead to stronger predation risk overall, unless *N* is very small or unless $\omega \approx 0$; only in these cases is selection relaxed, because the

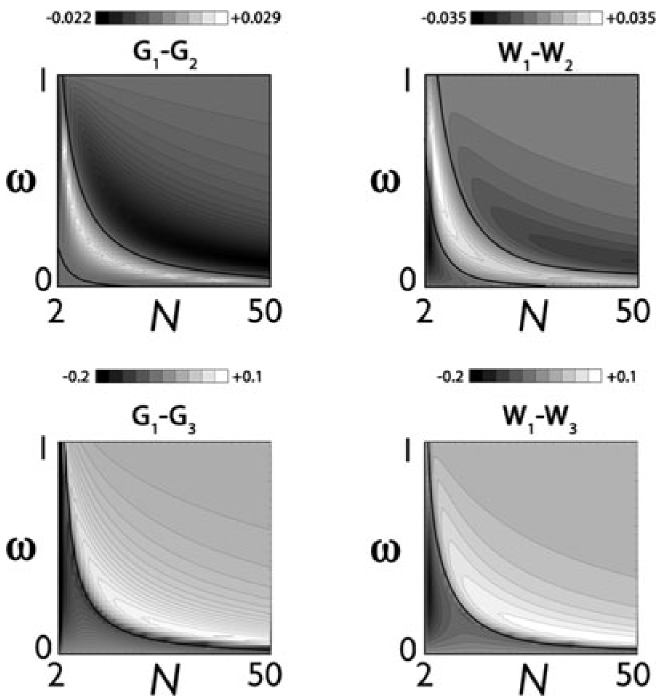

Figure 2. Difference between the three models. The difference in public good (G_i) or fitness (W_i) as a function of group size (*N*) and vigilance (ω) between the model in which the number of witnesses is common knowledge ($i = 1$) and the approximate method ($i = 2$); or the model in which only average vigilance is common knowledge ($i = 3$); $c = 0.3$ (the cost of volunteering); $a = 1$ (the benefit of the public good); the bold lines show where the difference is zero.

Table 2. Optimal vigilance, fitness, and public goods; differences between the three models. For each value of *N* (group size) the table shows the stable vigilance level, fitness, and the probability to produce the public good with that vigilance level when the number of players that are actually vigilant (*M*) is common knowledge (ω_{eq}, W_{eq}, G_{eq}; see Table 1); when vigilance is common knowledge but the number of players that are actually vigilant (*M*) is unknown ($\omega^\bullet$, $W^\bullet$, $G^\bullet$); using the approximate method (ω_A, W_A, G_A); and with perfect vigilance (ω=1, W_1, G_1). *a*=1; *c*=0.3.

	N					
	5	10	20	30	50	100
ω_{eq}	0.49	0.28	0.15	0.10	0.06	0.03
$\omega^\bullet$	0.47	0.22	0.11	0.07	0.04	0.02
ω_A	0.55	0.33	0.18	0.12	0.07	0.04
W_{eq}	0.775	0.796	0.806	0.810	0.813	0.815
$W^\bullet$	0.823	0.856	0.863	0.868	0.860	0.862
W_A	0.749	0.781	0.799	0.806	0.811	0.815
W_1	0.7	0.7	0.7	0.7	0.7	0.7
G_{eq}	0.851	0.833	0.825	0.822	0.820	0.819
$G^\bullet$	0.958	0.917	0.891	0.887	0.870	0.867
G_A	0.840	0.828	0.824	0.822	0.821	0.820
G_1	0.778	0.738	0.718	0.712	0.707	0.704

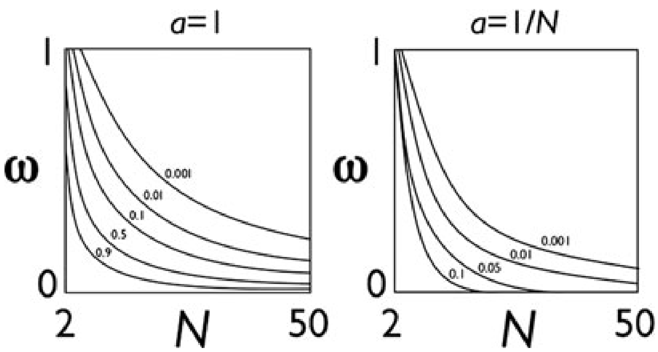

Figure 3. **The optimal vigilance level. The vigilance level (ω) at equilibrium as a function of group size (*N*) for different values of *c* (the cost of volunteering; the values next to each curve) and *a* (the benefit of the public good). The number of witnesses *M* is common knowledge.**

bystander effect is weaker than the many eyes effect. Note that this is true both when the risk of predation is diluted and when it is not (Fig. 1); therefore the selfish herd effect is not necessary to explain relaxed selection when N is very small or $\omega \approx 0$ (the many eyes effect is enough).

Result 2: With reduced vigilance, irrespective of the dilution of risk (the selfish herd effect), increasing group size relaxes selection against predation only if group size is very small or the probability to detect the predator is very low: only in this case is the many eyes hypothesis a possible explanation for the group size effect.

OPTIMAL VIGILANCE

So far we have assumed that individual vigilance (ω) is a fixed parameter; let us consider now what happens if ω can evolve. What is the optimal level of vigilance? An individual that reduces his own vigilance will pay the cost of volunteering c less often; the optimal probability of volunteering, however, will change according to the average vigilance, which includes the reduced vigilance of the mutant, and will affect the probability that the public good is produced. Figure 3 and Table 1 show some of the results. It is clear that the optimal individual vigilance level is lower than 1 and decreases with N and c.

As we have seen, relaxed selection is not a tenable explanation for the reduction of individual vigilance; moreover we have assumed that vigilance per se has not direct costs, therefore there is no direct benefit from reducing it. Why does individual vigilance decrease then? The explanation is the following. Reducing one's vigilance is equivalent to reducing the effective group size (the number of individuals that see the predator and therefore are actually able to volunteer); reducing group size, as we have seen, *increases* the production of the public good (the contrary of the bystander effect: increasing group size reduces the probability that someone gives the alarm). Therefore reducing individual vigilance can be a deliberate strategy to induce other players to volunteer more often (rather than the result of relaxed selection as posited by the "selfish herd" and the "many eyes" hypotheses).

Why does not vigilance decline to zero then? Because if vigilance is too low the probability that nobody sees the predator increases; reducing vigilance therefore is risky—it induces others to volunteer more often, thereby increasing the probability that someone gives the alarm, but if vigilance is too low there is a risk that nobody actually sees the predator. The equilibrium value, therefore will be lower than 1 but not zero (Fig. 3).

Result 3: The optimal level of individual vigilance is lower than 1 and declines with group size because reducing vigilance reduces the number of individuals that can raise the alarm, which *increases* the probability that at least one gives the alarm. Vigilance stops declining when a further reduction makes the risk of not detecting the predator too high.

BRINKMANSHIP

As we have seen a deliberate increase of risk (by reducing individual vigilance) is favored by selection because it induces everybody else in the group to volunteer more often. This provides a strategic interpretation of the group size effect for vigilance. The result, however, goes beyond the explanation of the group size effect and has broader implications for the theory of public goods.

Note that the highest probability to produce the public good (the alarm) is obtained with few individuals and high vigilance and that the highest fitness is obtained with many individuals and low vigilance (Fig. 1); the maximum self-interest and the maximum public good are not obtained for the same combination of parameters (N and ω), as is expected for a social dilemma. Note, however, that for a given N a reduction of vigilance, while it evolves for the benefit of the individual, incidentally also leads to a benefit for the group. With $N = 30$ and $c = 0.3$, for example (see Table 1), with perfect vigilance ($\omega = 1$) fitness is 0.7 and the probability that someone gives the alarm is about 71%; if vigilance can evolve it will reach the optimal value (that maximizes fitness) $\omega = 0.10$; with this level of individual vigilance fitness is 0.81 (16% higher than with $\omega = 1$) and the probability to produce the public good is about 82% (15% higher than with $\omega = 1$). The social dilemma is not solved completely but the production of the public good improves.

Result 4: Reducing one's ability to volunteer in a volunteer's dilemma can evolve by natural selection as a strategic behavior to maximize individual fitness, but incidentally it also improves the production of the public good.

Discussion

IMPLICATIONS FOR THE GROUP SIZE EFFECT

The group size effect (the fact that vigilance against predators decreases when group size increases) has been traditionally explained invoking relaxed selection: when group size increases, individuals are allowed a lower level of individual vigilance because the risk of predation is diluted (the "selfish herd" hypothesis) or because more individuals can detect the predator more efficiently (the "many eyes" hypothesis). These two hypotheses are still debated. I have shown that they are not theoretically grounded, except for a specific set of parameters, based on the results of the volunteer's dilemma and the bystander effect: the more individuals witness the predator approaching the *less* likely it is that someone volunteers to give the alarm; therefore selection is actually not relaxed when group size increases, which undermines both the "selfish herd" hypothesis and the "many eyes" hypothesis as explanations for the group size effect.

I have then shown that reduced vigilance can be explained, instead, as a deliberate strategy (brinkmanship) to induce other players to volunteer more often; increasing the risk that the predator is not detected (by reducing vigilance) *increases* the probability that someone will give the alarm. This is not intuitive, but it is the same logic as the bystander effect. Individual vigilance decreases because of strategic reasons, not because being less vigilant has direct energetic benefits (note that we have assumed no cost for vigilance). Reducing vigilance, therefore, is not a behavior that individuals are allowed to adopt when selection is relaxed (as assumed by the "selfish herd" or by the "many eyes" hypothesis) but rather may be interpreted as a deliberate strategy to improve the probability that the public good (the alarm) is produced. Clearly this need not be a rational decision; it simply means that selection will favor a reduction in individual vigilance if this reduction induces other group members to give the alarm more often.

The group size effect is commonly observed in many species, but not all (Beauchamp 2008). Explaining this diversity is an open question. My suggestion is that this could be due to different degrees of within-group relatedness. As I have shown previously (Archetti 2009a) if group members are related, the bystander effect disappears above a certain group size N_T(which is a function of relatedness); this means that reducing individual vigilance would not be an effective way to improve alarm calls if group members are highly related, and therefore that a reduction of vigilance would not be favored by natural selection in that case. An extension on the present game that included within-group relatedness would allow to produce testable predictions for this idea.

IMPLICATIONS FOR THE PRODUCTION OF PUBLIC GOODS

The volunteer's dilemma is a social dilemma; like the prisoner's dilemma (Tucker 1950) it is Pareto-deficient (although in the volunteer's dilemma the equilibrium is mixed). In social dilemmas, a free-rider problem arises because individuals pursue self-interest and can exploit the public good produced by others (Olson 1965); cooperation in public goods games can be improved by introducing relatedness between group members (which allows kin selection; Hamilton 1963, 1964, 1996; Frank 1998) or repeated interactions (which allow reciprocation, reputation effects, and punishment; Axelrod and Hamilton 1981, Doebeli and Hauert 2005). In the case of brinkmanship, instead, the public good improves by means of a selfish, strategic behavior that does not require relatedness nor repeated interactions.

The idea of increasing risk as a strategy (brinkmanship) to induce the opponent to adopt a certain behavior dates back at least to Schelling's discussion of credible commitments in the cold war (Schelling 1960) and is well known for the game of chicken (Rapoport and Chammah 1966). Here, I have shown that brinkmanship can also improve cooperation in a social dilemma; in particular, in the case of alarm calls against predators, brinkmanship corresponds to a reduction of vigilance; in general, for social dilemmas that can be modeled as a volunteer's dilemma (i.e., in any case in which the costly contribution of one volunteer is enough to produce a public good), brinkmanship corresponds to reducing one's own ability to volunteer. The idea could be extended to social dilemmas in which more than one individual is necessary to produce the public good and tested in controlled behavioral experiments (M. Archetti, unpubl. ms.).

We study the theory of public goods to understand why cooperation exists in nature and to design policies that may improve the production of public goods in our societies. Although relatedness and repeated interactions can explain many examples of cooperation in nature, they do not provide practical ways to improve cooperation in situations of conflict: relatedness cannot be increased and repeated interactions cannot be imposed by the players. Brinkmanship, instead, improves cooperation because of the strategic behavior of self-interested individuals, which incidentally also leads to a benefit for the group.

Selfish, rational individuals will not choose to pay too large a cost to produce a public good; if they are fully rational, however, they can remain selfish and, by adopting the correct strategic behavior, increase their payoff and at the same time improve the benefit for their group.

LITERATURE CITED

Archetti, M. 2009a. The volunteer's dilemma and the optimal size of a social group. J. Theor. Biol. 261:475–480.

———. 2009b. Cooperation as a volunteer's dilemma and the strategy of conflict in public goods games. J. Evol. Biol. 22:2192–2200.

Axelrod, R., and W. D. Hamilton. 1981. The evolution of cooperation. Science 211:1390–1396.

Beauchamp, G. 2003. Group-size effects on vigilance: a search for mechanisms. Behav. Proc. 63:111–121.

———. 2008. What is the magnitude of the group-size effect on vigilance? Behav. Ecol. 19:1361–1368.

Clutton-Brock, T. H. 1999. Selfish sentinels in cooperative mammals. Science 284:1640–1644.

Diekmann, A. 1985. Volunteer's Dilemma. J. Conflict Resolution 29:605–610.

Doebeli, M., and C. Hauert. 2005. Models of cooperation based on the prisoner's dilemma and the snowdrift game. Ecol. Lett. 8:748–766.

Frank, S. A. 1998. Foundations of social evolution. Princeton Univ. Press, NJ.

Hamilton, W. D. 1996. Narrow roads of gene land, vol. 1. Freeman, Oxford, UK.

———. 1963. The evolution of altruistic behavior. Am. Nat. 97:354–356.

———. 1964. The genetical evolution of social behaviour. J. Theor. Biol 7:1–52.

———. 1971. Geometry for the selfish herd. J. Theor. Biol. 31:295–311.

McNamara, J. M., and A. I. Houston. 1992. Evolutionarily stable levels of vigilance as a function of group size. Anim. Behav. 43:641–658.

Olson, M. 1965. The Logic of collective action. Harvard Univ. Press, Cambridge, MA.

Parker, G. A., and P. Hammerstein. 1985. Game theory and animal behaviour. Pp. 73–94 *in* P. J. Greenwood, P. H. Harvey, and M. Slatkin, eds. Evolution: essays in honour of John Maynard Smith. Cambridge Univ. Press, Cambridge.

Pulliam, H. R. 1973. On the advantages of flocking. J. Theor. Biol. 38:419–422.

Pulliam, H. R., G. H. Pyke, and T. Caraco. 1982. The scanning behavior of juncos: a game-theoretical approach. J. Theor. Biol. 95:89–103.

Roberts, G. 1996. Why individual vigilance declines as group size increases. Anim. Behav. 51:1077–1086.

Rapoport, A., and A. M. Chammah. 1966. The game of chicken. Am. Behav. Scientist 10:10–28.

Schelling, T. C. 1960. The strategy of conflict. Harvard Univ. Press. Cambridge, MA.

Searcy, W. A., and S. Nowicki. 2005. The evolution of animal communication. Princeton Univ. Press, Princeton, NJ.

Tucker, A. 1950. A two-person dilemma. Pp. 7–8 *in* E. Rasmusen, ed. Readings in games and information, 2001. Blackwell, Oxford.

Associate Editor: M. Doebeli