How to Choose Your Victim

KLAUS ABBINK, GÖNÜL DOĞAN[1]

We introduce the experimental mobbing game. Each player in a group has the option to nominate one of the other players or nominate no-one. If the same person is nominated by all other players, he loses his payoff and the bullies gain. We conduct three sets of experiments to study the effects of monetary gains, focality and fear of being bullied. We find that subjects frequently coordinate on a victim, even for only modest gains. Higher gains make mobbing more likely. Richer and poorer players are equally focal. We find no evidence that fear of becoming the victim explains bullying.

1. Introduction

Bullying inflicts severe consequences such as depression, loneliness, anxiety, and low self-esteem on the victims. Moreover, bullying is highly prevalent both at schools and in the workplace, but also in social networks in the internet. The estimates of students involved as a victim, perpetrator, or both range from 20% to 56% in the US[2], whereas between 3% and 15% of all employees are estimated to be victims of serious bullying in the workplace (Zapf et al. (2010)). The economic damages from workplace bullying are enormous. The Australian Productivity Commission (2010) estimates the costs to employers to be between 6 and 36 billion Australian dollars per year; the latter figure corresponds to roughly 2.5% of the country's GDP. Finding strategies against bullying is therefore important, and understanding its causes is a primary prerequisite for it.

In the field, bullying behaviour is notoriously hard to observe, which explains the wide margins in the estimates cited above. Much of it takes place below the surface, not least because victims often feel too embarrassed to come forward, and perpetrators cover their tracks. Questionnaire studies deliver valuable insights, but do not allow full control over the accuracy of the data, nor do they allow replicating cases of bullying. To complement the existing literature, we suggest a novel approach to the analysis of bullying. In a laboratory experiment, we can, under controlled conditions, construct a game that captures essential features of a potential bullying scenario. By varying the conditions across treatments, we can identify factors that contribute to bullying or prevent it to occur, and test hypotheses about the motivations underlying the observed behaviour.

In this study we look at mobbing, a particular form of bullying in which an individual is harassed by a group of tormentors. This is a more sophisticated setting than a one-on-one encounter, as it requires the coordination of several individuals, some of whom may be reluctant to take part. In fact, some might only participate to avoid challenging the group and not

---

[1] Abbink: Department of Economics, Monash University, Wellington Road, Clayton, VIC 3800, Australia. klaus.abbink@monash.edu. Doğan: CREED, University of Amsterdam, Roetersstraat 11, 1018WB Amsterdam, the Netherlands. gonul2517@gmail.com.

[2] See Centers for Disease Control and Prevention, (2011), Borowsky et al. (2013), or Kowalski and Limber (2013).

become a victim themselves. Situations can easily arise in which groups become locked in – bullying continues because no-one can simply leave the group without risk. We capture this group dynamic in a simple experimental game, which we call the *mobbing game*. In a group of players, each player can, but is not forced to, nominate a victim among the group members. If players nominate different individuals or none, then there is no victim, and players receive their default payoffs. If all other players nominate the same individual, then this individual becomes the victim, his payoff is taken away and the bullies receive an additional payoff. What we obtain is a new, darker incarnation of an experimental coordination game – players do not coordinate to find the best common choice, but to conspire against a fellow group member. Our new paradigm allows us to study the coordination problem involved in mobbing, and thereby address some important questions: (1) How do individuals coordinate on choosing a victim, (2) Who is most likely to become a victim, and (3) What motivates individuals to choose a victim?

Mobbing in the workplace is typically detrimental to economic efficiency, but there are situations in which groups coordinate on a victim to actually enhance efficiency. A drastic example, though now of mainly historical significance, is the *Custom of the Sea*. This informal, but once universally accepted code prescribed that shipwrecked sailors were to draw lots once they were out of supplies. The loser was killed and eaten to preserve the chances of survival for the rest of the group. The Custom demanded a random draw; however, in the few recorded cases the choice of the victim was far from random.[3] More mundanely, think of a group of friends choosing a driver for their weekly pub crawl. The loser suffers a cost having to drive everybody home, but the final outcome is more efficient than the uncoordinated solution in which everybody drives on their own. In our experiment, we look at both efficiency-enhancing and efficiency-reducing mobbing scenarios.

We conduct three sets of experiments, comprising altogether eight treatments. The first set uses a symmetric setting without any focal players or coordination devices. We observe that nevertheless, players find it remarkably easy to pick a victim, and are surprisingly inclined to do it even if the gains for the bullies are modest. In the second set we introduce an asymmetry, making one of the players either richer or poorer than the others. The player who is different is by far the most likely to become the victim, no matter whether the odd one out is richer or poorer by default. Finally, a third set of experiments is designed to disentangle possible motives behind bullying. We find that the pursuit of material gain drives bullying, not the fear of becoming the victim oneself.

To our knowledge, there are no studies that propose a similar game to ours. The closest studies to the present work are coordination games, which are due to the work of Van Huyck et al. (1990, 1991). In these games, players have to coordinate on one of multiple Pareto ranked equilibria. In their experiments subjects did poorly – they typically coordinated on the worst equilibrium. Such coordination failure has been replicated in various experiments and is

[3] For example, the *Mignonette* sank in 1884 with four crew members on open sea. When supplies ran out, the cabin boy was chosen by the other three to be killed and eaten. We do not know for certain whether he was selected because he was already ill, or because he was the lowest-ranked crew member. What we do know is he was not chosen randomly (see Hansen (2000)).

found to be dependent on the attractiveness of the efficient outcome, the time horizon, the group size, the deviation costs, refined action space, and communication (for a meta-analysis on coordination experiments see Devetag and Ortmann (2007)). Cooper et al. (1990, 1992) introduced the stag-hunt game, in which one equilibrium is Pareto dominant and the other is risk-dominant, and find frequent coordination failure. Isoni et al. (2013) and Poulsen et al. (2013) analyse solutions to coordination problems with respect to their saliency. In all these studies, there is no victim; all players benefit from successful coordination. The model most akin to ours is the collective resistance game by Cason and Mui (2007), in which two followers must coordinate if they want to challenge a transgressing leader. The results show that followers are very successful in coordinating their resistance, unless the leader is allowed to use divide-and-conquer strategies, i.e. target individual followers rather than both. The endogenous victim selection, which is at the heart of our study, is not part of their research agenda.

2. The experimental design

All experiments were conducted at the Center for Research in Experimental Economics and Political Decision Making (CREED) of the University of Amsterdam. Subjects were recruited via the online recruitment system of CREED and were mostly undergraduates from a wide variety of majors. Each subject could participate in only one session, and all treatments were across subjects, i.e. in each session only one treatment was run. There were 8 treatments conducted across 21 sessions. Average pay was 13.24 euros including 7 euros show-up fee. At the time of the experiment, the exchange rates to other major currencies were approximately US-$1.30, UK-£0.85, JP-¥105 and CN-¥8.50 for one euro. The experiments lasted between 40 minutes and 1 hour 15 minutes including the time spent on reading the instructions.

The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). The experiments were in English, the language of instruction at the university. The instructions used neutral language; a copy of the instructions for the baseline treatment as explained below is included in the Appendix. It was explained in the instructions that each group consisted of four players (or five, in treatments we describe later) who were randomly selected at the beginning of the experiment and that the group composition stayed the same throughout the whole experiment (partners design). Subjects had no way of knowing which of the other participants were in their group. In the instructions and the experiment, the players were denoted as M, T, G, and P, a seemingly random set of consonants.[4] Player labels were randomly assigned at the beginning of the experiment and remained fixed throughout the experiment.[5] Subjects were informed that the task would be repeated for 20 periods and that their final earnings comprised of the total points they earned in the experiment converted at a rate of

[4] The labels are also a hidden homage to the inmates Mather, Travers, Greenhill and Pearce, who escaped from a Tasmanian prison camp in a group of eight in 1822, only to get lost in the forest. When food ran out, the four conspired to apply the Custom of the Sea to the others. When no-one else was left, they turned to killing and eating one another, until only Pearce survived. All victims were chosen in decidedly non-random ways. This story is one of the great Australian foundation myths, and it was an inspiration for this study (for a dramatic reconstruction see *Van Diemen's Land* (2009)). We are confident that none of our Amsterdam subjects made that connection.

[5] In the treatments with five players we also used the label R, with no historical reference.

1.25 euro cents per point plus the show up fee. After the instructions were finished, subjects played one practice period with full feedback in order to familiarise them with the procedure and the screens.

Since subjects received no feedback from other groups and the groups remained the same, each group constitutes a statistically independent observation.

3. The first set of experiments: The symmetric baseline

The basic game is very simple: Each member of a group of four players[6] decides simultaneously to nominate one of the other players in their own group, or not to nominate anyone. Nominating as such is neither costly nor beneficial. If three players nominate the same fourth player, then the nominee loses all of his earnings whereas the nominators earn additional money. The game is repeated for 20 periods with the same players. After each period, subjects were given feedback about the number of nominations each player had received, but not from whom.

Table 1. Payoffs in the symmetric treatments

Situation	Treatment	Player 1	Player 2	Player 3	Player 4	Total
Default	All	24	24	24	24	96
		Victim	Player 2	Player 3	Player 4	
Bullying	High	0	48	48	48	144
	Medium	0	32	32	32	96
	Low	0	25	25	25	75

In the three treatments of the first set of experiments, all four players are identical. They have the same payoffs in absence of bullying [7] (which we refer to as the default payoff), and the same prospective payoffs in case that they become victims or bullies. Table 1 shows the payoff distributions. In all three treatments, all players receive 24 tokens if there is no bullying, i.e. if three players do not nominate the fourth player either by refraining from nominating a victim or by miscoordination. Further, in all three treatments the victim receives zero payoff. The treatments differ in the payoffs the bullies receive in case of successful coordination, and hence in the efficiency of bullying (measured, as usual in experimental economics, as the sum of the payoffs of all players in a group). In the High treatment, bullying is efficiency-enhancing, i.e. the sum of all players' payoffs is greater than the sum of all default payoffs. In the Medium treatment bullying is efficiency-neutral, while in the Low treatment the victim's loss far outweighs the minor gains successful bullies can achieve.

Assuming own-payoff maximisation in the one-shot game, the game theoretic analysis of the mobbing game is straightforward. Any outcome in which no player receives two nominations (more generally n-2 nominations if n is the number of players) is an equilibrium outcome. In all these constellations no player can unilaterally improve their payoff. Only the case of one player being nominated exactly by two other players is not an equilibrium outcome, since the

[6] Zapf et al. (2010) find that the typical number of bullies in mobbing cases is between two and four – our choice of groups of four , hence three bullies and one victim, is well in line with this figure.

[7] In the following, we will use mobbing and bullying interchangeably. Mobbing is generally used as a subcategory – a group harasses an individual, while bullying includes other forms, like harassment by an individual.

remaining non-victim can join the bullies and improve his payoff. Hence coordination on a victim is an equilibrium just as well as non-coordination.

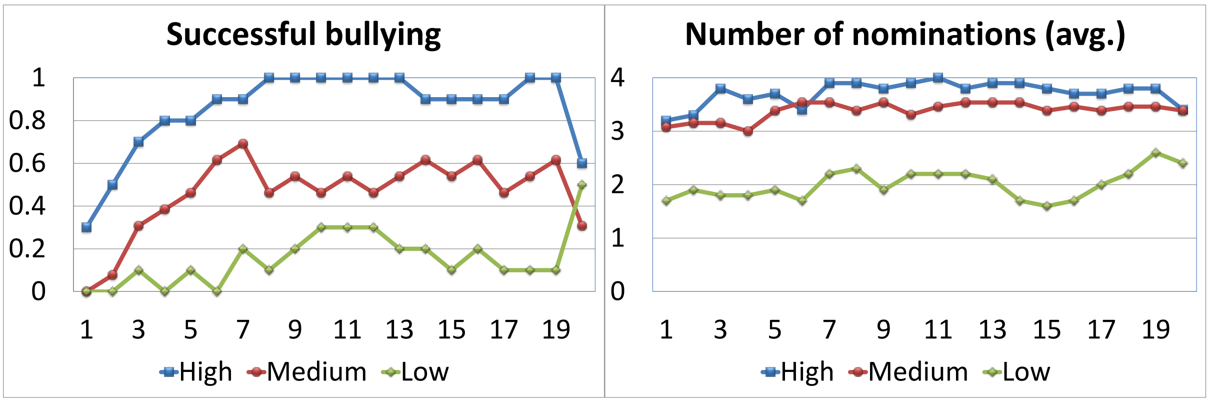

Figure 1

The left-hand panel of figure 1 shows the relative frequency of successful bullying, i.e. the fraction of groups which coordinated on the same victim, over the 20 periods of the experiment. The right-hand panel shows the frequency of nominations, i.e. the fraction of individuals who nominated a player, regardless of whether coordination could be achieved.[8] Table 2 below shows the overall average rates of successful bullying across all groups and periods. We can make three main observations. First, bullying becomes the more frequent the higher the individual gains from it are. The overall frequencies of successful mobbing are 86.0% in High, 47.3% in Medium, and 13.5% in Low. We can reject the null hypotheses of no difference for all pairwise comparisons.

Table 2. Bullying rates in the first set of experiments

	Bullying rate (%)	Pairwise comparisons*
High N=10	86.0 (11.7)	p=0.010 (M-H)
Medium N=13	47.3 (37.0)	p=0.018 (L-M)
Low N=10	13.5 (21.4)	

*Mann-Whitney exact test, 2-tailed.

A second observation is that if there is a will to choose a victim, coordination is remarkably easy to accomplish even though all players are identical and have no coordination devices at their disposal. Coordination happens quickly: If we look at the groups in which there existed a victim for at least three periods[9], we see that it took on average 2.8, 3.9 and 5.0 periods in the High, Medium and Low treatments, respectively. Such quick coordination on a victim may seem surprising, but it turns out that in the repeated game context, the coordination problem is not as hard as it seems. If every player arbitrarily nominates another player, it is actually quite likely that one player receives two nominations by chance. It then becomes easy for the remaining non-

[8] Note that nominating a victim is inconsequential if there are less than two other nominations on the same victim, so we cannot say for certain that every nomination is carried out with the intention of forming a mobbing group. Thus, while nominating behaviour is certainly informative, we focus predominantly on successful bullying in our analysis.

[9] There are 10, 9, and 3 such groups in the High, Medium and Low treatments, respectively. Our choice of a threshold of three periods is of course somewhat arbitrary.

victim to jump on the bandwagon. Once full coordination is achieved, the incentive for any bully to jump off again is limited – it not only leads to an immediate loss of earnings, but also bears the risk of becoming the next victim. The victim himself is defenceless.[10]

Finally, we can observe that subjects show a surprisingly strong propensity to engage in bullying. Even in the Low treatment, where the gain from bullying is minuscule compared with the damage it does to the victim, on average about half of all subjects nominated a victim. This indifference towards a fellow group member's distress is somewhat at odds with traditional experimental findings that subjects care much about fairness and equality (see the vast literature on dictator games since Hoffman et al (1994) and Forsythe et al. (1994)). Two motives could drive the high frequency of bullying. Subjects may be motivated by the additional money they can gain by coordinating on a victim. Or alternatively, they might nominate a player in order to avoid becoming a victim themselves. Note that it requires all three non-victims to establish mobbing, thus each player has veto power against mobbing any other player. If a player does not nominate, and the other players are inclined to bully, then this player might fear that he eventually becomes the only possible victim. Hence, even in absence of material motives, fear of getting into this situation may drive subjects towards nominating.[11] In our third set of experiment (section 5) we design treatments that disentangle these possible motives.

In the High treatment, it could be a sensible group behaviour to rotate the victim's role. This would capture the efficiency gain without victimising a particular individual. However, we hardly observe any such behaviour. All groups in this treatment coordinated on one victim more than half of the periods, and 8 out of 10 groups coordinated on the same victim at least 80 per cent of the time. In the other treatments, no collective gain can be achieved through taking turns, and as expected, we do not see much of it either. In Medium, 5 out of 13 groups coordinated on the same victim 80 per cent of the time; in Low, there were two groups that coordinated on the same victim in 7 and 13 periods, respectively. All these figures correspond well to the overall mobbing rates, so if there is mobbing, the victim's role typically falls to the same person.

4. The second set of experiments: The odd one out

A plethora of studies suggests that people often become bullying victims because they are different. When asked, school bullying victims state being different in appearance or in other things as the most important reasons for being bullied (Erling and Hwang (2004), Frisén et al. (2008)). In the workplace, envy seems to be an important reason for bullying (Zapf and Einarsen (2010), and references therein). In the second set of experiments we capture such

[10] It is of course possible that the easy coordination was due to the labelling of players, which may have made one player particularly focal. However, our data do not show a strong bias. On average across all three treatments, player M received 0.77 nominations, player T 0.58, player G 0.74 and player P 0.98 (we look at nominations because focality through labelling should express itself in more nominations). All the differences except the one between M and G are significant and might contribute to the quick coordination.

[11] Abbink and de Haan (2013) find fear to be a powerful driver of behaviours that are hurtful to others.

Table 3. Payoffs in the asymmetric treatments

Treatment	Situation	Focal	player 2	player 3	player 4
Poor	Default	16	24	24	24
	Bullying Focal	-8	32	32	32
	Bullying other	24	32	32	0
Rich	Default	32	24	24	24
	Bullying Focal	8	32	32	32
	Bullying other	40	32	32	0

asymmetries. We use the same payoffs as in the Medium treatment of the first experiment[12], with the exception that we modify the default payoff of one player. By changing this payoff we not only make this player naturally stand out, but we also change the relative position of earnings. We design two new treatments: In the Rich treatment, all payoffs for one player are increased by 8 tokens, in the Poor treatment all payoffs are reduced by the same amount. The additional payoff from bullying remains the same in Rich, Poor, and Medium, so they remain comparable in this respect. Table 2 summarises the possible payoff constellations.

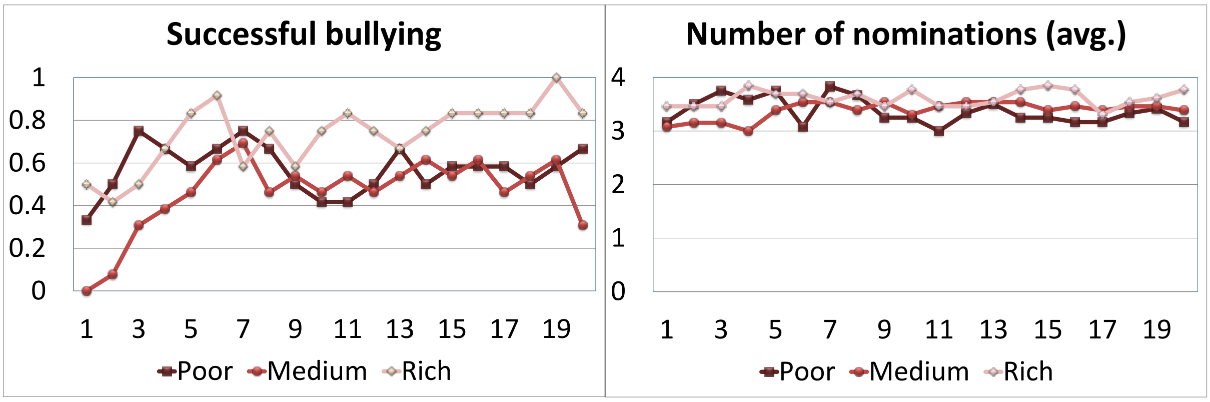

Figure 2

Table 4. Bullying rates in the second set of experiments

	Bullying rate (%)	Pairwise comparisons
Medium	47.3	p=0.024 (M-R)
N=13	(37.0)	p=0.443 (M-P)
Rich	74.2	
N=13	(29.4)	p=0.188 (R-P)
Poor	57.1	
N=12	(32.4)	

*Mann-Whitney exact test, 2-tailed.

Figure 2 and Table 4 show the frequency of bullying over periods and the average bullying frequency per treatment, in the same fashion as figure 1 and Table 2 of the first set of experiments. Coordination occurs even faster in the asymmetric treatments than in the symmetric case: The average number of periods before the first coordination on a victim were 2.5 in the Rich and 2.3 in the Poor treatments[13] (the difference between Rich and Medium is significant at p=0.038, and between Poor and Medium marginally significant at p=0.095, both according to two-sided exact Mann-Whitney tests). So focality matters. This

12 We use the Medium treatment as the baseline because this generated mobbing rates of about one half – this means the results leave most room for treatment effects in both directions.

13 Again, we include all groups that coordinated on the same victim at least three times.

finding is clearly corroborated when we look at who is most likely to become the victim in the asymmetric treatments. Figure 3 shows that in three quarters of all cases of successful mobbing the focal player is the victim. This is regardless of whether the focal player is richer or poorer than the others; the fractions are very similar in both treatments. The observed percentages are highly significantly different from 25 per cent (which would be the outcome of randomly choosing the victim) according to the two-sided Wilcoxon signed-rank test with corresponding p values of 0.004 and 0.009. Note that there is only one focal, but three non-focal players; this means that a focal player was indeed nine times more likely to become the victim.

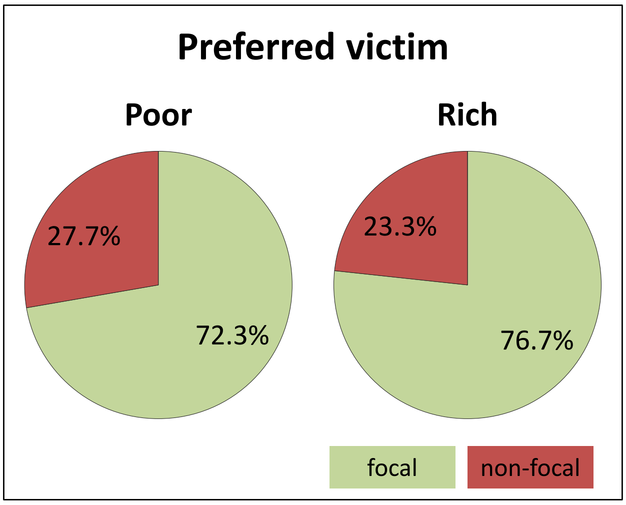

Figure 3

While the fraction of focal victims is very similar in both treatments, there is some evidence that the relative payoff of the focal player affects mobbing frequencies. Overall rates are similar in Medium and Poor (47.3% versus 57.1%, not statistically significant), but if the focal player is richer, bullying rates rise to 74.2%. This is significantly higher than in the Medium treatment. Envy towards the richer player seems to play some role in the likelihood of the richer player becoming a victim, while pity towards the poor does not.

5. The third set of experiments: Greed or fear?

The results of the first and second set of experiments leave an important question open: Do players engage in bullying because they want to pursue material gain from mobbing, or out of fear of becoming the victim if they do not? To shed light on this question, we designed three new treatments. In the first of the new treatments, the set-up is the same as in the Medium treatment of the first set, with one difference: One of the four players is designated as the "safe" player. We call this the Safe-4 treatment (the '4' stands for the number of players, which we will vary). If the other three choose him as the victim, all players still receive their default payoffs. Trying to mob the safe player is simply ineffective (think of the safe player as the kid who is so big and strong that even all others united cannot hurt him). For the safe player, the fear of becoming the victim is ruled out as an explanation for bullying behaviour. As an indirect effect, the presence of a safe player should also calm the non-safe players' fears. If the safe player is less likely to bully, they are immediately less likely to be bullied, since no action can be taken against them without the safe player (recall that each player has veto power against mobbing somebody else). So if fear plays a major role, we should observe less mobbing in the treatment with the safe player compared to the Medium treatment.

This manipulation is open to the objection that the introduction of a safe player also reduces the number of victim candidates from four to three. This might make the task of coordinating on a victim easier, and this could work against the reduction of bullying through the introduction of the safe player. Hence, by comparing mobbing rates in Safe-4 and Medium we could identify the presence of a fear effect (if we observe a drop in bullying, we do so *despite* having made bullying easier), but not its absence. If we do not observe any reduction in mobbing we could not clearly infer that fear plays no role, because a possible fear effect could be neutralised by easier coordination. To prepare our test for both scenarios, we introduce another treatment in which we also include a safe player – hence partially removing fear –, but make bullying harder than in Medium. We do so by introducing a fifth player to the game. One of the five is a safe player, the other four are not (we refer to this treatment as Safe-5). This sets the number of victim candidates back to four. Coordination on a victim should thus no longer be easier than in Medium. In addition, the number of bullies required increases to four, adding another difficulty to the coordination task. Hence, if we do not observe a drop in bullying rates *despite* having made it harder, we can dismiss the hypothesis of bullying out of fear with some confidence.

Table 5. Payoffs in the safe player treatments

Treatment	Situation	Safe Player	player 2	player 3	player 4	player 5
Safe-4	Default	24	24	24	24	
	Bullying	32	32	32	0	
Safe-5	Default	24	24	24	24	24
	Bullying	30	30	30	30	0

Figure 4 shows the frequency of successful mobbing in the safe player treatments. These figures do not include cases in which players coordinated on the safe players – since this is ineffective, we do not count it as mobbing. In neither of the new treatments, the introduction of the safe player has any effect in reducing the incidence of mobbing, not even in the Safe-5 treatment, where the effect might be assisted by a greater difficulty to coordinate. Mobbing rates are very similar, and statistically indistinguishable, in Medium, Safe-4 and Safe-5. Hence, we do not find evidence that fear of becoming the bully explains the high mobbing rates in our experiment.

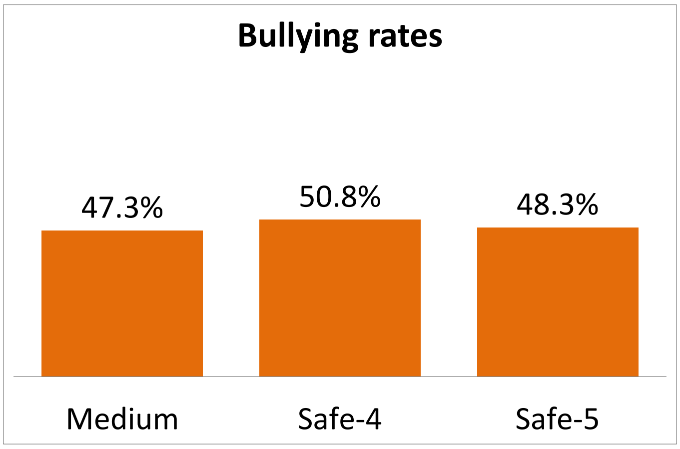

Figure 4

Finally, we analyse whether we can detect mobbing out of fear at the individual level. Although we do not observe a drop in mobbing frequencies through the safe players, they may still be less likely to engage in bullying. This question cannot be answered with the existing treatments. By design, any successful bullying in Safe-4 and Safe-5 must involve the safe player, since it requires all players except the victim to coordinate. To identify differences in bullying behaviour between safe players and others we design a new treatment, which we call Safe-5-3. In this setting there are five players, but only three nominations on the same player are sufficient to establish a bullying group. This way it is possible to have bullying groups that do not include the safe player, and we can detect whether the safe player is more or less likely to be part of a mob than the non-safe players. Table 4 shows the payoffs for the case without mobbing, and with bullying groups that consist of either three or four players.[14]

Table 6. Payoffs in Safe-5-3

Treatment	Situation	Safe Player	player 2	player 3	player 4	player 5
Safe-5-3	Default	24	24	24	24	24
	Bullying (3 Nom)	32	32	32	24	0
	Bullying (4 Nom)	30	30	30	30	0

In the new treatment bullying rates are very high at 76.9%, significantly higher than in both other safe treatments. Perhaps this is unsurprising, since it requires only three out of five players to be inclined to bullying in order to establish a mob. The more important question, though, is whether safe players are more likely to be part of a mobbing group than others. To answer this, we look at all rounds in which a mob consisting of exactly three players has formed. This is the case in which bullying groups with or without the safe player are possible. We observe that the safe player is a member of a three-player mob in 82.4% of all instances. This is even higher than the chance rate of 75%, though not significantly so.[15] Hence, in none of our tests do we find any evidence that removing fear of being bullied from the game reduces the frequency of mobbing in any way. Such fear therefore does not seem to be a likely driver of the frequent bullying we observe in the previous treatments.

6. Conclusion

We introduce the mobbing game, the first experimental paradigm to study the endogenous formation of bullying groups. The eight treatments of our experiment allow important insights into the nature of mobbing. Our results show that spontaneous formation of such groups is easy, and surprisingly many subjects are inclined to engage in bullying, showing little concern for the consequences to the victim. Nominations of victims are frequent even if the gains from it are minimal, and individuals do not hesitate to pick a fellow group member who is already poorer than themselves. In this case, the victim's relative poverty only serves as a focal point and a coordination device. Although fear of becoming a victim might seem a likely explanation for such 'aggressive' behaviour, our experiments do not produce any evidence for this.

[14] In the five-player treatments, we chose to keep the sum of all payoffs constant to ensure best possible comparability. This implies that each bully receives a payoff of 30 rather than 32 if a group of four bullies emerges.

[15] There are 20 possibilities to randomly form a group of three bullies and one victim out of five players. In 12 of them the safe player is part of the bullying group. In four of the remaining eight the safe player is the victim, thus these groups do not count as bullying groups. Consequently, by chance the probability to randomly draw a group with the safe player is 12/16=75%.

Our results underline the importance of studying the group dynamics that lead to bullying, especially keeping in mind that our experimental subjects are from a very conventional subject pool, not one that would be naturally driven towards aggressive behaviour. Of course, our findings are not the last word in the matter. In this study we can touch only a few aspects of the strategic environment underlying bullying. However, our model can easily be extended to capture other potentially important variables as well, or to be applied to more specific subject pools.

References

Abbink K, de Haan T (2013): Trust on the Brink of Armageddon: The First-Strike Game. Working paper, University of Amsterdam.

Auf der Heide J (Dir.) (2009). *Van Diemen's Land.* Madman Entertainment. Film.

Borowsky IW, Taliaferro LA, McMorris BJ (2013): Suicidal thinking and behavior among youth involved in verbal and social bullying: risk and protective factors. *Journal of Adolescent Health*, S4-S12.

Cason T, Mui V-L (2007): Communication and coordination in the laboratory collective resistance game, Experimental Economics, 10, 251.

Centers for Disease Control and Prevention (2012): Youth risk behavior surveillance United States, 2011. MMWR, Surveillance Summaries 2012;61(no. SS-4). www.cdc.gov/mmwr/pdf/ss/ss6104.pdf.

Cooper R, DeJong D, Forsythe R, Ross T (1992): Selection Criteria in Coordination Games. *American Economic Review* 80, 218-233

Cooper R, DeJong D, Forsythe R, Ross T (1992): Communication in Coordination Games. *Quarterly Journal of Economics* 107, 739-771.

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

Erling A, Hwang P. (2004):.Swedish 10-year-old children's perceptions and experiences of bullying. *Journal of School Violence* 3, 33-43.

Forsythe R, Horowitz J, Savin N, Sefton N (1994): Fairness in simple bargaining experiments. *Games and Economic Behavior*, 6, 347–369.

Frisén A, Holmqvist K, Oscarsson D (2008): 13-year-olds' perception of bullying: Definitions, reasons for victimisation and experience of adults' response. *Educational Studies*, 34, 105-117.

Hansen N (2000): The Custom of the Sea. Hoboken:Wiley & Sons.

Hoffman E, McCabe K, Shachat K, Smith V (1994): Preferences, Property Rights, and Anonymity in Bargaining Games. *Games and Economic Behavior*, 7, 346-380.

Isoni A, Poulsen A, Sugden R, Tsutsui K (2013): Focal Points in Tacit Bargaining Situations: Experimental Evidence, *European Economic Review*, 2013, 59, 167-188

Kowalski RM, Limber SP (2013): Psychological, Physical, and Academic Correlates of Cyberbullying and Traditional Bullying. *Journal of Adolescent Health* 53, S13-S20

Poulsen A, Poulsen O, Bett Z (2013): How Salient is an Equal but Inefficient Outcome in a Coordination Situation? Some Experimental Evidence. Working paper, University of East Anglia.

Productivity Commission (2010): Performance Benchmarking of Australian Business Regulation: Occupational Health & Safety. Australian Government Research Report.

Van Huyck J, Battalio R, Beil R (1990): Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure, *American Economic Review*, 234-248.

Van Huyck, J, Battalio R, Beil R (1991): Strategic Uncertainty, Equilibrium Selection Principles, and Coordination Failure in Average Opinion Games. *Quarterly Journal of Economics*, 106, 885–910.

Zapf D, Einarsen S (2010): Individual Antecedents of Bullying. In Einarsen S, Hoel H, Zapf D, Cooper C (Eds.): Bullying and harassment in the workplace: Developments in theory, research, and practice. Crc Press.

Zapf D, Escartín J, Einarsen S, Hoel H, Vartia M (2010): Empirical Findings on Prevalence and Risk Groups of Bullying in the Workplace. In Einarsen S, Hoel H, Zapf D, Cooper C (Eds.): Bullying and harassment in the workplace: Developments in theory, research, and practice. Crc Press.