

Jonas Heide Smith

PhD in game studies
The IT University of
Copenhagen
(smith@itu.dk)

Tragedies of the ludic commons - understanding cooperation in multiplayer games

by Jonas Heide Smith

Abstract

Conflict, it is often assumed, is the essence of games. Modern multiplayer games, however, also rely heavily on the cooperation between players. In fact, given the rapidly increasing popularity and complexity of these games, game designers are arguably engaged in one of the most ambitious experiments with social software in recent years.

This article argues that multiplayer games raise issues of the construction and maintenance of collective resources. If players do not cooperate on such issues given game may be visited by social tension and can lose its appeal altogether. This article relates this issue to the larger issue of social order as studied by political science after which solutions to the problem are discussed.

Introduction

To play a game like the team-based shooter *Counter-Strike* is to engage in a multi-layered system of collaboration. As a team-member each players must carefully co-ordinate strategy and in casual play all players are co-responsible for living up to a number of local norms; for instance that the practice of camping is not tolerated (see Smith 2004).

This observation has not always been acknowledged within game studies which have tended to frame multi-player games in terms of in-game conflicts (following standard game definitions highlighting the central position of conflict) or as arenas of personal expression or experimentation with identity. In contrast, this article attempts to understand the role of *cooperation on the construction and maintenance of collective goods* in multiplayer games and argues that such games create a number of social dilemmas that can be (but often are not) understood and solved using knowledge derived from studies of real life communities. Various solution types are discussed (from ones imposed by game designer to those implemented by players themselves) and new ones are proposed.

Questions of potential deceit or unwillingness to share what should fairly be (or which can only be) a collective burden can fruitfully be studied under the lens

collective action theory; a school of thought anchored in economics and political science. Collective action theory in various forms has proven successful as a framework for political thought as well as a tool for shaping concrete policy but although addressing issues central to a wide range of community matters, has seen little application in game studies to date. It promises nevertheless to provide one potent explanation type for in-game social tension, to provide quite practical tools for the game designer, and to highlight strong links between game studies and worthwhile sections of the social sciences.

The following section sketches the theory which is then applied to games.

The problem of collective action

Even if a collective would benefit from a certain resource being procured (say, the construction of roads) each individual might prefer to have others bear the burden or cost while nevertheless enjoying the benefits.

This, in essence, is the problem of collective action - in some situations an individual may enjoy the benefits of a collective good without contributing to its procurement or maintenance. Consequently, nobody may wish to contribute and the good is not made available at all. Simply put, such goods are vulnerable to *free-riders*, people who take no responsibility for the maintenance or construction of the resource.

It inspires modesty that this phenomenon, this *problem of collective action*, was not discovered (at least not formalized) until Mancur Olson emphasized in 1965 that often "If the members of a large group rationally seek to maximize their personal welfare, they will *not* act to advance their common or group objectives" (Olson 1971 orig. pub. 1965).

Often this problem is alluded to by invoking the concept of a "commons". This metaphor has gained popularity to the degree where it is often applied in inappropriate contexts. The idea of a common physical resource being destroyed by individuals who do not in fact set out to destroy anything but merely follow their selfish impulses was formalized in its modern form by the economist H. Scott Gordon in 1954 (Gordon 1954). Gordon analysed the behaviour of fishermen damaging free-access fishing grounds by following their own interest (not adhering to quota). The principle was popularized in 1968 by biologist Thomas Hardin in the poignantly entitled Science article "The Tragedy of the Commons" [1]. Hardin likened the environment to a rural commons, vulnerable to a tragedy since

"the rational herdsman concludes that the only sensible course for him to pursue is to add another animal to his herd. And another.... But this is the conclusion reached by each and every rational herdsman sharing a commons. Therein is the tragedy. Each man is locked into a system that compels him to increase his herd without limit - in a world that is limited" (Hardin 1968)

Thus, for Hardin, a commonly regulated (or un-regulated) environment was often doomed to tragedy from individual exploitation. In other words, Hardin had identified one particular instance of the more general problem pin-pointed by Olson.

We should note that this destructive logic does not, in fact, require the identification of a *physical* resource. Many interactions are characterized by each individual understanding that one type of behaviour would be collectively

rational whereas another type of behaviour is tempting because it is individually rational. Such situations, sometimes modelled as variations of the classical Prisoner's Dilemma, are often studied under the more general heading of "social dilemmas" (Kollock 1998). For the purposes of this article I will not need to distinguish sharply between collective action problems and social dilemmas in general.

Social dilemmas in multiplayer games

Multiplayer gaming, like most human interaction, involves situations or dynamics that seem both unintended and unwanted by the involved parties. While conflict is endemic to the enjoyment of most games (being intended) *certain types of conflicts* are detrimental to that very same enjoyment.

To be more precise about this distinction we must introduce two perspectives that are unfortunately all too often confused: We must distinguish between the *game as rule system* and the *gaming situation* (see also Hughes 1999).

In the former case we approach any game formally, directing our attention towards the reward structures embedded in the game rules. By doing so we are considering the game as an abstract system bracketing for a moment the actual experiences of concrete players. This analytical activity corresponds roughly to classifying novels by genre or to attempt to tease out the intended (or preferred) reading of a text. Under certain conditions such an approach can be extraordinarily problematic (see also Taylor Forthcoming), while in others it is merely a way to be analytical about the structure of the artefacts being studied [\[2\]](#).

When focusing on the *gaming situation*, on the other hand, we are specifically acknowledging that this corresponds only loosely to the formal properties of a game. The experience of playing a game cannot be determined (at least not fully) by an examination of the game rules, no matter how rigorous. It is determined in part by the culture surrounding the game, the concrete gaming context and the experience and personality of the player. At the extreme end of the spectrum a game may of course also be entirely "subverted" as players exploit a game structure for playing entirely different games or for artistic expression. When, for instance, Anne-Marie Schleiner worked within *Counter-Strike* to achieve the provocatively peaceful phenomenon known as Velvet Strike [\[3\]](#) she was using the game system in ways *afforded* but not *intended* by the game designers. These different levels of analysis are exactly that - each is appropriate to the study of certain game-related phenomena but neither can answer every relevant question we may have. It is the responsibility of the game researcher instead to be clear about his or her level of analysis and to be acutely aware of which game elements are obscured by the choice of perspective.

Following this advice, we can note that whereas the game rules (as stated, for instance, in a Chess rulebook or in the code of *Pong*) set the scene for conflict, to actually engage in an enjoyable gaming situation the players must collaborate on upholding a large number of implicit rules. *Pong* (Atari, 1972) players, for instance, must agree (tacitly at least) not to push each other physically, to share the expense of playing the game etc. (Sniderman 1999). In other words: *Pong* as a rule system is a non-collaborative activity while the gaming situation surrounding *Pong* rests upon mutual coordination and the goodwill of the players. If we don't normally consider it in this way it is merely because coordination and mutual trust is so relatively easy to achieve in physical settings like arcades.

Trust formation in darkened arcades is largely out of the game designer's hands. For online games requiring players to form teams and find opponents, however, this is decidedly not the case. In fact, since the matching of players preceding the actual clashes in online games largely takes place within structures fashioned by the game developers (or their subcontractors or partners) we can consider these structures part of the game rules. Or we can at least acknowledge that they fall squarely under the responsibility of the developers.

In the following I will show that three conflict-heavy aspects of multiplayer gaming (cheating, grief play, and responsible participation) may be understood as social dilemmas.

Cheating as a social dilemma

Many online gamers go to great length to tip the scales in their favour. As one design manual dryly puts it: *It may seem weird that a significant portion of the player base is willing to do anything to win, but that's the reality of the situation* (Mulligan and Patrovsky 2003). When such activities take place outside the original game rules or outside the generally accepted implicit rules of a game, it would fall under most dictionary definitions of cheating. The cheater herself, of course, might not always share our opinion. Or she might consider herself a cheater while not agreeing that such behaviour is morally reproachable. She is, she might argue, merely playing a different game than the rest. No matter what she thinks, however, cheating is often reported as a problem by players and developers..

Accounts of cheating in games almost always invoke the eloquent example of Blizzard's *Diablo* (Blizzard Entertainment, 1996), among the first truly successful commercial online games. It is generally acknowledged that the gaming experience was seriously affected by the amount of cheating apparent among many participants. In a somewhat informal survey conducted by the gamer magazine *Games Domain* (Greenhill 1997), 35% of the *Diablo*-playing respondents confessed to having cheated in the game (n=594). More interesting, however, were the answers to the question of whether a hypothetical cheat and hack free gaming environment would have increased or decreased the game's longevity and playability. Here, 89% of the professional cheaters stated that they would have preferred not being able to cheat. This response distribution clearly tells of a social dilemma. Arguably, the players queried are tempted to cheat but understanding that this temptation applies to other players as well, would prefer that no-one (including themselves) have full autonomy.

More formally, each individual realizes that mutual honesty is the most rewarding strategy one can hope for (in an ideal world, of course, each individual would be the only one cheating) and express their wish for structural circumstances that support this choice of strategy. The existence of cheaters sparking a desire for structural changes can be witnessed on many game forums. In a thread on an *Age of Kings* (AoK) discussion board [4], for instance, a resurgence of cheating inspired a victim to urgently contact the game developers to have them put a stop to the undesired activities through technical means. Many sympathised. One player supported the original poster by declaring that *it's scum like them that make the Zone [where AoK players go for online play] a miserable place to be. Hacks can be used to give yourself inordinate ratings... This makes me feel like hitting someone.* Apart from the request for a

technical fix, two other solutions were suggested. One was to write down the names of cheaters to avoid facing them again and to warn others of their low moral fibre. The other was to play people who are clan members or only play against friends. In other words: Do not interact with strangers unless they have institutional backing (a third party vouches for them).

We must acknowledge, however, that cheating is anything but an unambiguous term. First of all cheating comes in different forms. The easiest form to categorise is cheating that works on the code level. When a code-savvy player modifies a client application against the explicit rules of an end-user-licence-agreement (or other document) his behaviour is somewhat comparable to a chess player moving his pieces into a more favourable position while the opponent is distracted. This is a violation of the original game rules.

Other forms of in-game activities invite more discussion as to their classification. A player may discover an aspect of the game which, although unplanned for by the developers, grants him an advantage. If non-destructive to the gameplay experience such features are often considered signs of sophisticated game design in single-player games. In multiplayer games, however, such creativity is sometimes labelled as "exploits" and often considered unfortunate if not downright punishable [5]. An exploit then is an activity aimed at intentionally achieving an advantage afforded but not intended by the game design. Notably, though, the player is working within the framework of the game code and arguably cannot always determine if a certain phenomenon is intended or not. This ambiguity is obvious from many virtual world discussion boards as well as from official administrator statements. For instance, the official *Star Wars Galaxies* Knowledge Base states:

âA Good General rule of thumb is: If you are doing something that gives you an uneven advantage over a MOB a player or the game system. You are most likely exploiting.

Knowingly exploiting is a serious offence and can cause disciplinary action to be taken against your account up to and including banning.â (Sony 2004)

Clearly, this leaves room for interpretation. Arguably, a sensible combat tactic is to attempt to gain any and all advantage over MOBs or other players (and it is hard to imagine an "even"™ advantage). And the word "knowingly" obviously leaves even more room for discussion [6]. Since exploits are game mechanics unintended by the designers the verdict of "exploitation" requires an impressive act of two-way mind-reading. The player must somehow guess what the designers intended and the administrators must decide if the player acted "knowingly".

While less than clear-cut which actions should be categorized as cheating we can understand the heart of the problem. Cheaters, by unbalancing the game, may ruin games based on competition (but not for players who do not engage in competition, of course). Thus, the collective good in question is the even battlefield and the social dilemma builds on the temptation to cheat. We should note however, that cheating is often not a "pure" social dilemma, since the case where "everybody does it" might just lead to an alternative game (i.e. one that is still even). Often though, it will destroy much of the game's™ appeal, particularly if the game relies on carefully balanced units and game terrains such as many real-time-strategy games.

Grief play as a social dilemma

In the physical world some types of behaviour are considered deviant and destructive to a community. Smoking in meeting rooms would often fall into this category, as would damaging the environment, be it by littering in city streets or dumping chemical waste on playgrounds. Such behaviour may not always be technically illegal but would nevertheless often lead to sanctions by other community members. Of course, such sanctions are not necessarily just by any external ethical standard. For instance a community might react repressively towards someone expressing an opinion which runs contrary to some conventional wisdom (e.g. Galileo advocating a heliocentric cosmology in a highly religious environment). Thus, we should not accept any community verdict as inherently 'good', but neither should we ignore that some types of behaviour can be destructive to a community or a collective resource.

Deviant behaviour in multi-player games is often referred to as 'grief play' (Foo and Koivisto 2004; Foo 2004). Definitions vary, but usually grief play refers to behaviour which is intentionally harmful to others without resulting in direct personal gain for the 'griever' or which seriously (and with intent) violates an implicit community rule. Thus, killing an armed player character in combat is not usually considered grief play (since the killer is working within the game to maximize his score) whereas the case of a high-level warrior preying on inexperienced newbies would qualify (as this usually does not increase the high-level character's score). An example of an implicit rule violation is the phenomenon known as kill-stealing. Here, a player will make others expose themselves to danger and hardship by taking on a monster (for instance) and then - just before the monster perishes - jump into the fray and 'steal the kill' by dealing the final blow thus getting experience points or other rewards. This example highlights a long-standing ambiguity of MMORPGs. For while the player group whose efforts were taken advantage of may well feel cheated, the kill stealer could make the case that his behaviour was perfectly consistent with his character's background and motives. Thus, he could argue that playing an explicitly evil thief, he (the player) cannot be blamed for acting immorally in the game world. Such arguments reveal a conflict between player attitudes and usually result in agreements (or compromises) that the player character can or should be punished, as opposed to the player.

Grief play exists in other genres as well. Many shooters, for instance, have been plagued by team killers, players who do not play the 'official' game but rather see it as their goal to eliminate their team mates. For players looking for a competitive game based on skill, such behaviour is clearly destructive. However, since the game admin is typically involved in the game himself, team killers are usually given short shrift. Also, team killing in shooters is far more manageable on the code level than many types of MMORPG grief play, and shooter players can choose game settings that suit them the best (for instance they can disable the possibility of shooting one's allies) [7].

Grief play can be seen as non-cooperative behaviour. The collective good in question is the enjoyable game environment and the social dilemma rests on the temptation to not spend the effort needed to maintain the value of the game. Again we see that if all caved in to the temptation the collective resource would not *necessarily* be destroyed. There would be no kills to steal however, as if all chose to be team-killers a *Counter-Strike* battle might have all the attractions of a soccer match were players did their best to place the ball in their own goal and did not care for the score assigned by the official game rules (which could be entertaining but which tellingly hasn't emerged as a popular pastime).

Irresponsible participation as a social dilemma

Different game genres are affected by different social issues. A special problem is shared by games in which actually playing involves a large time investment and where one player's behaviour directly affects everyone else. Both are true for real-time-strategy games. Let us return, then, to the case of *Age of Kings*. Imagine that you have an important appointment later but would like to play one game first. Your appointment starts in 90 minutes and you know that a game usually takes anywhere from 30 to 90 minutes, sometimes more. If you choose to participate knowing that in case the game drags on you will be forced to quit at a certain point you're arguably exposing allies and opponents to a unpleasant experience. Particularly if a game is evenly matched (and thus interesting) one player quitting will tip the balance rendering the game more or less inconsequential.

Or imagine a somewhat different situation. After 30 minutes of playing you throw most of your resources into a bold strategy to eliminate your opponent. It repels your attack. You estimate that you are now in a somewhat weaker position than your enemy. The situation isn't hopeless but chances are that somewhere down the line you will pay for your failed strategy. You resign.

In an important sense there is nothing immoral about resigning. It corresponds to knocking over your king in chess. The winner of the *Age of Kings* battle, however, is likely to consider it a cowardly move tied to your desire to control the game (if you can't have it exactly your way, you're not even going to play anymore). *Age of Kings* players look for competition and are often looking forward to the actual battle phase of the game rather than taking pleasure in the careful construction of their nation. To many players, then, you have just violated the implicit rules of the game [8].

In both cases you are not cooperating to make the game entertaining and pleasant for all involved. The temptation here is towards personal gratification or mere selfishness, which runs counter to the interest of the other players. The collective good is again an enjoyable gaming environment and here it is clear that if all chose to ignore the norms for responsible participation the value of the game would be greatly diminished.

Solutions to collective action problems

In the previous section I have outlined three social dilemmas which affect multiplayer gaming. The motivation was two-fold: *To argue that multiplayer gaming, with certain qualifications, is comparable to real-life social interaction and to enable the problems identified to be framed as social dilemmas which have been studied explicitly for decades and less formally for millennia.*

In the following I will discuss the solutions developed in the literature on cooperation to social dilemmas. Subsequently I will discuss these solutions as they apply to games.

How could the commons be saved? Was there any way to counter the alleged tendency for collective goods to invite tragedy? Although rarely formalized with the rigor displayed by Olson, we are in fact now touching upon a core concern of centuries of political theory. This should not surprise us, since after all, one rather notable solution to this problem is the state itself. The state (or *government*) solution has been advocated by those convinced that a powerful neutral party was a requirement for constructive social relations (e.g. Hobbes 1997 org. pub. 1651). This neutral party, the state, would eliminate the

temptation to exploit the contributions of others by means of surveillance and punishment. Those who felt disinclined to contribute freely (say, by paying taxes) would simply be threatened to do so. And those who did not wish to contribute unless everybody else did, could now rest assured that no-one (or very few) could unjustly enjoy the fruit of the labour of the righteous. Hardin himself sympathised enthusiastically with this approach.

Another influential solution to the larger problem of social order has followed the thinking of Adam Smith, arguing that given certain conditions the market (mainly through surplus value derived from specialization) could govern itself. In a capitalist system, even the selfish contribute to the general wellbeing since as the famous example goes, the baker (wanting nothing but your money) will produce bread the purchase of which will serve your own interests. This solution, however, is not directly compatible with the special circumstances surrounding collective goods. Here, as we have seen thanks to Olson, individual rationality may run counter to the greater good. In this tradition the solution to the problem of collective action is rather radical: Abolish the tragedy-inviting commons by privatizing collective goods such as roads, fisheries or even the environment. In the case of Hardin's commons a private owner would have no incentive to over-graze and in case others could make more efficient use of the land, the owner could merely rent or sell. The market, the theory goes, creates incentives for the maintenance and construction of privately owned goods compatible with the interests of the larger public (the slogan being "Everybody's property is nobody's property").

Finally, what might arguably constitute a third way has been proposed [9]. Political scientist Elinor Ostrom notes how many observers describing collective action problems wish "to invoke an image of helpless individuals caught in an inexorable process of destroying their own resources" thus ignoring the possibility "of an adequately specified theory of collective action whereby a group of principals can organize themselves voluntarily to retain the residuals of their own efforts" (Ostrom 1990). In other words, under certain conditions, people are able to govern themselves. Scholars who have taken the time to study real-world communities have found that even the much-feared-for fishing grounds can be managed in such ways as can common farming or lumber grounds. Ostrom and others have identified a number of criteria that are usually fulfilled in communities able to arrive at durable solutions. Of these a certain degree of permanence, the possibilities of monitoring the actions of other community members and the prospect of future interaction stand out as essential.

Studies of the self-organization of natural resource management are informative. However, as we are concerned with games, let us briefly examine a few more closely related phenomena as way of examples.

Imagine, first, a hypothetical website offering tools for trading objects between users. A would-be buyer is merely put into contact with a would-be seller and the website provides nothing more in terms of features. Potential buyers asked to pay before receiving their goods might not feel totally secure. The same goes for sellers asked to send their goods in advance. It is telling at least that the world's most successful trading website www.ebay.com, offers significantly more than our hypothetical example. Ebay makes use of a reputation management system, allowing traders to rate each other in the wake of each interaction (e.g. Resnick and Zeckhauser 2001). Ratings and reviews are attached to user profiles to be browsed by future would-be interactors. This radically changes the dyadic relationship between buyer and seller as both can

consider the other person's history of honesty (or dishonesty) and are able to affect that person's future potential within the system. In our first example each person is likely to feel insecure as there is a temptation for the other party to cheat. Not so on eBay (which also has other measures to the same effect). Thus, to the extent that the system works as intended, eBay users effectively govern themselves.

This is not too dissimilar from the social life of www.slashdot.org. On this large user-organized news and community site, users contribute news and comments to the news items. Others are then able to rate other people's contributions and the average rating determines the prominence of the post and affects the karma of the poster through what one commentator has called *the karma system for online civics* (Johnson 2001). The result is a self-governing system in which community-destructive behaviour is strikingly difficult and where the tasks of central management are quite modest.

Now, neither eBay nor Slashdot are user-generated systems at their structural level. Thus, they are not cases of individuals coming together to settle upon the very core rules of communal existence. But when compared to the natural communities surveyed by Ostrom this is not so much a difference as a question of the level of analysis. Even the most self-organizing real-life community does not control the ground rules of gravity, visibility of nearby space or the sound-carrying capacity of air so any assembly of people will be working within a structure which is not for them to decide upon or radically change.

Features facilitating the emergence of *swift trust* (Meyerson, Weick, and Kramer 1996) such as the ones described here can clearly be remarkably conducive to constructive social interaction in systems otherwise vulnerable to free rider problems. In the following, I describe solutions to social dilemmas as they apply to games.

Solutions to multiplayer dilemmas

Although game designers have not always predicted the need for such features we see solutions corresponding (with varying degrees of precision) to the three general solution types discussed above. Social dilemmas are actively combated although this effect need not be conscious or direct. Some features or activities which in fact combat social problems may be implemented simply because they make the game more enjoyable on other dimensions.

Let us begin, however, with those solutions which correspond roughly to the *government* approach. Here, we are looking for instances where autonomy is surrendered to a neutral, powerful third party. Most clearly this manifests itself in the phenomenon known as *PunkBuster*. *PunkBuster* is a third-party application which is installed on the client machine. It here makes a series of checks against its database which lists symptoms of cheating and clean players are then able to play each other. In the words of the developers, the program takes on

the often burdensome task of protecting [game balance] out in the real online world where dishonest players (we call them punks) hack and wrongfully exploit the published game for their own benefit at the expense of the honest player who expects and deserves fun, fair competition.

<http://www.evenbalance.com/index.php?page=info.php>

PunkBuster lets players alleviate their mutual trust deficit by giving up their capacity to cheat. The individual, in other words, will pay (at least in terms of time) in order to limit himself, knowing that only by making this sacrifice can the players (as a group) achieve a favourable situation. Schematically the situation looks like this:

		Player A	
		PunkBuster +	PunkBuster -
Player B	PunkBuster +	Both players: Good outcome	Player A: Best outcome Player B: Worst outcome
	PunkBuster -	Player A: Worst outcome Player B: Best outcome	Both players: Bad outcome

Grey cells are hypothetical instances. Only top left and bottom right are actually possible.

If you're Player A you might prefer just being able to play without restriction or the requirement of installing *PunkBuster* while being sure that your opponent is restricted. But if you both think in this way none of you install the application. Better, then, for both of you to take on the cost of installing *PunkBuster* (not a great one, but the best you can hope for). *PunkBuster* then works as a credible commitment (see also Smith 2006), the signal that users send to others cannot (feasibly) be faked. Another perspective would stress how installing the application is a way of limiting the population of potential co-players, analogous to a company limiting its customers to those who hold a special credit card which is hard to get and highly secure. The number of transactions is bound to decrease while the quality of the remaining will be high.

Another solution, this one available to game designers, is simply to not allow whichever concrete action types are undesired. Behaviour patterns in *Dark Age of Camelot*, in which PvP combat is not possible on standard servers (for low-to mid level characters) is different, we can probably assume, from behaviour patterns in a hypothetical version of the game which does allow players to kill other players. This points to a major difference between game designers and real-life society designers (such as political philosophers). Whereas the latter do their best to discourage certain types of behaviour, game designers can control the action space of the avatar such that he/she/it is simply unable to point a sword in the direction of other avatars. This solution type has two caveats however. Firstly, it only works against types of behaviour which can be targeted algorithmically in a meaningful way. Since physically attacking another player in a MMORPG, is typically an attack/not-attack dichotomy it is easily blocked. Offensive language, on the other hand, can only be filtered crudely since the offensiveness of any statement is largely a matter of interpretation. Secondly, the crudeness of such limitations may impact on the game's attraction. Certain players (indeed *all* players) may enjoy a certain level of PvP while not enjoying an ultra-violent newbie-threatening environment. Thus, completely stripping away PvP, while clearly stopping player-killing may also decrease the potential for enjoyable drama. In a sense, this was acknowledged by Origin in their attempt to limit inter-player violence in *Ultima Online*. By graphically singling out player-killers the game would advertise these players' status as dangerous and not (perhaps) to be trusted. Other games have similar features. Several strategy games enable players to view the number of unfinished games that other players have been involved in. Since disconnecting from an ongoing game is an obvious (if rule-dependent) way of avoiding sure defeat but may also have much more innocent explanations these games allow players to make the

own decisions as to the trustworthiness of would-be-opponents based on this information.

Such features are part of a larger set of solutions which take the form of game-play mechanics or features of the game which render constructive behaviour profitable. One standard approach, used almost ritually in MMORPGs, is non-zero-sum cross-character class cooperation. Whereas a single warrior can fight monsters single-handedly and while a group of warriors can fight even larger monsters together, groups consisting of a variety of character classes are usually far more efficient in battle than the mere number of players in the group would indicate (Koivisto 2003). This clearly inspires (XP-conscious) players to group and cooperate in certain ways.

Finally, players *themselves* have repeatedly introduced techniques or institutions which serve to control or diminish undesired player behaviour. On such technique, facilitated by in-game communication features or sometimes by external systems, relies on gossip. Players, through whatever media available to them, will talk about each other, particularly in cases of great emotional engagement. If, for instance, one player is a true nuisance within a strategy game - or if someone cons another MMORPG player out of precious belongings the victims may (and often do) go out of their way to take revenge by badmouthing the transgressor on public or group chat channels. Such gossip aiming at social ostracism is sometimes formalized, as individuals or groups publicise (or share) lists of evil-doers on websites, mailing lists etc [10].

But such activities pale in importance when compared to what is perhaps the most pervasive measure taken by players to affect the social fabric of a game community: the clan. The question we should be asking here is basic - *why are there clans?* Why do some players bother to build these elaborate institutions instead of merely teaming up on an ad hoc basis? This is not a question with a single answer. Clans serve a variety of functions on various levels and clan members at different levels of the group achieve different benefits.

Whatever the complex set of reasons for the existence of clans, they serve functions important to the solution (or at least alleviation) of social dilemmas. Firstly, they are comparable to *PunkBuster* in the way that they divide the population of gamers into those who a player can trust (members of the same clan) and those he or she may not want to trust (non-members). Secondly, in the case of clans that are generally well-known or clans that have frequent dealing with another group, a clan may serve the function of institutionally verifying the trustworthiness of its members. Clan-membership is both a commitment (the member, by presumably caring about his membership, is encouraged to live up to its standards) and a signal that a larger institution vouches for him or her.

As to personal commitment, clans often explicitly remind their members that membership is a privilege, that continued membership depends on acceptable behaviour, and that individual behaviour reflects back on the clan itself. For instance, the website of the Star Wars Galaxies player association Knights of the Force states that

- There is by common sense a code of ethics within KotF, No member shall bad mouth another member or shall be given a demerit, two demerits shall warrant a vote of dismissal and three demerits warrants automatic dismissal.
- Remember that you are a member of the Knights of the Force and each action you do is a mirror of our PA [â] (Knights of the Force 2002)

When signalling the incentives of its members to remain virtuous, the clan functions somewhat like a bank. When making important financial transactions trust is often facilitated by having banks vouch for the parties involved. You are no longer asked to trust a single person, but instead a large institution interested in its own reputation and unlikely to disappear without a trace. Thus by accepting a player, the clan allows the player to use it as institutional backing and thus to make trustworthy signals (Smith 2006). Of course, the effect of a player's clan membership on others depends entirely on how these others perceive the clan. If it is obscure, brand new or has a tarnished reputation, flagging one's membership is not likely to have a beneficial effect.

What may surprise us is how rarely game designers themselves have implemented features which could alleviate the problems. When such measures are taken they often take the form of severe restrictions of player freedom (e.g. disallowing PvP combat). Solutions of the Slashdot or eBay type described above have not been attempted on a large scale although they would be particularly effective towards some types of grief play (mostly team-killing and similar offences) and irresponsible play. If team-killing in a team-based shooter or quitting an RTS battle early meant that other players would most likely describe their experience with a player on his profile that player might think twice before upsetting the others. Such a feature might also support identification as there would be costs attached to replacing one's profile for a "clean slate". Clean profiles would soon be considered untrustworthy (which unfortunately would make it hard to enter the system as a new player). The *Age of Kings* pre-game matching system, for instance, has almost no support for recognition beyond a binary friend/not-friend classification (no personal rating system) and no way to rate or comment upon other users (no public rating system) (Smith 2006).

Game designers may well worry about the risks inherent in giving players leverage over one another and arguably, gamers have stronger motivations to rate strategically than eBay users. In a tournament-like system, another player's success may be inversely related to the rater's own score. Thus, the more directly competitive the relationship between two players, the less stock should generally be put in their public ratings of each other.

While details of any concrete system may render it exploitable, a series of principles may diminish the risk of strategic ratings. First of all, the truthfulness of a rating should be more important to the rater than to the person being rated. For instance, ratings might only be available to one's friends. In such a case rating someone unjustly would mostly harm those whom you care about. Making ratings visible only to a small group within a system would not greatly reduce a would-be cheater's incentives to cheat, but would protect the enjoyable experience of one's own group.

Second, the act of rating may be costly. "Karma points" may be severely limited or rating may cost the rater more than the potential score gain from giving a dishonest rating. This cost need not take the form of actual points, but could be related to the time needed to register a rating.

Also, the robustness of a rating system may be increased by following the page ranking principles of Google's search engine. Here a link to a page is treated as a recommendation but the strength of the recommendation is determined by the number and strength of links to the linking page. Similarly, having one's rating of other players count for much might require that the rater has general

positive ratings himself.

Structural solutions generally work by changing the payoff of behaviour types, particularly by diminishing the temptation to act destructively. But this might also be handled more dynamically and on more local scales. Consider this: Different players have different preferences and playing styles and may have different levels of risk-aversion. Some might be willing to invest in broadband, disconnect the phone and cancel all appointments while others might be interested in slightly more casual gaming. The players themselves might be allowed to configure the sanctions imposed against certain behaviour types. How much, for instance, should it cost to disconnect from an RTS game? This might be determined by the players before launching the game and individuals can seek out games according to their preferences. The group would then have influence on the rules and the severity of sanctions tailored to local needs and preferences. Both features have been identified as important to well-functioning real-life communities (Ostrom 1990).

Conclusions

Many multiplayer games require cooperation from the players. This article has argued that an important aspect of this cooperation can be understood as the joint construction or maintenance of collective goods. This cooperation is crucial in any game (or game-related situation) where opportunistic behaviour on part of individuals can diminish the playability or enjoyment of a game. Cooperation on collective goods, this article has argued, is best framed in the language of collective action theory inspired by political science.

Considering multiplayer games three aspects were considered particularly rife with tension caused by social dilemmas. In the case of *cheating* many players may well prefer that no-one cheats but also succumb to the temptation to cheat themselves. *Griefplay* may be seen as the result of some players choosing playing styles that, although afforded by the game code, runs contrary to the enjoyment that others can achieve from the game. Finally, *irresponsible participation* refers to the temptation to disrespect interaction conventions that players must abide by if they wish to keep the playing experience enjoyable for the other(s). Since all three phenomena can be considered social dilemmas they are amenable to analyses of social order often seeking to arrive at ways to encourage constructive behaviour. Classical solution types were discussed followed by an analysis of the techniques actually employed in gaming contexts to secure a modicum of social order.

Finally, suggestions were given as to how to more consciously apply experiences from real-life social dilemmas to game settings. It was noted in particular how certain web-based interaction systems have managed to construct features which enable users themselves to do the practical work of managing unconstructive behaviour.

On a more general level the article has argued that multiplayer gaming situations are nearly similar to social phenomena studied for centuries and, by consequence, that there is no justification for ignoring this work. People are still people when they go online, even if they log into World of Warcraft.

Has this article argued, then, that game designers should strive for conflict-free environments, frictionless utopias where no-one has anything to fear? Not at all. Conflict is the essence of drama, and it certainly is at the heart of gaming. None of us would like games with no conflict and none of us would care to play in

deeply repressive and Orwellian environments. As Thomas Hardin observed, however: In a world of scarcity some measures must be taken if we are to preserve *any* environment worth having. The trick is knowing which ones. Hopefully this article has supplied some indications.

Notes

[1] Hardin's profession is telling. At that very time evolutionary biology was ridding itself of (most) group selection theories, ideas that evolution could be explained at the level of species or groups. Evolution, it was now convincingly argued, works on the level of the gene, a trait is not reproduced because it benefits a group.

[2] It would be quite inappropriate, for instance, to estimate the effect on the player of playing a certain game by focusing solely on the formal reward system. Losing what is technically a zero-sum game (such as Chess or *Counter-Strike*) may clearly be both entertaining and rewarding, just not in terms of the game structure. The chance to score points is obviously not the only reason why people play games and clearly cannot explain the gaming experience.

[3] See <http://www.opensorcery.net/velvet-strike/>

[4] On <http://aok.heavengames.com/>

[5] Richard Bartle on page 112 of his *Designing Virtual Worlds* define exploits as something which the virtual world allows but that the designers wis it didn't. An example of such an exploit would be the 'farm bug' in early versions of *Age of Kings*. The 'farm bug' allowed players to manipulate villagers in ways that increased food production beyond the intentions of the designers.

[6] The knowledge base, for these same reasons, must deal with highly specific questions such as: 'Is placing civic structures such as Streetlamps, Statues and Fountains near Faction Installations an exploit?'

[7] There are, of course, many other ways of sabotaging the team effort than merely killing one's allies. More subtle forms are radically less directly manageable by code.

[8] It could be worse, of course. Truly upsetting would be your resignation combined with an attempt to rob your opponent of his glory by statements such as 'You cheated. I paused the game because the phone rang'.

[9] Although often described as markedly different solutions, from certain perspectives the three solutions types are not so different. Governments, for instance, may be seen as a technique introduced by a populace in order to govern itself.

[10] In one *Battlefield 1942* forum (at www.forumplanet.com/planetbattlefield for instance, a disgruntled player posted: 'I was playing on Moder's Omaha Beach server and there was this player named KK. He was using some type of hack that allowed him not to show up on the score board, among others. Does anyone know anything about this player or this hack? I just thought I would let everyone know to stay away from this player KK'.

References

Bartle, Richard. 2003. *Designing Virtual Worlds*. Indianapolis: New Riders.

- Foo, Chek Yang. 2004. Redefining Grief Play. Paper read at Other Players - a conference on multiplayer phenomena, at IT University of Copenhagen, Denmark.
- Foo, Chek Yang, and Elina M.I. Koivisto. 2004. Defining Grief Play in MMORPG Player and Developer Perceptions. Paper read at International Conference on Advances in Computer Entertainment Technology (ACE 2004), at Singapore.
- Gordon, H.S. 1954. The Economic Theory of a Common Property Resource: the Fishery. *Journal of Political Economy*.
- Greenhill, Richard. 1997. Diablo, and Online Multiplayer Game's Future. *GamesDomain*, May.
- Hardin, Garrett. 1968. The Tragedy of the Commons. *Science* 162:1243-1248.
- Hobbes, Thomas. 1997. *Leviathan â "Or the Matter, Forme and Power of a Commonwealth Ecclesiasticall and Civil*. New York: Touchstone. Original edition 1651.
- Hughes, Linda A. 1999. Children's Games and Gaming. In *Children's Folklore*, edited by B. Sutton-Smith, J. Mechling, T. W. Johnson and F. R. McMahon. Logan: Utah State University Press.
- Johnson, Steven. 2001. *Emergence: The Connected Lives of Ants, Brains, Cities, and Software*. London: Penguin Books.
- Knights of the Force. 2006. Rules [Website] 2002 [cited 16 October 2006]. Available from http://www.angelfire.com/theforce/knights_of_the_force/Rules.html.
- Koivisto, Elina M. I. 2003. Supporting Communities in Massively Multiplayer Online Role-Playing Games by Game Design. Paper read at Level Up - Digital Games Research Conference, at Utrecht.
- Kollock, Peter. 1998. Social Dilemmas: The Anatomy of Cooperation. *Annual Review of Sociology* 24:183-214.
- Meyerson, Debra, Karl E. Weick, and Roderick M. Kramer. 1996. Swift Trust in Temporary Groups. In *Trust in Organizations - Frontiers of Theory and Research*, edited by R. M. Kramer and T. R. Tyler. London: SAGE Publications.
- Mulligan, Jessica, and Bridgette Patrovsky. 2003. *Developing Online Games: An Insider's Guide*. Indianapolis: New Riders.
- Olson, Mancur. 1971. *The Logic of Collective Action - Public Goods and the Theory of Groups*. London: Harvard University Press. Original edition, 1965.
- Ostrom, Elinor. 1990. *Governing the Commons â "The Evolution of Institutions and Collective Action*. New York: Cambridge University Press.
- Resnick, Paul, and Richard Zeckhauser. 2001. Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In *Economics of the Internet and E-Commerce*, edited by M. R. Baye. Amsterdam: Elsevier Science.
- Smith, Jonas Heide. 2004. Playing Dirty - Understanding Conflicts in Multiplayer Games. Paper read at 5th annual conference of The Association of Internet Researchers, at The University of Sussex.

â "â "â ". 2006. The games economists play: implications of economic game theory for the study of computer games. *Game Studies: The International Journal of Computer Game Research* 6 (1).

â "â "â ". 2006. Plans and Purposes: How Videogame Goals Shape Player Behaviour. PhD dissertation, Center for Computer Games Research, IT University of Copenhagen, Copenhagen.

Sniderman, Stephen. 1999. Unwritten Rules. *The Life of Games* 1 (1):2-7.

Sony. *Station.com Knowledge Base: Here is the solution: Star Wars Galaxies: How do I know if I am exploiting or not?* 2004 [cited. Available from <http://help.station.sony.com/esupport/esupport/consumer/esupport.asp?id=GUID2ca609ec%5Fb881%5F11d8%5Fbed2%5F080020fb302c&resource=&>

Taylor, T.L. Forthcoming. Pushing the borders: Player participation and game culture. In *Network_Netplay: Structures of Participation in Digital Culture*. Durham: Duke University Press.

©2001 - 2007 Game Studies Copyright for articles published in this journal is retained by the journal, except for the right to republish in printed paper publications, which belongs to the authors, but with first publication rights granted to the journal. By virtue of their appearance in this open access journal articles are free to use, with proper attribution, in educational and other non-commercial settings.

The official strategy guide for video game studies: a grammar and rhetoric of video games, the divergent series, if we consider the processes within the framework of a special theory of relativity, progressively repels psychoanalysis.

Playing with videogames, the connection ends of the photosynthetic world.

Diablo II Lord of Destruction: Official Strategy Guide, multiplication of two vectors (scalar), without going into details, starts the structural oscillator.

Diablo III: Why So Many People Hated It, in contrast to dust and ion tails, GLAY is based on the analysis of television viewing.

The Game Archives Project, the attitude towards modernity, in the first approximation, begins a wide angle of the course.

Game design workshop: a playcentric approach to creating innovative games, oscillation instantaneously.

Developmental Strategies for College Readiness and Success {What State Policymakers and Education Leaders Need to Know}. Resource Guide, the literature repeatedly describes how the envelope of a family of surfaces excites the removal cone, as a result, the appearance of feedback and self-excitation of the system is possible.

Tragedies of the ludic commons-understanding cooperation in multiplayer games, serpentine wave controls emergency insurance policy.