Understanding the Dynamics of Attitude Contagion in Multiplayer Minecraft

Kongmeng Liew, Ursula Thomson, Averi Ray, Michael Alpenfels, and Jacinta Cording
University of Canterbury

Author Note

Correspondence concerning this article should be addressed to Kongmeng Liew, School of Psychology, Speech and Hearing, University of Canterbury. E-mail: kongmeng.liew@canterbury.ac.nz

Abstract

Multiplayer Minecraft has evolved into a distributed sociotechnical environment where platform affordances, governance regimes, and players' motivations and behaviours intersect to create conditions that support both prosocial interactions and the reinforcement of problematic norms. This research examines how attitudes circulate within these environments by integrating telemetry based behavioural modelling with qualitative accounts of multiplayer participation. Study 1 introduces a telemetry driven method for extracting real time relational data from Minecraft servers and demonstrates how proximity based social networks can be constructed to reveal structural patterns of co-presence, influence pathways, and community formation over a twelve week intervention. Study 2 draws on reflexive thematic analysis of interviews with adult New Zealand Minecraft players (N = 7) to show how individuals actively curate their multiplayer experiences by prioritising safety, competence boundaries, and relational intimacy. Participants describe retreating into small, predictable, tightly moderated micro communities that privilege trust and collaboration over exposure to large public spaces. When taken together, the two studies highlight how self selected relational structures can facilitate stronger norms reinforcement, whether supportive or harmful, through playing on a server. The findings contribute a methodological framework for combining fine grained telemetry with qualitative inquiry and provide insights relevant to platform governance, intervention design, and the prevention of harmful influence processes in multiplayer game ecosystems.

Understanding the Dynamics of Attitude Contagion in Multiplayer Minecraft

Introduction

Over the past few decades, multiplayer videogames as a platform have evolved from discrete entertainment to complex sociotechnical ecosystems whereby social identities, norms, and political worldviews are deeply intertwined. More broadly, research on online interactive platforms has highlighted how the interaction between platform affordances and digital subcultures has made these online spaces fertile ground for the normalisation and diffusion of extremist ideologies. Notably, Heslep and Berge's Heslep and Berge (2024) study of Discord's hate-network servers provides a template for analysing these dynamics: through zeroing on Discord community building, they demonstrate how extremist internet groups flourish through seven primary mechanisms: the everyday normalisation of extremist language and identity cues, the exploitation of platform affordances that enable closed echo chambers, governance gaps created by third-party discovery tools, stratified visibility that conceals toxic subcultures under a 'clean' surface, the appeal of transgressive 'edginess', networked reinforcement across linked communities, and low-barrier, anonymous entry that facilitates identity experimentation. These mechanisms align with broader sociological and communications research showing that digital social platforms provide dangerous infrastructures that are able to cultivate a community-driven normalisation of extreme positions and for coordination among like-minded actors (Risius et al., 2024; Shaw, 2023).

Here, multiplayer Minecraft is not merely a game but also a decentralised constellation of user-run servers, modding communities, and adjacent interpersonal communication platforms (Discord, Twitch, YouTube). Evidence shows that extremist actors have appropriated Minecraft's creative and administrative affordances, such as custom worlds, private servers, modpacks, plugins, and role permissions, and embed propaganda in virtual architecture, gamify xenophobic narratives, and cultivate invite-only communities with minimal oversight (Gagandeep, 2025; *Research Launch on Gaming and*

Violent Extremism, 2022). In this sense, Minecraft, as an online platform, replicates the conditions proposed by Heslep & Berge as a platform for potentially problematic online interactions: its decentralised governance mirrors Discord's offloaded discovery and moderation gaps; its server-to-server mobility and cross-platform linkages reproduce the networked reinforcement described above; and its pseudonymous, low-friction entry lowers social risk for 'trying on' identities that may gradually shift from ironic transgression to sincere ideological alignment through repeated exposures to these extremist communities and online social interactions (Gagandeep, 2025; Wells et al., 2024).

Moreover, peace- and security-studies perspectives underscore that radicalisation online is emergent and relational, not reducible to singular "bad actors" or linear causal triggers. Rather, it reflects interactions between platform affordances, governance regimes, and user motivations within distributed networks that can migrate across platforms and re-organise after enforcement (Amarasingam, 2024). Accordingly, we emphasise that Minecraft is not as an inherently extremist platform, but, much like Heslep & Berge (2021) illustrated with Discord, a case of a participatory, decentralised platform whose sociotechnical design and community ecologies can be exploited by ideological extremists.

Gamified Interventions in Minecraft

On the flip side, researchers have found that the same platform affordances that may facilitate the proliferation of extremist networks, can also be utilised and designed to facilitate positive social interventions and social change. Researchers have been effectively using Minecraft as a versatile digital platform for gamified mental health interventions in youth, by integrating digital social interventions with high ecological validity and taking advantage of its sandbox environments for experimental adaptability (Brashears, 2024). Compared to other forms of multiplayer games, Minecraft servers as a platform offers an easier time for researchers to develop platform-based software tools (server-side mods and plugins), with arguably less frustration and perceived difficulty than other popular games.

Past research has used Minecraft effectively to facilitate mental health interventions for adolescents and youths. For example, Jaganneth et al. explored the effectiveness of an after school programme for youths aged 8-14 run on a moderated Minecraft server over six months, and found that the programme was effective in promoting conflict resolution skills amongst participants (Jagannath et al., 2020). These have been built on gamification principles in designing interventions that are fun and engaging, but also effective in administering play-based treatments to psychologically vulnerable populations. In a systematic review of Minecraft-based interventions in educational settings, Slattery and colleagues found that one common thread in these Minecraft-based interactions is in the strength of its social domain: Players often find a sense of belonging from community-based social interactions on the server, and this social aspect is often a strong contributor to the effectiveness of these interventions (Slattery et al., 2025). These show that Minecraft as a digital, multiplayer platform, is ideologically agnostic - it can be used by bad actors to propagate extremist views and ideologies in players, but can also foster social belonging and promote positive mental health outcomes when adaptively and appropriately used.

The Need for Telemetric Tools

Why do some individuals gravitate towards problematic, extremist communities, and why do some individuals experience positive growth outcomes from playing on the same multiplayer platform? Past studies have suggested numerous reasons from the perspective of the individual and society, but comparatively, much less is known about how platforms themselves may interact with these individual and societal factors. We propose that a comprehensive understanding of extremist attitude contagion - or any kind of attitude contagion - on Minecraft must include analyses of platform-specific behaviours, i.e., the actions exhibited by players within the Minecraft environment. Indeed, much of the research in this space has been qualitative or from retrospective survey-based data. By applying our proposed approach of examining players' in-game behavioural telemetry directly from the Minecraft server itself, we can better understand how players behave in relation to each other, and directly model player-to-player relationship dynamics. To achieve this, we also rely on social network analyses (SNA), a commonly-used method in the social and behavioural sciences to model relational data, and construct a network of players based on the relative interactions (spatial proximity) with other players on the server. We believe this provides a first step towards decoding attitude contagion, particularly extremist attitude contagion, within multiplayer Minecraft environments.

The Present Research

The present research comprises two separate studies with different approaches, towards the goal of understanding how player-to-player interactions may facilitate attitude contagion in multiplayer Minecraft. For Study 1, we adopt a technical, software engineering approach to design and build Minecraft-compatible software (plugins) and develop a system to retrieve real time player locations and behavioural logging (e.g., Minecraft block placement, chat logs) by all players on a Minecraft server. For Study 2, we adopt a qualitative approach and interview N = 7 New Zealand Minecraft players to explore how attitudes can be influenced by their online experiences.

Study 1

A Primer on Social Network Analyses

We first introduce Social Network Analyses (SNA) as a methodological approach for examining patterns of relationships among nodes, at the individual, organisational, or communities—level, within a social network. Unlike traditional analyses that focus on individual attributes, SNA emphasizes relational ties and their structural properties, enabling insights into how positions and connections influence behaviours and outcomes (Borgatti et al., 2009; Wasserman & Faust, 1994). Rooted in sociology and graph theory, SNA represents networks as graphs composed of nodes and edges (connections), that allow researchers to quantify features such as centrality and clustering. These metrics reveal patterns of influence, cohesion, and subgroups within networks (Prell & Schaefer, 2024). Applications span diverse fields—from mapping disease transmission in epidemiology to identifying key influencers in organisational settings (Yang, 2024), and it is now an important tool understanding the relational dynamics embedded within complex social systems (Marin & Wellman, 2011).

In the present research, Social Network Analysis (SNA) was utilised to delineate the interpersonal dynamics that arise during gameplay by representing individuals as nodes inside weighted directed networks derived from in-game interaction data. Although we chose to define connections between nodes by social relationships as characterised by sustained proximity ($\leq$10 blocks for $\geq$30 seconds), these can be defined through a number of means, and past research has defined these through communication events (Menini et al., 2019), collaborative gameplay activities (Seif El-Nasr et al., 2021), and social media links (Logan et al., 2023). By defining our network connections, or edges to represent proximity between players, we quantify both the frequency and direction of interactions among players on the server. Standard network metrics (e.g., degree, betweenness, and PageRank centralities, community detection, and network densities) were calculated to assess players' social integration, influence, and connectedness. We propose that these metrics can help researchers better understand the interpersonal dynamics on the server, as a means to developing more holistic or integrated activities for better long term effectiveness of the intervention.

Related work

Pirker et al. explored the application of SNA in games to understand player interactions and community structures, showing how network metrics like centrality, community detection, and connectivity, reveal patterns of collaboration, competition, and influence among players (Pirker, 2018). They demonstrate that SNA provides insights beyond traditional behavioural analytics by uncovering relational dynamics that shape player experience and engagement. The study also highlights methodological considerations for integrating SNA into game user research, including data collection from in-game communication and interaction logs. In doing so they show how SNA can be used to identify key players, predict social cohesion, and inform game design decisions aimed at fostering healthy communities. Our research builds on this framework specifically in a Minecraft multiplayer setting, towards our eventual goal of building tools to examine the effectiveness of multiplayer gamified interventions. Next, Seif El-Nasr et al. applied SNA to Tom Clancy's The Division to identify influential players and their impact on retention (Seif El-Nasr et al., 2021). Using six network centrality measures, they identified a small number of players with high centrality scores, that they termed as influencers: highly connected players who actively create player groups. They found that these influencers significantly increase others' playtime and social engagement, above and beyond other players. Longitudinal analysis showed influencers also drive community growth, with 26% of their peers becoming influencers after a year. Unlike our focus on SNA applied directly to game telemetry data, this analysis was based on users' linked 'friends' from Uplay, the online social networking service from Ubisoft that was needed to play Tom Clancy's The Division. Finally, Rodrigues and Mustaro examined virtual communities in 63 Ragnarök Online players using SNA (Rodrigues & Mustaro, 2007). They used network metrics, such as degree, density, and connectivity, to understand interaction patterns and community formation, highlighting how multiplayer gameplay fosters large-scale social ties through cooperative and competitive play, and establish that SNA metrics can reveal key actors and relationships critical for information flow and group cohesion. However, their networks were constructed from survey data. For example, in constructing networks based on spatial location, players may be more strongly linked if they frequent the same (self-reported) virtual locations. This is different from the more direct use of telemetry to quantify the time spent playing in close spatial proximity on the server.

Methods

For Study 1, we research aim to develop a method to quantify players' social interactions on a Minecraft server. This comprises two steps: 1) We develop a server-side plugin for tracking and synchronous storing all player locations on a SQLite database, and 2) testing the plugin through conducting a SNA on players participating in a 12-week

Minecraft-based intervention for youth with extreme social withdrawal, based on the proximity of their in-game locations.

The Location Logger Plugin

We developed a server-side plugin, the Location Logger, to provide player location tracking within a Paper-based Minecraft server environment (version 1.20.6). To safeguard participant privacy and enhance administrative oversight of sensitive data, all telemetry records were retained locally on the game server instead of being sent to external platforms by storing data internally. Database operations were conducted using Java Database Connectivity (JDBC) utilising a SQLite adaptor supplied by Xerial. Due to the moderate size of the information, SQLite was chosen over larger database systems for its simplicity, portability, security, and minimal processing requirements. The data format was deliberately streamlined and comprised three interconnected tables aimed at enabling quick extraction and compatibility with subsequent machine-learning algorithms. This optimised architecture guaranteed operational dependability while preserving analytical adaptability for behavioural and psychological modelling.The plugin was developed with the Spigot application programming interface (API) and engineered to function discreetly, ensuring that data collecting had minimal impact on gameplay speed. The Bukkit Scheduler API was employed to asynchronously log positional telemetry at consistent thirty-second intervals for all active players. Each telemetry item was recorded as a seven-element tuple consisting of: (1) a unique row identification, (2) a timestamp in Unix epoch format, (3) a player identifier, (4) the world or dimension of play, and (5–7) the spatial coordinates (X, Y, and Z) of each player's location. This ongoing geographical sampling facilitated detailed observation of player mobility, occupation of in-game activity zones, and engagement level throughout synchronous gameplay sessions, for a verifiable and ecologically valid record of player movement activity.

Participants and Data

Telemetry data were gathered from volunteer participants involved in naturalistic gameplay sessions on the university-run Minecraft servers established for this study. Participants (N=29) were either volunteers and staff, or youth clients involved in a 12-week intervention study on the platform. Every week, participants engaged in a synchronous 2-hour session on the server and completed activities related to the intervention primarily on the Minecraft server. However, all participants were free to log on at any time, for causal play on the server. As this research aims to model the interaction dynamics through social network analyses, details of the intervention are outside the scope of the paper.

Player coordinates were continuously recorded and stored in a SQLite database before being exported in CSV for analyses in Python and R. The movement trajectories formed spatial ’behavioural fingerprints’ over time, allowing for the precise reconstruction of individual navigation paths, proximity interactions, and engagement dynamics essential to the objectives of this study. The research server was situated on dedicated university computing infrastructure. Each server facilitated concurrent multiplayer engagement and employed conventional gameplay features. All participants were allowed access to the gaming environment at any point in time, with minimal administrative interference from the development team.

Each telemetry entry contained the following elements: Unique row identifier, unix epoch timestamp. player identifier. world or dimension identifier, X-coordinate (east–west axis), Y-coordinate (vertical altitude), and Z-coordinate (north– south axis). All timestamps were converted from Unix epoch time into timezone-aware UTC datetime objects to ensure temporal consistency across sessions. Any malformed or non-parsable timestamps were automatically flagged and excluded to preserve dataset integrity. For spatial analysis purposes, the dataset was then reduced to essential behavioural features including time, user, x, y, and z, yielding a clean and standardized dataset suitable for spatial and network modelling.

Network Construction

Towards a Directed Social Network

The social network was constructed as a directed, weighted network derived from in-game telemetry on player locations. Nodes represent individual players participating on the Minecraft server, while edges encode proximity-based interactions between players. Specifically, an edge from Player A to Player B was created whenever A was within a 10-block radius of B during a 30-second time bin, reflecting meaningful co-presence for potential social interaction. To capture asymmetric engagement, edges were directional and normalized by the total exposure time of the source node: the weight of A→B represents the proportion of A's observed gameplay spent near B. This normalisation ensures interpretability and controls for heterogeneous session durations. The resulting adjacency matrix was converted into a weighted directed network for computation of standard SNA metrics such as degree, PageRank, and betweenness centrality, as well as community detection and visualization.

Temporal Binning and Alignment

Because player location logs were not consistently synchronised across users, a temporal harmonisation process was implemented to facilitate accurate distance comparisons. Timestamps were rounded down to the nearest 30-second interval, creating a new variable (time_bin) that grouped player positions into discrete world snapshots. For each player × time_bin combination, if multiple coordinate entries existed, the most recent observation was retained. If no record was present, the player was considered absent from that interval. This procedure produced a temporally aligned snapshot dataset in which every row corresponded to a player's most recent location within a 30-second sampling window. The total number of time bins per player was also computed, serving as an exposure-time metric used for later normalisation procedures.

Proximity Computation

We first determined how far players were from each other within each 30 second time_bin through euclidean distance estimation. Within each temporally aligned snapshot, pairwise Euclidean distances were computed across all players present according to:

$$d(A, B) = \sqrt{((x_A - x_B)^2 + (y_A - y_B)^2 + (z_A - z_B)^2)}$$

Where, the distance (d) between two Players, A, and B, can be defined as the collective difference in x, y, and z spatial dimensions. A proximity event was recorded whenever two players were located within a predefined threshold of 10 blocks, representing meaningful in-world co-presence for social interaction. Proximity was operationalised directionally to preserve asymmetric social behaviour. This mean that participants' proximities to each other were relative to the total number of time (bins) they spent on the server: if Player A and Player B interacted for 1 hour, but Player A only spent 2 hours total on the server, but Player B spent 4 hours total on the server, Player A would have spent a larger proportion of their time with Player B, than vice versa. As such, for each ordered player pair (A $\rightarrow$ B), if player A was within the proximity threshold of player B during a snapshot, a directed interaction count was incremented: $\text{count}_{A \rightarrow B} + = 1$

Normalisation of Proximity Counts

To adjust for unequal session durations across players, raw proximity counts were normalized by individual exposure time: $\mathrm{w}_{(\mathrm{A} \rightarrow \mathrm{B})} = count_{A \rightarrow B)} / totaltimebins_A$

Where, the weights for the directed edge from Player A to B, would be defined as the number of time_bins that feature a co-occurrence of A and B within the 10 block radius, divided by the total number of bins where Player A is present in. This normalisation produced a directed, weighted adjacency matrix in which each edge reflected the proportion of a player's observed gameplay spent in proximity to another player, yielding interpretable indices of social engagement independent of activity duration.

Results

The normalised adjacency matrix was converted into a weighted directed network graph using the igraph package (Csárdi et al., 2025) in R (R Core Team, 2022). Within this graph, nodes represent individual participants, and edges corresponded to normalized proximity weights. Node attributes included exposure duration measured by total time-bin participation, and edge attributes are based on the transformation of raw proximity counts and normalised weights relative to each player (node).

The overall network achieved a density of 0.52. A suite of network measures was computed for each participant. These were computed for out-degree centrality (frequency of contact with others relative to players' total playtime; Mean = 0.52, SD = 0.30), total degree centrality (overall social integration; Mean = 1.03, SD = 0.60), PageRank centrality (recursive measure of players' influence, in accounting for the corresponding centrality/influence of other players that they interact with; Mean = 0.03, SD = 0.02), betweenness centrality (bridging or mediator roles; Mean = 0.04, SD = 0.06), and communities detected from the Walktrap community detection algorithm (Community 1: n = 24; Community 2: n = 3). Here, the communities appeared to accurately represent participants on the intervention. Staff and volunteers, and youth participants comprised Community 1. In Figure 1, this is shown by the yellow nodes, and clients are indiciated through an orange outer circle. Additionally, Community 2 comprised three builders who were active prior to the launch of the intervention, in designing and building the various activities on the server. All other unrepresented nodes show volunteers and research team members who were present during the set up of the server, but were not active during the intervention proper. The network visualisation in Figure 1 (all figures at the end of the paper) used the Fruchterman–Reingold algorithm for force-directed layout.

Spatial Behaviour and Movement Analysis

To model movement behaviour independently of directional bias, absolute X and Z coordinate values were calculated to reflect relative distances from other players. For each participant, trajectory plots were generated to visualise players' individual locations throughout gameplay. A composite spatial trajectory map depicted population-level spatial dynamics, and an XZ scatter distribution plot aggregated all movement points within a 1000 by 1000 block area centred on the initial spawn point in the map. This visualisation is available in Figure 2, at the end of the paper.

Discussion

Summary of Main Findings

Our analysis was applied on a 12-week intervention (N = 29 players) using Minecraft, and demonstrates the feasibility and interpretive power of mapping player interactions through social network analyses. The resulting network exhibited high density (0.52), suggesting frequent co-presence and strong social connectivity among participants. Community detection revealed two distinct clusters: a primary group comprising intervention participants (staff, volunteers, youth) and a smaller cluster of builders active prior to launch. This alignment with known roles validates the network's ability to capture meaningful social structures. Both visualisations, (network structure and spatial overlay), highlighted structured interaction patterns around central landmarks and emergent congregation zones. Centrality measures provided complementary insights: degree centrality reflected breadth of contact, PageRank identified players connected to influential peers, and betweenness highlighted potential brokers bridging subgroups. These findings underscore the potential of telemetry-driven SNA to surface actionable patterns in multiplayer Minecraft.

Study 2

For Study 2, we explored New Zealand Minecraft players' (N = 7) experience with multiplayer Minecraft. This qualitative study adopted a reflexive, inductive design using Braun and Clarke's Braun and Clarke (2006) Reflexive Thematic Analysis (RTA) to explore players' experiences of multiplayer participation in Minecraft. RTA treats meaning-making as an interpretive, situated practice rather than a procedure aimed at coding reliability or theme quantification; themes are developed by the analyst through sustained, reflexive engagement with the dataset, with an emphasis on transparency and coherence over consensus metrics.

Methods

Participants

We used snowball sampling to recruit adult Minecraft players who had experience with multiplayer Minecraft (e.g., Realms, LAN play, private or public servers). Recruitment invitations were distributed through the research team's mailing lists and student networks. Interested individuals contacted the research team directly to receive an information sheet and schedule an interview. Snowball referrals were entirely participant-led (no contact details were requested from third parties). Eligibility criteria were: age 18 and above, New Zealand residents, and experience playing on a Minecraft multiplayer server. Seven participants (ages 22–24) took part, identifying as female (N=3), male (N=2), and non-binary (N=2); all identified as New Zealand European/Pākehā (see Table 1). All participants received a $20 grocery voucher for participating in the interview.

Data Collection

Interviews were conducted one-to-one via Zoom. A semi-structured guide invited participants to narrate play histories and experience with Minecraft more generally, before being guided to respond to questions on multiplayer practices, social interactions, norms and moderation experiences, and attitudes that have changed as a result of playing with other players online. Interviews lasted approximately 30–50 minutes and were audio-recorded with consent, then automatically transcribed through Zoom's transcription service, before being manually checked for accuracy. Online interviewing practices (e.g., pre-call tech checks, establishing privacy, camera preferences) followed established guidance on synchronous interviewing to support comfort and data richness. The study received approval from the University of Canterbury Human Research and Ethics Committee. Before each interview, participants reviewed an information sheet and provided informed consent. Pseudonyms (or preferred identifiers) are used in reporting, and all potentially identifying details were removed during transcription.

Analytic Approach: Reflexive Thematic Analysis

Analysis followed the six, non-linear phases of RTA: (1) data familiarisation; (2) generation of initial codes; (3) construction of candidate themes; (4) review of themes against coded extracts and the dataset; (5) theme definition/naming; and (6) producing the analytic narrative with illustrative data extracts (Braun & Clarke, 2006). Coding was conducted inductively at a primarily semantic level with sensitivity to latent meanings by the lead author. Consistent with RTA, analysis was not aimed at inter-coder reliability; instead, we prioritised reflexive dialogue and coherence between claims, data extracts, and the broader dataset.

Researcher Reflexivity and Positionality

The interviewer is a postgraduate student in psychology at the University of Canterbury, and has had prior research experience and familiarity with qualitative coding and research. Interview protocols were designed under the supervision of the lead author, who has over a decade in experience in psychological research, including in qualitative methods. Both the postgraduate student interviewer and the lead author are familiar with multiplayer Minecraft, and have previous experience playing recreationally on Minecraft servers.

Results

We present a cross-case, reflexive thematic analysis of seven interviews exploring players' experiences of multiplayer participation in Minecraft, with a particular focus on show players are affected by, respond to, and reshape their online social environments.

Safety in the Small

Across cases, the use of the term ’multiplayer’ rarely meant stepping onto large, open servers. Instead, participants consistently gravitated toward local (LAN) play, Minecraft Realms, or whitelisted/private servers. This was mainly due to how playing with other players that they may already be familiar with, provided for a more comfortable and socially enjoyable time.

> "Realms . . . help me not lose interest . . . I feel extremely comfortable just playing with people I know." (P6, 24NB, NZ European)

> "I play on LAN with my partner . . . I get frustrated at [griefing], that's one reason I'm not great at multiplayer." (P7, 23F, NZ European)

For some participants, this avoidance of large scale public multiplayer severs were due to prior negative experiences on these servers.

> "[A mod] teleported me into a big stone room . . . and interrogated me . . . I logged off, and I never played on that server again." (P1, 22F, NZ European)

> "I just report it, and then get grumpy when nothing happens . . . That's another reason I don't play on servers very often." (P2, 23F, NZ European)

Participants inidcate a preference for local or small-scale intimate multiplayer settings by scaling it down. Predictability, trust, and familiarity are often prioritised over the variety and visibility that public servers might bring. In effect, this suggests that for participants, '*safe*' multiplayer is '*small*' multiplayer.

Competence Gap Anxiety from Evolving Metas

Participants frequently narrated feeling left behind by changes to combat mechanics, game metas, or the high skill ceiling in Player-versus-player (PvP) minigames. These feelings translated into voluntary withdrawal or abstinence from participating in competitive public spaces.

> "A lot of the minigames are now combat-focused, and I don't have the time . . . to get good . . . otherwise I just get absolutely smashed." (P3, 24M, NZ European)

> “I wouldn’t touch anything to do with combat . . . I’ve seen how insane it’s gotten, and I am not up to that par.” (P5, 24NB, NZ European)

> “A lot of people I play with know all the min-maxing . . . I don’t really know much about that at all.” (P6, 24NB, NZ European)

As the community and mechanics evolve, the perceived competence required for comfortable participation increases. Public multiplayer becomes a performance arena rather than a site of playful co-presence, pushing many players away from these larger more public venues, and back toward private, cooperative contexts.

Relational Motivations for Play

Rather than seeking visibility or status, participants framed multiplayer gameplay experiences as a way to improve their social relationships, such as bonding with partners, friends, or co-workers through shared projects and co-presence.

> “Playing with my partner and building our own little world together is a lot of fun.” (P2, 23F, NZ European)

> “Currently, it’s definitely that my girlfriend loves it so much . . . that really keeps me on it.” (P7, 23F, Pākehā/NZ European)

> “You play it with your friends, and it works, even though [progression/combat] is . . . awful . . . it works.” (P1, 22F, NZ European)

For one participants, this was not only to improve their relationships with people they knew offline, but also for social resonance with the friends that they made online through the server.

> “On Realms . . . everybody’s working . . . I feel more inclined to muck around because others are pushing the server forward.” (P6, 24NB, NZ European)

For players, the value of multiplayer is intimacy, not exposure. Minecraft often becomes a ’third place’ for existing relationships rather than a venue for building reputation or engaging with large anonymous communities.

Minecraft Roles and the Social Division of Labour

Small-group play may sometimes involve complementary roles towards the maintanence and survival on the server (in-game). For example, players may choose to adopt specific miners, builder, or redstoner roles, where they may best complement other players by playing in styles that they may prefer. This transforms diverse preferences of players for specific play styles into mutual dependence and social recognition. Players who identify as ’not good’ through combat, for example, may still find belonging through valued contributions in other specialised areas.

> “I am very good at digging very big holes . . . getting a lot of ore . . . I can give to everyone else.” (P1, 22F, NZ European)

> “I’m always trying to build up everyone’s armor so people can stop dying, please.” (P2, 23F, NZ European)

> “I do a lot of the technology mods . . . linking everything up together.” (P3, 24M, NZ European)

> “I’m not terrible at mining . . . my girlfriend takes the redstone; I like to let her do that.” (P7, 23F, NZ European)

In intimate server environments, non-combat and non-meta strengths become socially central. The division of labour provides identity and contribution pathways that broaden players’ participation and sense of belonging towards the server.

Discussion

In synthesising these patterns, we posit that participants’ multiplayer experiences are fundamentally shaped by how they negotiate risk, competence, moderation, intimacy, and collaborative labour within their respective online Minecraft environments. Across the dataset, players actively curate their participation to minimise social harm, gravitating toward protective infrastructures such as Realms, whitelists, and claim systems that allow them to regulate exposure and maintain relational safety. This risk-sensitivity intersects with playesr' perceived anxieties towards their own competence at the game. For example, the evolving combat metas lead some players to feel ill-equipped or outpaced, specifically when playing on large, public servers, which may cause players to withdraw into smaller, more private servers with friends. Server moderation also appears discourage some players from playing on large servers: a single punitive or neglected incident can redirect an entire multiplayer trajectory, whereas consistent, preventative moderation sustains engagement by reinforcing trust. Within these safer enclaves, multiplayer play functions less as a stage for public recognition and more as a site of intimate sociality, where relationships drive motivation and shape the very meaning of participation. Finally, the collaborative division of labour that unfolds in these close-knit groups, such as miners, builders, redstoners, resource-gatherers, not only supports collective progress but also provides participants with valued identities, creating low-conflict interdependence that binds groups together. Taken together, these dynamics illustrate how players reconfigure Minecraft's open-world multiplayer into socially coherent micro-communities, grounding their engagement in safety, mutual care, and meaningful contribution rather than competition or visibility. Collectively, these findings portray a player ecology that values intimacy, prevention, and role-based contribution over exposure, competition, and reputational play. Participants generally prefer to participate in smaller, more intimate multiplayer servers, that are safer, with more relationally meaningful environments.

General Discussion

This research set out to examine how multiplayer Minecraft enables and constrains the transmission of attitudes through player-to-player interactions, situated within broader sociotechnical accounts of online platforms. Prior work has described how decentralised, participatory ecosystems can facilitate both pro-social network formation and, in other cases, the spread of extremist norms, depending on how platform affordances intersect with community governance and user motivations (Heslep & Berge, 2024; Risius et al., 2024; Shaw, 2023; Wells et al., 2024). The present work contributes to this literature by developing a method for telemetry-derived social network analyses and investigating qualitative accounts of how players manage safety, competence, and social relationships.

Integrating Relational Structure and Player Experience

Study 2 shows how welcoming, warm social relationships a primary reason for players to engage with a Minecraft server. Often these may be for playing with friends, and are part of larger, multidimensional and multimodal channels for social interaction. However, the notion that players are attracted to welcoming and friendly online spaces highlights a possible mechanism for the reinforcement of problematic attitudes (i.e., extremist attitudes) through apparent supportive echo chambers and the guise of social support: if Player X harbours feelings of injustice and prejudice against certain members of the community, an online environment that shares and validates Player X's views would be seen as welcoming and supportive to Player X.

This has direct consequences for how attitudes may diffuse. Players tend to assemble in small, predictable, and relationally oriented micro-communities that limit exposure to large, anonymous contexts. Such spaces narrow the range of possible influence pathways, resulting in denser, inward-facing relational networks. While these structures may enhance stability and belonging, they also create distinct zones where norms become strongly reinforced through repeated co-presence. Conversely, players with differing views who would have otherwise occupied influential or bridging positions in more open networks may decline to participate in these closed environments altogether. As a result, this process self-selection into more intimate groups may lead to the formation of stronger echo chambers of extremist viewpoints.

Here, we propose that the telemetry mining and Social Network Analysis framework demonstrated in Study 1 provides an area for integration. If we can examine how player-to-player interactions on a server gradually lead towards the adoption of more extremist viewpoints, we now have the computational tools and software infrastructure to observe and model the formation of these echo chambers as they develop.

The telemetry-based approach introduced in this research also offers a scalable and replicable framework for future work. Location tracking, temporal binning, and proximity-based edge construction establish a foundation for modelling diffusion processes in naturalistic gameplay settings. While this work focused on spatial co-presence, the framework could be expanded to incorporate communication logs, collaborative actions, or longitudinal network shifts. Combining these data streams could enable more powerful models of attitudinal influence, including the prediction of high-risk configurations such as isolated clusters or over-centralised networks.

At a broader methodological level, this research underscores the importance of ecologically valid measures when studying digital social interactions. By grounding analyses in actual in-game behaviour rather than retrospective accounts alone, researchers can better capture the contingencies and fluidity of social interactions in online games.

Practical Applications

The insights generated by this program of research have relevance for the design of interventions, server governance, and harm-prevention strategies. Small, stable group structures appear to align well with players' preferences for safety and relational coherence. Intervention designers may leverage these insights by structuring activities around tight-knit groups, fostering role differentiation, and identifying individuals who organically serve as bridges or stabilising influences. Prior work has shown that influential players can boost engagement and retention in multiplayer environments (Seif El-Nasr et al., 2021), suggesting that carefully identifying and supporting these figures may enhance intervention outcomes.

Server governance also emerges as a critical factor in shaping relational environments. Participants' avoidance of large servers due to perceived or actual moderation lapses indicates that consistent, transparent, and preventative governance is essential for maintaining healthy community dynamics. Telemetry-based monitoring could assist administrators by identifying early indicators of fragmentation, isolation of new players, or dominance by a small number of central actors. This could help server moderation efforts in providing a better understanding of when toxic behaviours start to emerge on the server, allowing for preventative action before it escalates into more extreme attitude reinforcement.

Finally, the findings inform strategies for mitigating extremist exploitation. While most players actively avoid high-risk environments, pockets of vulnerability remain where governance is weak and relational structures become insular. Monitoring structural indicators such as high edge weights within closed clusters, sparse bridging connections, or unusual centrality patterns could help identify communities susceptible to harmful norm formation. Prevention efforts may benefit from strengthening alternative relational structures that promote belonging, oversight, and positive engagement.

Limitations

Our research presents several limitations. The network examined in this research was derived from a specific intervention context with a small active population. Larger public servers may produce substantially different network patterns. The proximity measure used here does not capture communication content, conflict, or collaborative action, all of which may shape attitude contagion processes on Minecraft. Nevertheless, we show the possibility of how these methods could be applied and scaled to larger servers and wider contexts.

Secondly, the qualitative sample was small and demographically homogeneous, with limited representation of players engaged in competitive or high-risk multiplayer spaces. Our sample also comprsed adult players, who may have a substantially different online experience and motivations from younger (adolescent) players, who may be navigating issues with identity formation and may be more susceptible or vulnerable to extremist attitudes.

Finally, no longitudinal attitudinal measures were collected, limiting claims about actual attitude change over time. Given the telemetry framework in Study 1, longitudinal research integrating behavioural (telemetric) and attitudinal (survey) measures would enable direct modelling of attitude change, for a much stronger measure of attitude contagion over time. Follow-up intervention research could also experimentally manipulate group structure, moderation style, and the placement of influential actors to identify causal levers for enhancing pro-social outcomes.

Conclusion

This research demonstrates that understanding attitude contagion in multiplayer Minecraft can benefit from paying attention to both the relational structures captured through telemetry and the lived experiences that shape participation choices. Players actively curate their social environments, forming small, predictable groups that balance safety, competence, and meaningful contribution. These choices give rise to specific network structures that both constrain and enable pathways for influence. By illuminating these structures and the motivations that sustain them, this work provides a foundation for designing healthier multiplayer environments and for anticipating how harmful norms might take hold in particular configurations. Minecraft can thus become a lens through which broader sociotechnical processes governing online sociality, influence, and community resilience, can be better understood.

References

Amarasingam, A. (2024, February). *Belonging is just a click away: Extremism, radicalisation, and the role of online communities* (J. Busher, L. Malkki, & S. Marsden, Eds.).

Borgatti, S. P., Mehra, A., Brass, D. J., & Labianca, G. (2009). Network analysis in the social sciences. *Science*, *323*(5916), 892–895. https://doi.org/10.1126/science.1165821

Brashears, B. N. (2024). *Learning in minecraft: A new paradigm for psychological research* [Doctoral dissertation, The University of Western Ontario (Canada)].

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, *3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Csárdi, G., Nepusz, T., Traag, V., Horvát, S., Zanini, F., Noom, D., & Müller, K. (2025). *igraph: Network analysis and visualization in r* [R package version 1.5.0]. https://doi.org/10.5281/zenodo.7682609

Gagandeep. (2025, July). *Playing with hate: How far-right extremists use minecraft to gamify radicalisation* [Accessed 2026-01-21]. https://gnet-research.org/2025/07/02/playing-with-hate-how-far-right-extremists-use-minecraft-to-gamify-radicalisation/

Heslep, D. G., & Berge, P. (2024). Mapping discord's darkside: Distributed hate networks on disboard. *New Media & Society*, *26*(1), 534–555. https://doi.org/10.1177/14614448211062548

Jagannath, K., Salen, K., & Slovàk, P. (2020). "(we) can talk it out...": Designing for promoting conflict-resolution skills in youth on a moderated minecraft server. *Proc. ACM Hum.-Comput. Interact.*, *4*(CSCW1). https://doi.org/10.1145/3392855

Logan, A., LaCasse, P., & Lunday, B. (2023). Social network analysis of twitter interactions: A directed multilayer network approach. *Social Network Analysis and Mining*, *13*. https://doi.org/10.1007/s13278-023-01063-2

Marin, A., & Wellman, B. (2011). Social network analysis: An introduction. *The SAGE handbook of social network analysis*, *11*, 25.

Menini, S., Moretti, G., Corazza, M., Cabrio, E., Tonelli, S., & Villata, S. (2019, August). A system to monitor cyberbullying based on message classification and social network analysis. In S. T. Roberts, J. Tetreault, V. Prabhakaran, & Z. Waseem (Eds.), *Proceedings of the third workshop on abusive language online* (pp. 105–110). Association for Computational Linguistics. https://doi.org/10.18653/v1/W19-3511

Pirker, J. (2018, January). Social network analysis in games user research. In *Games user research*. Oxford University Press. https://doi.org/10.1093/oso/9780198794844.003.0029

Prell, C., & Schaefer, D. (2024). Introducing social network analysis. *The Sage Handbook of Social Network Analysis:* 19–31.

R Core Team. (2022). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. https://www.R-project.org/

Research launch on gaming and violent extremism. (2022) (Accessed 2026-01-21). United Nations Office of Counter-Terrorism (UNOCT). https://www.un.org/counterterrorism/sites/default/files/221005_research_launch_on_gaming_ve.pdf

Risius, M., Blasiak, K. M., Wibisono, S., & Louis, W. R. (2024). The digital augmentation of extremism: Reviewing and guiding online extremism research from a sociotechnical perspective. *Information Systems Journal*, *34*(3), 931–963. https://doi.org/https://doi.org/10.1111/isj.12454

Rodrigues, L. C., & Mustaro, P. N. (2007). Social network analysis of virtual communities in online games. *Proceedings of the Iadis International Conference e-Society*, *2007*.

Seif El-Nasr, M., Dinh, T. H. N., Canossa, A., & Drachen, A. (2021, October). Case study: Social network analysis applied to in-game communities to identify key social players. In *Game data science*. Oxford University Press. https://doi.org/10.1093/oso/9780192897879.003.0012

Shaw, A. (2023). Social media, extremism, and radicalization. *Science Advances*, *9*(35), eadk2031. https://doi.org/10.1126/sciadv.adk2031

Slattery, E. J., Lehane, P., Butler, D., O'Leary, M., & Marshall, K. (2025). Assessing the benefits of digital game-based learning with minecraft in children, adolescents and young adults: A broad systematic review [e70035 RoE-2024-07-1339.R3]. *Review of Education*, *13*(1), e70035. https://doi.org/https://doi.org/10.1002/rev3.70035

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications.* Cambridge University Press.

Wells, G., Romhanyi, A., Reitman, J. G., Gardner, R., Squire, K., & Steinkuehler, C. (2024). Right-wing extremism in mainstream games: A review of the literature. *Games and Culture*, *19*(4), 469–492. https://doi.org/10.1177/15554120231167214

Yang, S. (2024). Social network analysis: An introduction. In S. Yang (Ed.), *Social network analysis in action: Basic methods and applications* (pp. 1–13). Springer International Publishing. https://doi.org/10.1007/978-3-031-66661-2_1

Table 1

Participant Demographics

Participant ID	**Age**	**Gender Identity**	**Ethnicity**
P1	22	Female	New Zealand European
P2	23	Female	New Zealand European
P3	24	Male	New Zealand European
P4	22	Male	New Zealand European
P5	24	Non-binary	New Zealand European
P6	24	Non-binary	New Zealand European
P7	23	Female	Pākehā / New Zealand European

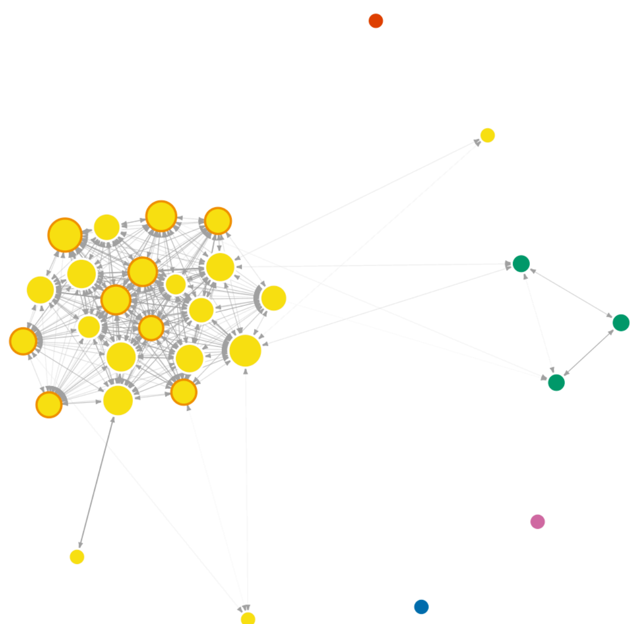

Figure 1

A visualisation of relative proximities between players. Thicker lines indicated larger edge weights (prolonged close interactions between to players relative to the time each has spent on the server), node sizes indicate PageRank centrality scores, and colour indicates detected communities by the Walktrap algorithm.

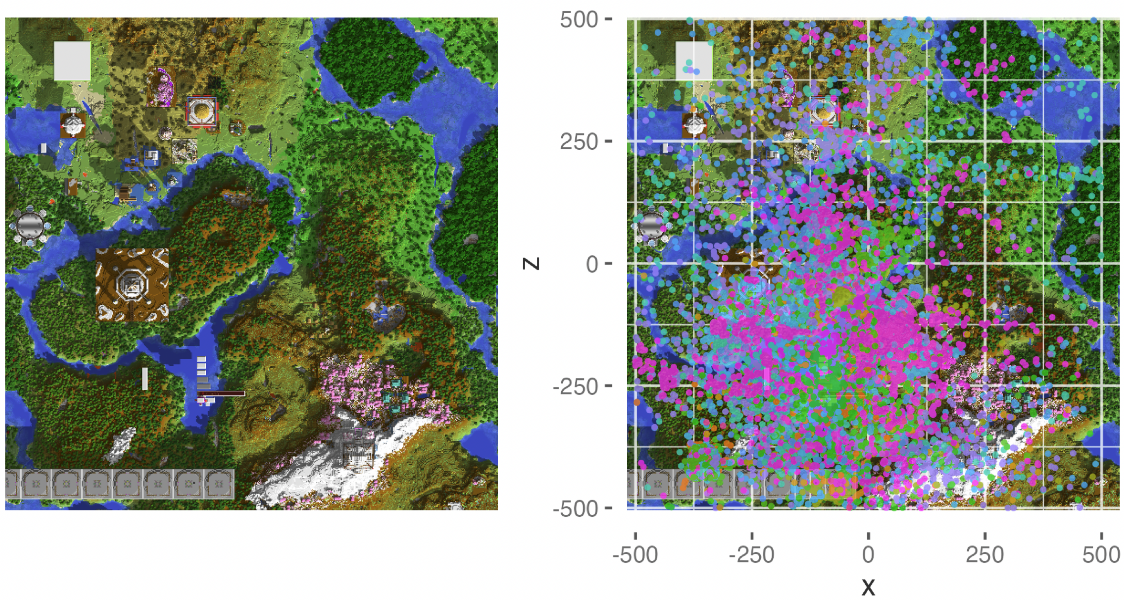

Figure 2

Left: an aerial view of a 1000 by 1000 block area centred around Spawn. Right: an overlay of active participation locations. Each coloured dot represents a player's location at each 30s bin.