

Viral vitriol: Predictors and contagion of online toxicity in World of Tanks

Cuihua Shen^{a,*}, Qiusi Sun^a, Taeyoung Kim^a, Grace Wolff^a, Rabindra Ratan^b, Dmitri Williams^c

^a University of California, Davis, USA

^b Michigan State University, USA

^c University of Southern California, USA

ARTICLE INFO

Keywords:

Toxicity
MMO
Online games
Contagion
Social network

ABSTRACT

Toxic behaviors are pervasive in online games and can be harmful to building a positive online environment. Guided by the social identity model of deindividuation, this study represents one of the first efforts to examine the antecedents of toxicity in team-based online games using longitudinal behavioral data. It fills two important gaps in existing research, by 1) exploring non-verbal and behavioral dimensions of toxicity, and 2) examining team-level in addition to individual-level predictors. Employing a large-scale behavioral dataset from the popular game *World of Tanks*, we found that, in general, experienced and skillful players are more likely to commit toxic behaviors. Teams that are losing, or have a high internal skill disparity among their members tend to breed toxicity. In addition, this study provides empirical evidence that toxicity is contagious among players, especially toxic behaviors in one's own teams and in clan battles.

Playing video games has become an essential part of digital life for hundreds of millions of people all over the world. 65% of Americans play video games daily and the total US consumer spending on video games reached \$43 billion in 2018 (ESA, 2019). Particularly popular are massively multiplayer online games (MMOs) where players can interact with each other through a wide range of actions and relationships, such as friendship, textual/verbal messaging, trading, mentoring, teamwork, and competition.

While opportunities to connect with other players contribute to the enjoyment of games, they also give rise to undesirable behaviors such as trolling, griefing, cyberbullying, flaming, harassment, and cheating, partially due to the anonymous, disinhibiting environment where bad behaviors have minimal consequences (Suler, 2004). Such behaviors reduce the enjoyment of gaming and fracture the community. Even worse, toxicity may have far-reaching impacts beyond games as they shape what the gamer population, especially younger generations, perceive as “normal” conduct, potentially contaminating other spheres, online or off.

The causes and remedies of online toxicity have become important topics for academic and industry research. In non-game spaces, a large stream of research focuses on verbal or textual manifestations of toxicity such as online comments, mostly in political communication contexts (e.g., Gervais, 2017; Rains et al., 2017). Data are often collected through

scraping online texts and self-administered surveys. The majority of game toxicity research has also focused on verbal forms of toxicity (e.g., Balci & Salah, 2015), and/or relied on self-reports such as surveys and interviews (e.g., Fox & Tang, 2017; Kou & Nardi, 2013). These descriptive studies provide rich details but may be limited in their accuracy, scale, generalizability, and examination of the situational factors leading to toxicity. One notable exception is the research conducted by Riot Games, the parent company of the popular game *League of Legends*, which experimentally tested various user interface design features to reduce toxicity within *League of Legends*, beyond mere observation of the type and extent of toxicity occurring in the game (Blackburn & Kwak, 2014; Maher, 2016).

Taking advantage of large-scale behavioral data from a popular online game, *World of Tanks* (WoT), the current study aims to explore the antecedents of toxicity at both the individual and collective levels in a team-based, massively multiplayer game context. This study focuses entirely on non-verbal and behavioral dimensions of toxicity, using highly granular and longitudinal data collected unobtrusively from the game servers. This approach enables a close examination of associations between gameplay patterns and toxic behaviors in online games. It recognizes that anti-social behaviors are not solely determined by individual traits. Rather, individuals may become toxic as a result of situational factors such as team dynamics and their exposure to other

* Corresponding author. Department of Communication, UC Davis, One Shields Ave, Davis, CA, 95616, USA.

E-mail addresses: cuishen@ucdavis.edu (C. Shen), qssun@ucdavis.edu (Q. Sun), tytkim@ucdavis.edu (T. Kim), ghwolff@ucdavis.edu (G. Wolff), rar@msu.edu (R. Ratan), dmitri.williams@usc.edu (D. Williams).

<https://doi.org/10.1016/j.chb.2020.106343>

Received 13 August 2019; Received in revised form 4 March 2020; Accepted 11 March 2020

Available online 12 March 2020

0747-5632/© 2020 Elsevier Ltd. All rights reserved.

perpetrators.

1. Theory and hypotheses

1.1. Online toxicity

Toxicity is an ambiguous term. Its designation ranges from “uncivil discourse,” “incivility,” “anti-social behaviors,” “verbal aggression,” and “griefing,” to general “degrading comments”. Uncivil discourse is defined as “claims that are deliberately disrespectful and insulting, or those presented in a hyperbolic nature” (Gervais, 2015, p. 169). While incivility has been operationalized as comments that include profanity, yelling by use of capital letters, and name-calling (Coe, Kenski, & Rains, 2014; Papacharissi, 2004), it can also include “political flaming”, which is generally defined as aggressive communication around politics in online spaces (Graham & Wright, 2015) that tend to attack another’s self-concept rather than the argument or opinion presented (Infante & Wigley, 1986). Researchers have also attempted to define sub-dimensions of incivility as related to the varying definitions of toxicity. Incivility generally consists of insulting language (including but not limited to mean-spirited, disparaging, and vulgar words), overloading of emotional language, and lying or exaggeration (Coe et al., 2014; Gervais, 2015). In most cases, the concept is defined exclusively in the context of online discussion. Most of this literature is focused at the individual level, examining why certain people in certain contexts may engage in toxic behaviors.

However, once we examine group-level predictors, there are more appropriate theoretical frameworks, especially when the concept of “team” becomes more or less salient to the player. The social identity model of deindividuation effects (SIDE) suggests that when people enter an online world, they feel more anonymous and detached from actual life. In turn, they become more likely to conform to group norms, including antisocial ones (Chesney et al., 2009). In a cooperative-competitive game such as *World of Tanks*, gameplay occurs exclusively between pairs of opposing teams. Some teams (ad hoc or pick-up teams) consist of strangers who have no prior history of interaction, while others are based on preexisting social networks such as clans and guilds. These significant variances in what team means, and when, allow for tests of SIDE effects in different contexts within one game title. In general, the process of deindividuation occurs when communicators’ awareness of own and others’ personal identity is decreased (Matheson & Zanna, 1988). Its effects are influenced by the type of computer-mediated communication and the amount of social presence or social cues available through the medium (Kiesler, Siegel, & McGuire, 1984; Short, Williams, & Christie, 1976). This lack of social cues or decreased social presence may potentially give rise to anti-social behavior. Correspondingly, the introduction of those cues may decrease those same behaviors. To date, most research on online toxicity has focused on analyzing patterns from textual data from online discussions (Coe et al., 2014; Gervais, 2015; Thorson, Vraga, & Ekdale, 2010; Vargo & Hopp, 2017). Although such textual analyses provide valuable insights into toxicity, data on online users’ behaviors offers additional insights by capturing what people do, rather than what they say.

1.2. Toxicity in massively multiplayer online games

Massively Multiplayer Online (MMO) games are persistent and immersive online worlds where millions of users engage in gaming activities and interact with other players, often through avatars (Williams, Yee, & Caplan, 2008). MMOs enable players to cooperate and/or compete strategically on a large scale with other players (player-versus-player or “PvP”) or non-playable characters (player-versus-environment or “PvE”) with several affordances that enable interpersonal communication, relationships, and virtual communities (Shen, 2014). Of the most frequent gamers, 56% play multiplayer games, spending on average 7 h per week with friends and strangers

online (ESA, 2019).

Similar to other online spaces, MMOs are not free of toxicity. Due to anonymity, players form impressions of others based on minimal social cues often surmised from voice/text team chat and/or visual representations, such as avatar gender, race, weight, and clothing choices (Lea, Spears, & de Groot, 2001). The over-attribution of these minimal cues produces exaggerated and stereotyped positive or negative impressions of communication partners (Hancock & Dunham, 2001). These exaggerated evaluations of others from minimal cues can foster identification with group norms, which can be both pro- or anti-social. The present research focuses on the latter, and thus the potential for deindividuation to lead to further alienation, targeting, and the use of toxic behaviors towards others. Of particular importance is the question of whether and when group identification will or will not be activated. Out-group hostility and in-group liking are well-studied concepts (e.g., de Zavala, 2011), but in dynamic online environments the “group” can be more or less salient as the context around the players shifts. Players playing on a temporary team compared to a more long-term one will likely place a different emphasis on the cues present. But before acknowledging the team- and cue-based predictors, we examine the existing literature on gaming toxicity, and on individual-level differences.

1.3. Antecedents of toxicity

Studies of factors leading to online toxicity have mainly focused on text-based exchanges, including political forums and social network sites, relying on participants’ self-reported survey data or textual data. For example, frequency of leaving comments was negatively associated with the amount of incivility on a news website (Coe et al., 2014) and Twitter users with lower socioeconomic background or education were more likely to express incivility in their Tweets (Vargo & Hopp, 2017). Another study found that users tended to use more incivility when they were dealing with harsher responses from the community and that irrelevant posts were more likely to contain toxicity (Cheng et al., 2017).

1.4. Individual-level predictors in MMOs

Several studies have examined incivility within online multiplayer gaming environments and the variables that predict toxicity. A prevailing form of toxicity in games is bullying, defined as repeated unpleasant behavior that is intended to harm the victim (Olweus & Limber, 2010). Researchers have pinpointed that a power difference between the actor and target, regardless of whether it is real or perceived, makes the victim more vulnerable to bullying (Ballard & Welch, 2017). In online games, there are a variety of cues that suggest the status of one player to another, such as ranks and achievement levels of a user, the amount of time a user has invested in the game, or the amount of monetary investment a gamer has made. Empirical research confirms that power differences among players is associated with toxic behavior. A self-administered survey on MMO players found that gamers’ dominance and power (as manifested in game rank) was the most significant predictor for cyberbullying other gamers (Ballard & Welch, 2017). Following this line of research, our study attempts to explore whether players’ gameplay characteristics and their toxicity are related.

RQ: What kind of players are more likely to engage in toxic behaviors?

1.5. Team-level predictors in MMOs

The predictive power of individual traits on toxicity is rather limited, however. The studies conducted by Riot Games on a hugely popular title, *League of Legends*, show that the majority of toxicity comes from occasional offenses committed by normal players, rather than consistent offenders (Maher, 2016). Machine learning models also reveal that the opponent team’s success is the most important predictor of toxicity (Blackburn & Kwak, 2014). These studies support the importance of

team-level effects in triggering negative behavior.

At the team level, toxicity can arise along with the amount of stress and frustration players experience as a team. Studies have found that acute stress is negatively associated with the information processing capacity of the team, leading to suboptimal team performance (Ellis, 2006). One significant source of team stress and frustration comes from outside threat, defined as “an environmental event that has impending negative or harmful consequences for the entity” (Staw, Sandelands, & Dutton, 1981, p. 502). In competitive game environments, fighting against the enemy team places considerable pressure on a team to perform at a standard equal to or above the enemy team because losing the battle incurs a nontrivial loss of time, resources, and status for every team member (Johnson, Nacke, & Wyeth, 2015).

While matchmaking algorithms in MMOs strive to facilitate “fair” matches where both teams have an equal probability of victory (Shores et al., 2014), this does not always occur in practice. Some skill discrepancy between the two teams can still exist, giving one team an advantage over the other. In game systems with no matchmaking mechanisms, tensions are even more likely, as teams assembled and matched at random can face significant skill disadvantages. Ethnographic accounts show that toxicity typically surfaces when a team is losing (Kou & Nardi, 2013; Märtens et al., 2015, pp. 1–6), or when expectations for performance are not met (Johnson et al., 2015). This can lead to the shaming and blaming of one or more team members for perceived poor performance. Initially, minor offenses such as flaming could quickly be reciprocated and escalate into more intense toxic behaviors (Kou & Nardi, 2013). By contrast, teams with skill advantages are presumed to perform better and are more likely to have a relaxed mood, a positive climate and lower levels of toxicity (Neto, Yokoyama, & Becker, 2017). Therefore, we hypothesize:

H1. Toxic behaviors are negatively associated with a team’s skill advantage over the opponent team such that teams with a greater skill advantage display less toxicity.

While it may be frustrating for players when their own teams are not performing well against the enemy team, skill disparity among one’s own team may also increase the level of stress. By design, competitive team-based MMOs have an intricate division of labor among various character roles and classes. A successful team not only depends on the overall team skill level, but also the even distribution of such skills among team members, a requirement seldomly met in reality (Vella et al., 2016). Team imbalance in skills and experience may result in poor coordination and performance. For instance, gamers with better skills must put more effort to compensate for the deficit caused by the inexperienced team members. Skill and experience disparity may also cause misunderstandings among team members when they must act and react collaboratively, leading to conflict and hostility toward lower-status members.

Another potential explanatory mechanism is the black sheep effect, which asserts that groups expect members to contribute to the accomplishment of collective group goals (Pinto et al., 2010). When those expectations for normative behavior are met, in-group members upgrade their appraisals of that in-group member. When expectations for normative group behavior or performance are not met, in-group members will lower their appraisals of that in-group member (Marques, Abrams, & Serôdio, 2001). This exaggerated evaluation based on skill and perceived performance (or lack thereof) of another in-group teammate, as well as the discrepancy of skill amongst team members may lead to toxic reactions. This leads us to hypothesize:

H2. Toxic behaviors are positively associated with skill disparity among a member’s in-group teammates.

1.6. Team-based predictors of toxicity

Another important predictor of toxicity may stem from the nature of

the team. SIDE theory suggests that text-only communication is generally bereft of cues, and that without more information individuals do not identify others as individuals themselves (Walther, 2006). As will be demonstrated below, *World of Tanks* features both forms of group formation and communication—one mode with only text and another which encourages audio communication. These modes correspond to typical competitive game environments, which generally support two types of teams. First, ad hoc or pick-up teams are created by game algorithms with strangers who have little prior history of interaction. These players use only text-based communication and cues. Second, players may form teams based upon pre-existing social structures, such as friend lists, guilds or clans, which facilitate the organization and coordination of teams with players who are already connected to each other online and/or offline. These teams typically play with audio communication, and so have far more cues. These two team types differ fundamentally in ways that could influence the calculation of cost and benefits associated with in-game actions. Prior experimental research among gamers has shown that adding voice and its cues to a team can insulate players from negative social outcomes (Williams, Caplan, & Xiong, 2007).

There are of course other theoretical mechanisms to predict poor behavior that pre-date online group research. Classic game theory presents social dilemmas such as the “tragedy of the commons” or “prisoner’s dilemma” where rational individuals can benefit more if they defect, instead of cooperate, in group efforts (Axelrod & Hamilton, 1981). When it is guaranteed that two parties will never see each other again, selfish defection should be the unavoidable ultimate outcome. An empirical study shows that cooperative behavior among individuals is increased when future interaction is expected and frequency of contact is anticipated to increase (Heide & Miner, 1992). In other words, individuals who anticipate no future interaction with the same partner will more likely to “defect” rather than “cooperate”, eventually resulting in an all-defect population. This phenomenon can explain anti-social behaviors online, as there are limited social cues, little or no prior history of interaction and no expectation of future interactions. There is no “shadow of the future” (Axelrod & Hamilton, 1981), and so no social cost for defecting from good behavioral norms, such as being shunned by friends and guilds. If, by contrast, players anticipate future interactions, they are more likely to cooperate and behave prosocially (Van Lange, Klapwijk, & Van Munster, 2011). Therefore:

H3. Toxic behaviors are more likely to occur in battles with strangers than those with friends.

1.7. Contagion of toxicity

In online settings, the affordance of anonymity deindividuates players and reduces their accountability, increasing their susceptibility to being influenced to act anti-socially (Lea et al., 2001). Social contagion describes the process in which information, social norms, diseases, and collective behaviors propagate through contacts in social networks (Christakis & Fowler, 2009). Two contagion mechanisms, group norms and emotional contagion, are particularly relevant for the diffusion of toxic behaviors, which can be risky and controversial for the perpetrator.

The first mechanism is the establishment of group norms that anti-social behaviors are in fact acceptable. According to Social Cognitive Theory (SCT), individuals develop mental models about the rules that govern which behaviors are desirable and which are not by viewing others’ behaviors (Bandura, 1994). Behaviors that are shown as having positive consequences are more likely to be modelled by viewers, whereas behaviors with negative consequences are avoided (Bandura, 1994). This modeling can be problematic when anti-social online behaviors do not receive negative consequences and become increasingly normalized. Toxic behaviors are viewed as common occurrences, “a part of life in the world of modern multiplayer games” (Blackburn & Kwak,

2014, p. 877), with such behaviors considered amusing or annoying by players. In a survey study, most gamers did not have a problem with flaming, which was considered normative communication by gamers with a longer playing history compared to novices (Elliott, 2012). This normalization of toxicity is potentially exacerbated by the notion that virtual spaces are not real, in which interactions within them do not have a “real” impact (Fox & Tang, 2017) and thus subsequent in-game behaviors are inconsequential. When a player acts independently according to their own self-interest with abusive and toxic behaviors rather than cooperating for the good of the group, and others see and emulate this norm. Through repeated exposure and replication, the toxic norm is collectively reinforced, thereby harming overall gameplay and the community.

The second mechanism is emotional contagion, a well-established process in online social networks (Kramer, Guillory, & Hancock, 2014). The influence to act anti-socially has been shown to be contagious and spread, much like other instances of emotional contagion that infuse groups with more positive or negative moods which can influence cognition, behaviors, and attitudes (Barsade, 2002). For example, a study on trolling found that both a negative mood and exposure to troll posts significantly increased the probability of a user trolling (Cheng et al., 2017). Therefore we hypothesize:

H4. Toxic behaviors are contagious.

Studies have shown that toxic behaviors in games typically occur when a team is losing and one or more players shame and blame another player for perceived poor performance (Kou & Nardi, 2013; Märtens et al., 2015, pp. 1–6). Even if flaming occurs between only two players, it can upset everyone on the team. This inevitably leads to less cooperation and worse performance when the team is unmotivated and no longer in the mood to play (Kou & Nardi, 2013). While it may be normalized and not considered or viewed as disruptive, toxic behavior hurts the in-group team more than out-group enemy team.

Research suggests that incivility when targeting one’s in-group generates anger and that exposure to disagreeable incivility induces feelings of anger and aversion, reduction of satisfaction and increased incivility (Gervais, 2015, 2017). Neto et al. (2017) found that teammates of the toxic player are more affected by the antisocial behavior than the opponent team and that opponent groups are less vulnerable to contamination. This is in part due to the greater number of in-group communication channels than outgroup channels at the perpetrators’ disposal. Indeed, in *WoT*, players can only chat in real time with their own team. Perpetrators are more likely to diffuse anti-social behaviors through verbal and nonverbal communication channels within their own team rather than the opponents’ teams. Therefore.

H5. Toxic behavior is more likely to spread within the perpetrator’s own in-group team than to or from the opponent team.

Finally, as mentioned previously, the shadow of future interaction could potentially influence the likelihood of toxicity contagion (Van Lange et al., 2011). Playing with already known friends cooperatively leads to improved team and individual performance, possibly due to greater assistance and less instances of betrayal (Mason & Clauset, 2013). By contrast, teams made up of strangers who assume they will never play together again can exacerbate the detrimental consequences of anonymity and the normativity of toxicity (Shores et al., 2014). As previously explained by SIDE theory, when shrouded in anonymity and playing with strangers compared to members of a clan or friends, one may deindividuate and target another teammate based off minimal cues acquired in an interaction, thus becoming more likely to engage in disinhibited toxic behaviors. Therefore.

H6. Toxic behavior is more likely to spread in battles with strangers than those with friends.

2. Methods

2.1. *World of Tanks*

We situate our study of toxic behaviors in the game *World of Tanks* (*WoT*), classified as an MMO.¹ It can be characterized as a hard-core, team-based player-versus-player game in which individuals drive a tank in a bounded first-person environment, similar to a shooter game. Gameplay mainly revolves around matches between two teams. In each separate 15-min match (called a battle), players use a single tank, while players’ overall profiles reflect an accumulation of these matches. Tanks are of different tiers, ranging from one to ten, with increasing power and toughness. Tanks are also of different types (e.g., light or heavy), roughly analogous to the different roles seen in shooter or team-battle games. These tanks are unlocked through play or can be purchased with real money. Through a matching algorithm, matches pit players against another team by balancing tank tiers as well as the different tank types. A critical difference between *WoT* and similar games is that it does *not* balance teams by skill. Thus, the potential for imbalanced matches, and the frustration they create, is relatively high.

In the most common game format, called “random battle,” players are chosen by the game algorithm to be on one of two teams. Because of the automatic team assignment process, players do not have prior interactions, and do not expect future interactions with others they encounter in random battles. The random battle mode is analogous to “pick up” or ad hoc teams found in other MMOs. Analogous to other online games, players can join larger, more permanent groups called “clans,” complete with rosters, ranks, and the usual trappings found in MMOs and other team-based games (Williams et al., 2006). Clans therefore represent semi-permanent social structures which provide a stable backdrop for many in-game activities. More advanced players will join clan-based matches where the team sizes can be 7-, 10- or 15-players, and where the tank tiers are fixed at the higher levels. An in-game chat and voice system helps coordinate team actions, but is less common in random battles than in clan-based battles.

2.2. Data

Data were supplied by the game’s publisher, Wargaming.net, and consisted of behavioral data from the game’s server for players in North and South America from March 1 to March 31, 2018. The behavioral data included two tables. The first table contained time-stamped information on the more than 1.3 million battles (excluding battles for training purposes) fought during the 31-day period. It included the type of the battle, the length of the battle, the result of the battle (a team may win, lose or tie), the members of both teams, the skill level of each member at the time of battle, the type and level of tanks used in the battle, etc. The second table recorded time-stamped information on player-reported toxicity incidents. It included the ID of the reporter and the accused perpetrator, the type of toxicity incident, the ID of the team, and the battle where the toxic behavior occurred.

2.3. Measures

Toxicity. *WoT* provides two ways to report toxic behavior. The first and most common way is through an automatic reporting system during battles. Anyone can report others during a battle and specify the type of toxic behavior committed by the player. The server database automatically captures the battle identifier, the perpetrator’s ID, the reporting player’s ID, and the type of toxic behavior. The second way of reporting is through a customer service complaint after the battle is completed. Like the automatic system, the customer service ticket system records the reporter and perpetrator’s game ID. However, given that reports

¹ <https://www.mmorpg.com/world-of-tanks>.

Table 1
Descriptive statistics of major study variables.

	Mean	SD
Individual players		
Prior battle count	8925	11584.6
Prior skill	3035	2270.08
Battle count during observation	156.2	254.21
Average days in clan	166.8	402.13
Average tank level	5.51	2.27
Toxic (yes = 1)	0.23	
Reported others (yes = 1)	0.39	
Teams (size = 15)		
Average team skill	4227	1051.6
Team skill advantage over opponent team	0.09	617.55
Ingroup team skill disparity	1602.8	350.38
Battle time (seconds)	386	129.88

Table 2
Welch's *t*-test comparing toxic and non-toxic players.

	Toxic	Non-toxic	df	t	<i>p</i>
Prior battle count	16152.24	7588.97	78,389	137.34	***
Prior skill	4399.18	2898.43	99,926	155.7	***
Battle count during observation	377.84	104.73	72,909	193.15	***
Average tank level	6.79	5.42	139,210	170.77	***
Reported others (yes = 1)	0.73	0.28	99,581	218.19	***
Average days in clan	226.13	163.85	93,225	31.5	***

Note. ****p* < 0.001.

happen after battles are completed, the toxic behavior was not readily linked to any battle ID, so we could not identify the specific people and team involved, other than the reporter and the perpetrator. Being reported as toxic is costly to the player, as five reports will result in a suspension of the account.

When reporting a toxic incident through either automatic reporting or a customer service ticket, the reporter must identify the type of toxicity incident out of the four categories: unsportsmanlike conduct (41%), inaction/bot (41%), inappropriate behavior (14%), and offensive nickname (4%). Because some of these categories are rather vague (e.g., inappropriate behavior) and no explicit explanations were provided by *WoT*, we decided to conduct our analysis at the level of toxicity incidents regardless of type. We created a dummy variable to capture toxicity (toxic = 1, non-toxic = 0) at the team level (H1-3) as well as the individual level (H4-6). An individual is considered toxic if s/he was reported at least once during the 30-day study period. A team is considered toxic if any player within the team was reported at least once during the battle.

Player skill. *WoT* utilizes a specific algorithm to measure player skill, called Player Rating (PR), which considers the number of battles a player has played, battle results, tank level, damage dealt, and other measures of in-game performance. One's PR can increase or decrease depending on the battle result. As a skill metric, PR is visible to every player outside of the match. We measured the individual PR as well as the average team PR by taking the mean of every team members' PR at the time of the battle. This measure was normalized before analysis.

Team skill advantage over opponent team. The extent to which two teams are well-matched regarding skills was measured through the difference of average team member PRs. We first calculated the team PR by averaging every member's PR, then subtracted the enemy team's PR from the focal team's PR. The larger the metric, the better one's own team's average skill in relation to enemy team's average skill. This measure was normalized before analysis.

Ingroup team skill disparity. Player skill imbalance within a team was measured by taking the standard deviation of all team members' PRs at the time of battle. A large standard deviation indicates that

members' skills are rather unevenly distributed in the team. This measure was normalized before analysis.

Battle type (team size). *WoT* has various types of battles, which determine the type and size of teams participating in these battles. The most popular form is random battle, or pick up battles, where random strangers are assigned into opposing teams following the game's team matching algorithms. Clan-based battles, on the other hand, take advantage of the semi-permanent in-game communities (clans) and are often arranged based on pre-existing social structures. Players generally anticipate future interactions with others in the clan. We aggregated all random battle types and clan battle types and created a dummy variable (random battle = 1, clan battle = 0). The most popular battle size is 15 against 15. Therefore, in testing the effects of team type (H3, random versus non-random), we kept team size constant at 15 to avoid confounding of team size and battle type, though we note that for the clan-based battles the 15-player versions are only accessible to the most experienced players overall.

Exposure to toxicity. This measure captures every player's exposure to toxic behavior, as reflected in the number of reported toxicity instances during battles (there could be more than one instance every battle). We further distinguished a player's exposure to toxicity from their own team (in-group team exposure) and their enemy team (opponent team exposure) and from random battles as well as clan battles.

2.4. Control variables

Battle time was measured by the length of each battle (in seconds). This variable was included as a control and was normalized before analysis. Prior battle count was the player lifetime battle count prior to the study period. Battle count was the total number of battles in which a player was involved during a particular day within the 31-day study period. Table 1 presents descriptive statistics of major study variables.

3. Results

3.1. Antecedents of toxicity

The exploratory research question asked about the general characteristics of toxic players. We combined information on toxicity from the in-game and out-of-game reporting systems. There were close to 500,000 instances of reported toxicity, committed by less than 20% of all active players during the study period.² Using a combined sample of toxic and non-toxic players, we ran a series of Welch's *t*-tests. Welch's *t*-test was chosen as it is a better alternative than the original *t*-test for samples with unequal sample size and variance (Ruxton, 2006).

As shown in Table 2, players who were identified as toxic at least once during the study period were the most experience and involved: they had played significantly more battles prior to the study period, had higher skills prior to the study period, played more battles during the study period, had higher level tanks, and had spent more time in clans than non-toxic players. Interestingly, we also found that those who were reported as toxic during the study period were more likely to be reporting others. Fig. 1 showed the share of toxic players at each PR tier.

To test H1 through H3 about antecedents of toxicity, we ran logistic regression models at the team level (Table 3). The dataset was aggregated by team, with each row representing a team in a particular battle, so every battle had two entries, one for each team involved. The dependent variable was a binary variable for toxicity, with 1 indicating

² The precise numbers of active players were removed to protect Wargaming's business interest.

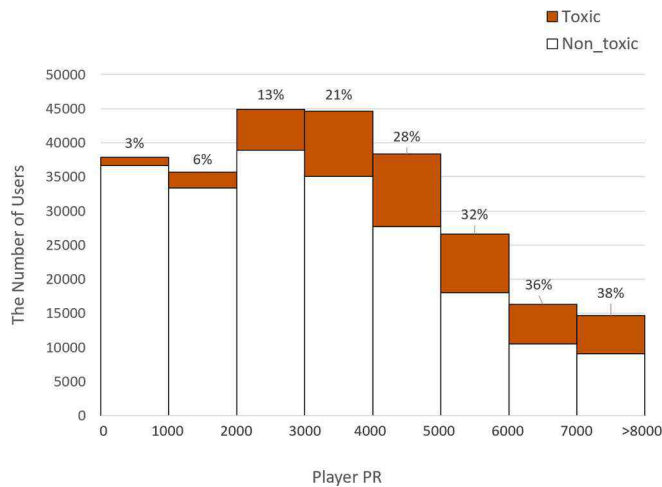


Fig. 1. Percentage of toxic players by player skill (PR).

Table 3
Logistic regression model predicting team toxicity (N=2689130).

	OR	SE	p
Intercept	138.82	0.07	***
Average team skill (z)	1.16	0.003	***
Team skill advantage over opponent team (z)	0.91	0.003	***
Ingroup team skill disparity (z)	1.01	0.003	***
Battle time (z)	1.36	0.002	***
Random battle (yes = 1)	11.50	0.07	***

Note. *** $p < 0.001$. Robust standard errors were used to adjust for battle-level effect.

at least one reported toxic behavior within the team during the battle.³ The independent variables included team skill advantage over opponent team (H1), ingroup team skill disparity (H2), and whether the battle was a random battle versus a clan battle (H3). We also included two control variables—average team skill and battle time. Because every battle involved two teams, we controlled for battle-level effect using robust standard errors through the *rms* R package (Harrell, 2018).

Results showed that both control variables were positive and significant as expected, suggesting that the more skilled the team was on average (OR = 1.16, $p < 0.001$), and the longer the battle was (OR = 1.36, $p < 0.001$), the more likely toxic behaviors were to emerge in the team.

H1 predicted that toxic behaviors are less likely to happen in teams where the players' own team held a skill advantage than the enemy team. This hypothesis was supported. Team skill advantage (own team PR minus opponent team PR) had a negative and significant effect on the likelihood of team toxicity in battles (OR = 0.91, $p < 0.001$).

H2 predicted that toxic behaviors are more likely to happen in teams where members' skills are unevenly distributed. This hypothesis was also supported, as the standard deviation of team PR increases the odds of toxic behavior within the team (OR = 1.01, $p < 0.001$).

H3 predicted that toxic behaviors are more likely to emerge in random battles rather than clan-based battles. This hypothesis received strong support, as toxicity was 11.50 times ($p < 0.001$) more likely to emerge in random battles than clan battles, when controlling for other factors.

³ Because only automatic reporting system captures battle and team information, we included only those toxicity incidents reported through the automatic system and excluded toxicity reports through customer service (approx. 25% of all cases).

Table 4
Event history analysis predicting the hazard of becoming toxic (N = 6933006).

	Model 1		Model 2		Model 3	
	Coeff.	SE	Coeff.	SE	Coeff.	SE
Prior battle count (z)	0.19	0.003	0.16	0.003	0.19	0.003
Prior skill (z)	0.47	0.004	0.48	0.004	0.47	0.004
Battle count for the day (z)	0.11	0.001	0.08	0.001	0.11	0.001
Total daily exposure to toxicity	0.04	0.000				
Exposure to toxicity in own team			0.03	0.000		
Exposure to toxicity in opponent team			0.09	0.001		
Exposure to toxicity in random battles					0.04	0.000
Exposure to toxicity in clan battles					0.20	0.006
Likelihood ratio	77,384, df = 4	***	80,611, df = 5	***	79,590, df = 5	***

3.2. Contagion of toxicity

To test Hypotheses 4–6 about the contagion of toxicity, we employed event history analysis to estimate the effects of various factors on the spread of toxic behaviors within the game. In event history analysis, the outcome variable is the likelihood (hazard) of the occurrence of a particular life event, such as death, at time t , given that the event has not occurred before t . Here, the event of interest was a player becoming toxic for the first time during the study period, and the time unit was one day.

The dataset was aggregated and arranged at the player-day level, with each row representing a player's exposure to toxic behavior during a particular day. Every player has a maximum of 31 observations (days) if the player was never reported as toxic during the study period, or as many observations as the first day when the player turned toxic (e.g., if the player turned toxic on day 5, then there were 5 observations). Independent variables include a player's total exposure to toxic behaviors in battles (regardless of team in which the toxic behavior occurred), which was also further broken down into exposure to toxicity in their own team and toxicity in opponent team, as well as exposure from random battles and clan battles. We also controlled for the player's prior battle count and skill before the observation period, as well as the number of battles played for each day in the period. The independent variables were time-varying, while the control variables (except battle count each day) were time-independent (Table 4).

All control variables were positive and significant, as we expected. The battles played in the past (Coeff. = 0.193, $p < 0.001$), player's skill (Coeff. = 0.469, $p < 0.001$), and the battles a player played during that day (Coeff. = 0.110, $p < 0.001$) all contributed positively to the hazard that player turned toxic on that day.

H4 predicted that toxic behavior is contagious. This hypothesis was supported, as the total exposure to toxicity had a significant and positive effect on the hazard of becoming toxic (Coeff. = 0.039, $p < 0.001$). Exposure to one more toxic incident is associated with a 3.98% (exp(0.039) = 1.0398) increase of the hazard ratio for a player to turn toxic themselves.

H5 predicted that toxic behaviors are more likely to spread in the ingroup team than the opponent team. The results from Model 2 showed that both in-group exposure and outgroup exposure were significantly

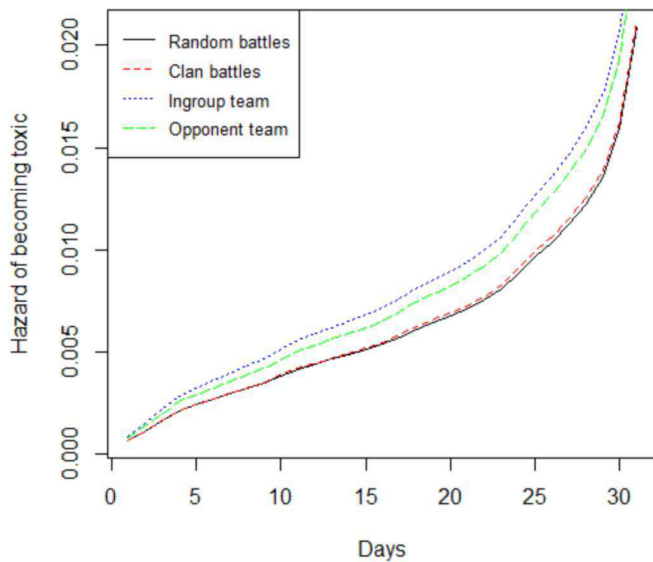


Fig. 2. Hazard of a player becoming toxic over time.

and positively associated with contagion. Exposure to opponent team toxicity (Coeff. = 0.09, $p < 0.001$) actually had a stronger impact than in-group toxicity on contagion (Coeff. = 0.03, $p < 0.001$). The difference between these two coefficients is significant ($p < 0.001$).⁴ Thus, the hypothesis was not supported.

H6 predicted that toxic behavior from random battles are more contagious than those from friends. The hypothesis was not supported. The results from Model 3 showed that toxicity was contagious in both random battles (Coeff. = 0.039, $p < 0.001$) and clan battles (Coeff. = 0.197, $p < 0.001$), but the effect from exposure in clan battles was five times stronger than that from random battles ($p < 0.001$). Thus, H6 was not supported. Fig. 2 shows the hazard ratio of a player becoming toxic over time, differentiated by exposure to toxicity in various team contexts.

4. Discussion

Toxic behaviors are pervasive in online games, and can be harmful to building a positive online environment. Guided by the SIDE model and classical theories of group behaviors, this study represents one of the first systematic efforts to examine the antecedents of toxicity in team-based online games using longitudinal behavioral data directly obtained from the game server. It fills two important gaps in the existing research, by 1) exploring non-verbal and behavioral dimensions of toxicity, and 2) examining team-level predictors in conjunction with individual characteristics. Employing a large-scale behavioral dataset from the popular online game *WoT*, we found that, in general, experienced and skillful players are more likely to commit toxic behaviors. Teams that are losing or have a high internal skill disparity among their members tend to breed toxicity. In addition, this study provides empirical evidence that toxicity is contagious among players, especially toxic behaviors in one's own teams and in clan battles.

4.1. Veterans are more toxic than newbies

Regarding the individual characteristics associated with toxic players, results suggest that veteran gamers who are more experienced, more skilled, and spend more time in clans are more likely to be offenders. All these status characteristics require significant time

⁴ The test of significance between these two coefficients was achieved through a likelihood ratio test between the fit of Model 1 and fit of Model 2.

investment from the players and also indicate a comprehensive socialization process preceding these offenses. This finding is consistent with previous toxicity studies in games (Ballard & Welch, 2017; Chesney et al., 2009). As SIDE predicts, when only given text-based cues, the other players on one's team are not seen as real individuals. So, the inhibitions from behaving poorly towards them are very low.

There are other mechanisms that may also explain why veterans are more toxic than newbies. First, past cyberbullying research generally points to a status or rank differential between the perpetrator and the victim (Ballard & Welch, 2017; Olweus & Limber, 2010). Consistent with that research, a more experienced and skillful gamer may feel a sense of superiority over low-ranked gamers and such status advantage may trigger toxic behavior. Additionally, given the need for agile team coordination in a fast-paced game with a steep learning curve, highly experienced players are often impatient towards newbie teammates who may inadvertently jeopardize the team's success due to their lack of experience. When surviving newbies reach higher levels, they, too, are likely to feel the same sentiment towards less experienced players. This mechanism implies that there are specific victims of such toxic behavior, however our dataset only provided information on the perpetrators, not the victims. We could not verify whether the status differential was indeed one of the triggers of toxicity.

The second explanation lies in a social contagion process. After playing for a longer period of time, gaining more experience through battles, and having more opportunities for social interactions, players are more likely to be exposed to other gamers' toxic behaviors. Contagion of toxic behaviors may happen through the establishment of group norms (Bandura, 1994) as players observe the forms and scale of toxic behaviors and gradually accept them as a "normal" component of online gaming – as well as the emotional contagion of moods that negatively impact cognition, behaviors, and attitudes (Barsade, 2002). This contagion process was confirmed, as our survival analysis showed a positive and significant association between exposure to toxicity and the hazard of becoming toxic oneself.

The third and related explanation recognizes that the current study focused on player-reported, rather than objective, toxic behaviors. It is possible that newbies are just as toxic as more experienced players, but veterans are more likely to get reported, possibly due to more competitive gameplay and a less tolerant environment as players level up. Indeed, our analysis showed a significant reciprocity with regard to reporting — those who were reportedly toxic were also more likely to report others, which attests to a highly vengeful, tit-for-tat dynamic in *WoT*. Yet, given the self-reported nature of the toxicity data, this study could not verify whether reported incidents align proportionally with actual incidents across the player lifecycle.

4.2. Context matters

One important finding of this study is that various contextual factors are also associated with toxicity, beyond individual characteristics. Here, we tested two intertwined contextual variables: outside threat and player stress/frustration. When one's own team is losing (measured by an average skill disadvantage to the opponent team), a gamer is likely to experience negative affect such as frustration and anger. The player's specific appraisal of the situation leads to hostile and toxic behaviors. On the contrary, when one's team is winning, the gamer is likely to experience positive affect and have a more positive outlook of the situation. Toxicity is in turn less likely to occur. Similarly, when the skill level is unevenly distributed within a team, coordination failures and in-group conflicts happen more frequently. Both increase the level of stress and frustration, evoking more toxicity.

Our findings also show that gamers exhibited toxicity significantly more when the battles were convened randomly. When there is little anticipation of future interactions, a self-interested individual may defect because doing so incurs minimal social cost. This phenomenon is exacerbated by online deindividuation, which makes gamers more likely

to dehumanize other players. Unless intended and planned, efforts to improve teamwork are limited by time constraints in random matches. As this pattern persists, gamers will be “contaminated” by toxicity and more gamers will become toxic.

4.3. Toxicity is contagious

We found that exposure to toxicity is associated with becoming toxic oneself. This may partially explain our earlier finding that more experienced and skilled players are also more likely to adopt toxic behaviors. As such, the pervasive normativity of toxic behaviors is reinforced the more experiences players accumulate. This is also consistent with SIDE theory in that in cue-poor environments like the anonymous battles, players not only don't see other players as real people, but they also are more likely to adopt group norms. Observing toxic behavior is therefore likely to create a sense of social proof that is more likely to be acted upon than in a corresponding offline context: toxicity is normal and acceptable here, so I should engage in it as well. These findings carry implications not only in the context of *WoT*, but potentially for many other MMOs and online spaces, as the more time a player dedicates, the more likely they are to espouse toxic behaviors and “contaminate” other players. Such a trend is difficult to reverse, as research suggests that negative events and behaviors produce larger contagion effects than positive events and behaviors (Rozin & Royzman, 2001). This creates a vicious cycle in which those who play more often and are highly skilled become toxic and contaminate the community, disincentivizing new players and/or contributing to the spread of toxicity with repeated exposure that affects the entire game community as whole.

We further broke down total exposure to toxicity into exposure from in-group and opponent teams, as well as exposure from random battles and clan battles. All these categories contributed significantly to contagion, but the effect sizes varied. Contrary to our expectation, exposure to an enemy team's toxicity had a stronger effect on contagion compared to an in-group's team toxicity. This could potentially be explained by a tit-for-tat reaction to bad behaviors by the opponent team. Still, as mentioned previously, the current study lacks information on the target of any toxic behavior, only the behavior itself, so we could not ascertain whether the motivation to “get even” at the perpetrators was indeed the underlying mechanism. More fine-grained data, with information on the type of toxic offense and the specific victims, is needed to further investigate how in-group team and opponent team toxicity may differ in their spreadability. Our last hypothesis that toxicity is more likely to spread when playing with strangers (random pick-up battles) compared to friends (clan battles) was also not supported. Interestingly, toxicity was more contagious in clan battles in comparison to random battles, even though toxicity was much less likely to happen in clan battles than random battles. In other words, friends are much more civil when playing with each other, but when toxicity does occur, it can set up a firestorm. This could be due to social learning in which deviant or toxic behaviors are more likely to be reinforced and imitated when playing with friends (Shores et al., 2014). If friends or clan members further reinforce the normativity of toxicity, these behaviors would be more pervasive and contagious than when playing with strangers in which norms are not as salient due to a less cohesive group identity (Lea et al., 2001).

4.4. Curbing toxicity

Taken together, the current study demonstrates that toxicity becomes more prevalent as players level up, is more likely to occur in teams under stress, and is contagious. These findings echo recent observational and experimental work on trolling behaviors in discussion communities, which shows that, contrary to the belief that trolls are born, not made, deviant behaviors are better explained by mood and discussion context rather than innate factors (Cheng et al., 2017). Given that online games are an integral part of digital life for millions of

people, toxicity in the gaming domain, if left unchecked, could potentially define what we perceive as “normal” online behavior. This perception could well permeate to other non-game online spaces, gaining momentum as negativity begets negativity. How to curb the downward spiral, therefore, remains an imperative question for practitioners and policymakers.

Fortunately, learning the situational factors of toxicity suggests potential practical measures to minimize it. First, because SIDE theory suggests a cue-poor environment is more likely to lead to these behaviors, there is a strong incentive to increase those cues either in the random battles, or by incentivizing more clan formation and membership. As the findings here suggest, clans are far less likely to create toxicity so more time spent in them is likely to reduce toxicity overall. Second, reducing players' stress and frustration could mitigate toxic behavior induced by negative affect. Although competitive gameplay always has winners and losers, the team matching algorithms can be better designed to ensure team skills are evenly matched so that teams do not lose due to an unfair disadvantage. Third, an automatic toxicity detection function to identify uncivil messages could work in conjunction with a selective message filtering system that either delays or mutes uncivil messages during battle. Players may also be subject to a temporary messaging limit or freeze when a flame war is about to happen. These measures aim to defuse the tension when small offenses are about to escalate into more severe toxic behaviors. Additionally, it pays dividends to promote civil behaviors and highlight the sanctions of toxicity, especially early on when players just start the game, as newer members are more likely to adopt and uphold what they perceive as community norms than veterans with entrenched beliefs.

4.5. Limitations and future research

This study has a number of limitations. First and foremost, our study relied on player-reported toxicity rather than behaviorally identified toxicity. A study using toxicity reporting data from *League of Legends* showed that most players did not report toxicity despite exposure to it, so overall toxicity was underreported (Kwak, Blackburn, & Han, 2015). Therefore, the toxicity data in the current study are rather conservative estimates of reality. It is also possible that reported toxicity incidents were systematically different from unreported toxicity due to selection bias. A promising future research direction is to develop and validate a behavioral measure of toxicity independent from player reports. Still, a big challenge lies in the generalizability of such a measure, as it may rely heavily on the genre and gameplay mechanisms of specific game titles.

Second, due to the nature of the reporting system, we did not know whether these toxic behaviors were targeted to specific victim(s), the in-group, the opponent group, or nobody in particular. As characteristics of the victim(s) are critical to understanding the various antisocial behaviors, we were unable to pinpoint specific mechanisms of toxicity. Instead, this study discusses toxicity in broad strokes, describes the individual and team level predictors, and conjectures about potential mechanisms. Further research is warranted to unpack different categories of toxic behavior in games and test theoretical mechanisms underlying each category.

Additionally, our behavioral dataset is limited in various ways. *WoT* is known for its steep learning curve and hardcore gameplay culture, which may have contributed to the prevalence of toxicity. The dataset did not include chat records, which could shed light on verbal toxicity in conjunction with behavioral toxicity. The timespan of our dataset is limited to one month. Although we conducted a longitudinal survival analysis, the dataset captured neither players' exposure to toxic behaviors outside of the observation period, nor player retention dynamics after the observation. All these limitations call for replication studies using longer and more complete sets of behavioral data from various other game worlds.

Author contributions

CS carried out study design, data analysis and wrote the first draft. QS, TK, and GW contributed to literature review and data analysis, RR and DW contributed to study design, data analysis and manuscript revisions.

Acknowledgement

We thank Wargaming for providing data access, and Jeremy Balenger and Eugene Kislyi for their helpful feedback. The last author received funding from Wargaming for unrelated consulting work, but Wargaming did not provide funding for the study reported here.

References

- Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, *211*, 1390–1396.
- Balci, K., & Salah, A. A. (2015). Automatic analysis and identification of verbal aggression and abusive behaviors for online social games. *Computers in Human Behavior*, *53*, 517–526.
- Ballard, M. E., & Welch, K. M. (2017). Virtual warfare: cyberbullying and cyber-victimization in MMOG play. *Games and Culture*, *12*, 466–491.
- Bandura, A. (1994). The social cognitive theory of mass communication. In J. Bryant, & D. Zillmann (Eds.), *Media effects: Advances in theory and research* (pp. 61–90). Hillsdale, New Jersey: Erlbaum.
- Barsade, S. G. (2002). The ripple effect: Emotional contagion and its influence on group behavior. *Administrative Science Quarterly*, *47*, 644–675.
- Blackburn, J., & Kwak, H. (2014). STFU NOOB!: Predicting crowdsourced decisions on toxic behavior in online games. In *Proceedings of the 23rd international conference on World wide web* (pp. 877–888). Seoul, Korea: ACM.
- Cheng, J., Bernstein, M., Danescu-Niculescu-Mizil, C., et al. (2017). Anyone can become a troll: Causes of trolling behavior in online discussions. *CSCW : Proceedings of the Conference on computer-supported cooperative work. Conference on Computer-Supported Cooperative Work*, 1217–1230, 2017.
- Chesney, T., Coyne, I., Logan, B., et al. (2009). Griefing in virtual worlds: Causes, casualties and coping strategies. *Information Systems Journal*, *19*, 525–548.
- Christakis, N. A., & Fowler, J. H. (2009). *Connected: The surprising power of our social networks and how they shape our lives*. New York: Little, Brown and Co.
- Coe, K., Kenski, K., & Rains, S. A. (2014). Online and uncivil? Patterns and determinants of incivility in newspaper website comments. *Journal of Communication*, *64*, 658–679.
- Elliott, T. (2012). *Flaming and gaming—computer-mediated-communication and toxic disinhibition*. University of Twente.
- Ellis, A. P. (2006). System breakdown: The role of mental models and transactive memory in the relationship between acute stress and team performance. *Academy of Management Journal*, *49*(3), 576–589.
- ESA. (2019). *Essential facts about the computer and video game industry*. Available at: https://www.theesa.com/wp-content/uploads/2019/05/ESA_Essential_facts_2019_final.pdf.
- Fox, J., & Tang, W. Y. (2017). Women's experiences with general and sexual harassment in online video games: Rumination, organizational responsiveness, withdrawal, and coping strategies. *New Media & Society*, *19*, 1290–1307.
- Gervais, B. T. (2015). Incivility online: Affective and behavioral reactions to uncivil political posts in a web-based experiment. *Journal of Information Technology & Politics*, *12*, 167–185.
- Gervais, B. T. (2017). More than mimicry? The role of anger in uncivil reactions to elite political incivility. *International Journal of Public Opinion Research*, *29*, 384–405.
- Graham, T., & Wright, S. (2015). A tale of two stories from “below the line”: comment fields at the guardian. *International Journal of Press/Politics*, *20*, 317–338.
- Hancock, J. T., & Dunham, P. J. (2001). Impression formation in computer-mediated communication revisited: An analysis of the breadth and intensity of impressions. *Communication Research*, *28*, 325–347.
- Harrell, F. E. (2018). *rms: Regression modeling strategies*, *5.1*.
- Heide, J. B., & Miner, A. S. (1992). The shadow of the future: Effects of anticipated interaction and frequency of contact on buyer-seller cooperation. *Academy of Management Journal*, *35*, 265–291.
- Infante, D. A., & Wigley, C. J. (1986). Verbal aggressiveness: An interpersonal model and measure. *Communication Monographs*, *53*, 61–69.
- Johnson, D., Nacke, L. E., & Wyeth, P. (2015). All about that base: Differing player experiences in video game genres and the unique case of MOBA games. In *Proceedings of the 33rd annual ACM Conference on human factors in computing systems. Seoul, Republic of Korea* (pp. 2265–2274). ACM.
- Kiesler, S., Siegel, J., & McGuire, T. W. (1984). Social psychological aspects of computer-mediated communication. *American Psychologist*, *39*, 1123–1134.
- Kou, Y., & Nardi, B. (2013). Regulating anti-social behavior on the internet: The example of League of Legends. *iConference 2013 Proceedings*, *7*.
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*, 8788–8790.
- Kwak, H., Blackburn, J., & Han, S. (2015). Exploring cyberbullying and other toxic behavior in team competition online games. In *Proceedings of the 33rd annual ACM Conference on human factors in computing systems. Seoul, Republic of Korea* (pp. 3739–3748). ACM.
- Lea, M., Spears, R., & de Groot, D. (2001). Knowing me, knowing you: Anonymity effects on social identity processes within groups. *Personality and Social Psychology Bulletin*, *27*, 526–537.
- Maher, B. (2016). Can a video game company tame toxic behaviour? *Nature*, 568–571.
- Marques, J., Abrams, D., & Seródio, R. G. (2001). Being better by being right: Subjective group dynamics and derogation of in-group deviants when generic norms are undermined. *Journal of Personality and Social Psychology*, *81*, 436–447.
- Märtens, M., Shen, S., Iosup, A., et al. (2015). 2015 Toxicity detection in multiplayer online games. In *International workshop on network and systems support for games (NetGames)*.
- Mason, W., & Clauset, A. (2013). Friends FTW! Friendship and competition in halo: Reach. In *Proceedings of the 2013 conference on Computer supported cooperative work* (pp. 375–386). San Antonio, Texas, USA: ACM.
- Matheson, K., & Zanna, M. P. (1988). The impact of computer-mediated communication on self-awareness. *Computers in Human Behavior*, *4*, 221–233.
- Neto, I. R., Yokoyama, K. M., & Becker, K. (2017). Studying toxic behavior influence and player chat in an online video game. In *Proceedings of the international conference on web intelligence* (pp. 26–33). Leipzig, Germany: ACM.
- Olweus, D., & Limber, S. P. (2010). Bullying in school: Evaluation and dissemination of the Olweus bullying prevention program. *American Journal of Orthopsychiatry*, *80*, 124–134.
- Papacharissi, Z. (2004). Democracy online: civility, politeness, and the democratic potential of online political discussion groups. *New Media & Society*, *6*(2), 259–283. <https://doi.org/10.1177/1461444804041444>.
- Pinto, I. R., Marques, J. M., Levine, J. M., et al. (2010). Membership status and subjective group dynamics: Who triggers the black sheep effect? *Journal of Personality and Social Psychology*, *99*, 107–119.
- Rains, S. A., Kenski, K., Coe, K., et al. (2017). Incivility and political identity on the internet: Intergroup factors as predictors of incivility in discussions of news online. *Journal of Computer-Mediated Communication*, *22*, 163–178.
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, *5*, 296–320.
- Ruxton, G. D. (2006). The unequal variance t-test is an underused alternative to Student's t-test and the Mann-Whitney U test. *Behavioral Ecology*, *17*, 688–690.
- Shen, C. (2014). Network patterns and social architecture in massively multiplayer online games: Mapping the social world of EverQuest II. *New Media & Society*, *16*, 672–691.
- Shores, K. B., He, Y., Swanenburg, K. L., et al. (2014). The identification of deviance and its impact on retention in a multiplayer game. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (pp. 1356–1365). Baltimore, Maryland: USA: ACM.
- Short, J., Williams, E., & Christie, B. (1976). *The social psychology of telecommunications*. London: Wiley.
- Staw, B. M., Sandelands, L. E., & Dutton, J. E. (1981). Threat rigidity effects in organizational behavior: A multilevel analysis. *Administrative Science Quarterly*, *26*, 501–524.
- Suler, J. (2004). The online disinhibition effect. *CyberPsychology and Behavior*, *7*, 321–326.
- Thorson, K., Vraga, E., & Ekdale, B. (2010). Credibility in context: How uncivil online commentary affects news credibility. *Mass Communication & Society*, *13*, 289–313.
- Van Lange, P. A. M., Klapwijk, A., & Van Munster, L. M. (2011). How the shadow of the future might promote cooperation. *Group Processes & Intergroup Relations*, *14*, 857–870.
- Vargo, C. J., & Hopp, T. (2017). Socioeconomic status, social capital, and partisan polarity as predictors of political incivility on twitter: A congressional district-level analysis. *Social Science Computer Review*, *35*, 10–32.
- Vella, K., Klarkowski, M., Johnson, D., et al. (2016). The social context of video game play: Challenges and strategies. In *Proceedings of the 2016 ACM conference on designing interactive systems* (pp. 761–772). Brisbane, QLD, Australia: ACM.
- Walther, J. (2006). Nonverbal dynamics in computer-mediated communication, or :(and the net :(s with you, :) and you :) alone. In V. Manusov, & M. Patterson (Eds.), *The Sage handbook of nonverbal communication* (pp. 461–480). Thousand Oaks, CA: Sage.
- Williams, D., Caplan, S., & Xiong, L. (2007). *Can you hear me now? The social impact of voice on internet communities*. Human Communication Research.
- Williams, D., Ducheneaut, N., Xiong, L., et al. (2006). From tree house to barracks: The social life of guilds in world of warcraft. *Games and Culture*, *1*, 338–361.
- Williams, D., Yee, N., & Caplan, S. E. (2008). Who plays, how much, and why? Debunking the stereotypical gamer profile. *Journal of Computer-Mediated Communication*, *13*, 993–1018.
- de Zavala, A. G. (2011). Collective narcissism and intergroup hostility: The dark side of ‘in-group love’. *Social and Personality Psychology Compass*, *5*, 309–320.