

The Prosocial Turn in Video Games: A Twenty-Year Shift in the Motivations Games Simulate

Alexandre Javelle^{a,*}, Benjamin Dahan^a, Toni Zian Alabdeen Ali^a,
Valérien Chambon^{a,†}, Edgar Dubourg^{a,b,†}

^a Institut Jean Nicod, Département d'études cognitives, École normale supérieure, Université PSL, EHESS, CNRS,
75005 Paris, France

^b Paris School of AI (PSAI), Université PSL, 60 Rue Mazarine, 75006 Paris, France

[†] Co-last authors

* Corresponding author: ajavelle1@gmail.com

Abstract: As video games have grown from a niche pastime into a major cultural medium, they have reached broader audiences and transformed the social context of play, whether alone or with others. We ask whether the motivations engaged by the games themselves have changed over time. In Study 1, we used a large language model, validated against human coding, to annotate 11,359 commercially released games from 1993 to 2016 on 48 motivational ingredients, including status-seeking, revenge, affiliation, compassion, and curiosity. We normalized within each game to represent its relative motivational composition. Over this period, the average competitive index (the share of a game's motivational content devoted to competition, dominance, and aggression) fell from 41% to 26%, while the prosocial index (affiliation, warmth, and care) rose from 11% to 24%. This reallocation held within genres and franchises, in both single- and multiplayer titles, and under all 642,560 admissible definitions of the two indices. In Study 2, 4,498 players reported their favorite games and completed a Big Five measure. Higher Agreeableness predicted a balance favoring prosocial over agonistic content, driven primarily by lower agonistic content, whereas higher Extraversion, particularly Sociability, predicted a balance favoring agonistic content. This contrast suggests that a disposition towards social engagement is distinct from a preference for caring social content. Together, the studies document a “prosocial turn” in the motivational content of video games: a historical reallocation from competition and dominance towards affiliation, warmth, and care. The direction of this shift aligns with contemporary personality differences in content preference.

Keywords: video games; cultural evolution; prosociality; motivation; personality; large language model annotation

1. Introduction

For many players, playing video games is a deeply social activity. Contemporary games increasingly support interaction with other players (Granic et al., 2014), and online worlds can foster lasting friendships, social capital, and a sense of belonging (Cole & Griffiths, 2007; Domahidi et al., 2014; Trepte et al., 2012; Ferguson & Olson, 2013; Raith et al., 2021). Players also commonly identify social connection as a reason for playing online games (Yee, 2006; Yee et al., 2012). Although greater involvement in online gaming has sometimes been associated with

lower perceived social support, the direction of this relationship remains unclear (Kowert et al., 2014). Taken together, this literature portrays gaming not simply as a solitary pastime but as a social context in which people meet, cooperate, form relationships, and develop a sense of belonging.

Much of this work, however, concerns the *social context* of play – with whom players play, communicate, and form relationships – rather than the *motivational content* of games themselves: what games are about and which motivations they engage. As video games developed from a niche pastime into a global medium, networked technologies transformed this context by creating new opportunities for players to interact through games. Yet changes in how games are played need not imply corresponding changes in what games are about. These two dimensions are distinct and can vary independently: a competitive shooter played with friends is social in context but agonistic in content; conversely, a single-player game centered on tending a village and caring for its inhabitants is solitary in context but prosocial in content. The transformation of gaming’s social context therefore raises a parallel historical question about the games themselves: did the motivations they engage change as well? To our knowledge, no study has examined this question at sufficient scale and over a sufficiently long period to determine whether, and in what direction, such a shift occurred.

Such a change could matter for at least two reasons. First, the content of a popular medium may offer a window onto the preferences of its audience. Cultural artefacts partly reflect the psychological traits and concerns of the people who produce and consume them and can therefore serve as “cognitive fossils” of past preferences and motivations (Baumard et al., 2024). Their content has been used to trace historical changes ranging from the increasing prominence of cooperation in early modern fiction (Martins & Baumard, 2020) to changing displays of trustworthiness across six centuries of European portraiture (Safra et al., 2020). A historical shift in the motivational content of video games could similarly provide indirect evidence about the preferences to which the medium catered over time. By the same logic, differences among contemporary players should be reflected in the motivational content of the games they prefer.

Second, game content may influence players’ attitudes and behavior beyond play. Previous research suggests that the social framing of this content may matter: games organized around helping and cooperation have been linked to greater helping and empathy (Gentile et al., 2009; Greitemeyer & Osswald, 2010; Prot et al., 2014), whereas content centered on harm and interpersonal conflict has sometimes been associated with more aggressive responding and less prosocial behavior (Anderson et al., 2010). Although the magnitude and robustness of these effects remain debated (see Ferguson, 2015, for a critical assessment), even modest effects could matter when exposure is sustained and widespread.

We address these questions in two complementary studies. Study 1 tracks historical changes in the motivational content of 11,359 commercially released games between 1993 and 2016. Study 2 turns from historical change at the level of the medium to variation among contemporary players, asking, in an independent sample of 4,498 players, whether personality predicts the motivational content of the games they name as favorites. Together, the studies examine both

whether the motivational content of games has changed and whether the direction of that change aligns with personality-related differences in players' content preferences.

2. Study 1: The motivational content of games shifted from competitive to prosocial

2.1. Method

2.1.1 Data

We drew on the VGChartz catalog, which records commercially released video games together with their release year, genre, platform, and regional and global sales. Because a title released on multiple platforms appears once for each platform, we aggregated these platform-specific entries into a single game release, defined as a unique title-year combination. For each release, we retained a single content annotation, averaging the ingredient scores across platform entries when these differed (6.6% of the resulting scores are therefore non-integer). We summed sales across platforms and retained the number of platforms as a covariate. Remakes and reissues sharing a title but released in different years were treated as distinct releases. To ensure adequate annual coverage, we restricted the analysis to years containing at least 30 titles, yielding a contiguous period from 1993 to 2016. The final corpus comprised 11,359 game releases with complete content annotations and associated metadata.

2.1.2 Motivational Ingredient Framework

To characterize the motivational content of games, we drew on a framework that represents cultural products as combinations of the evolved human motivations they engage (Dubourg, Thouzeau, et al., 2026). The framework builds on the fundamental-motives tradition in cognitive science, according to which human motivation comprises a set of distinct systems shaped to address recurrent adaptive problems (Kenrick et al., 2010). These include acquiring food and water, avoiding predators and pathogens, gaining status, forming and defending coalitions, attracting and retaining mates, and caring for kin (Kenrick et al., 2010; Schaller et al., 2017; Neel et al., 2016). Each system is proposed to have its own eliciting conditions and associated rewards and to operate with some degree of independence from the others.

From this perspective, cultural products owe part of their appeal to their capacity to engage these motivational systems (Dubourg, Thouzeau, et al., 2026; Dubourg & Baumard, 2022). A work can therefore be characterized not only by its genre but also by the particular set of motivations it engages – its motivational “recipe”. The framework distinguishes 48 such ingredients, spanning basic drives (e.g., hunger, thirst, Sleepiness), epistemic motivations (e.g., curiosity, explanation-seeking), and social, moral, and status-related concerns (e.g., affiliation, warmth, compassion, status-seeking, revenge, moral blame). Works within the same genre may have different motivational recipes, whereas works from different genres may share many of the same ingredients.

Video games are particularly well-suited to this framework because they can be described in terms of the opportunities they offer players to pursue or fulfill different motivations (Dubourg & Chambon, 2025). A survival game such as *Don't Starve*, for example, engages motivations

related to hunger, thirst, and fear of predators; a farming-and-village simulation such as *Stardew Valley* engages affiliation, friendship, warmth, and parental care; a competitive shooter such as *Call of Duty* engages status-seeking, coalitional affiliation, and aggression. Individual ingredients can therefore recur across otherwise very different games: curiosity, for example, may be engaged by both a puzzle game and an open-world adventure. Conversely, a single game may combine numerous ingredients, as when a role-playing game engages curiosity, friendship, moral choice, and status simultaneously. This common motivational vocabulary allows games to be compared across genres and provides a basis for examining whether the relative prevalence of these ingredients has changed over time.

2.1.3 Annotation

We annotated games using a large language model. Recent studies suggest that, for some annotation and classification tasks, large language models can achieve performance comparable to or better than that of trained human coders, often at substantially lower cost and with greater consistency (Gilardi et al., 2023; Rathje et al., 2024; see also: Pei et al., 2023; Bongini et al., 2023). This LLM-based approach has also been used and validated in previous applications of the motivational-ingredient taxonomy (Dubourg et al., 2024; Dubourg, Thouzeau, et al., 2026; Dubourg & Thouzeau, 2026).

Each game was rated on all 48 ingredients using OpenAI’s GPT-5.4. The model was prompted to act as an expert in psychology and video game analysis and to treat the experience of the human player – not that of the in-game character – as the unit of analysis. For each ingredient, it rated both its presence and its importance to the player’s experience on an anchored scale ranging from 0 to 6. Selected anchors included 0 = Absent, indicating that the game offered no opportunity to fulfill the motivation; 3 = Supported, indicating that occasional moments or mechanics allowed the player to fulfill it; 4 = Emphasized, indicating that fulfilling the motivation was integral to the experience; and 6 = Defining, indicating that removing it would fundamentally alter the meaning of the game. Scores were restricted to motivations directly engaged through the player’s own perceptual, affective, or bodily experience, regardless of whether the player consciously pursues them. For each ingredient, the model generated a one- to two-sentence justification before assigning its score.

Each game was annotated in a single call covering all 48 ingredients, with the definition of each ingredient included in the prompt. The model returned a structured JSON object containing a justification and a score for each ingredient. If the model lacked sufficient confidence to score a particular ingredient for a title, that value was recorded as missing. All annotations were processed through the OpenAI Batch API. Missing scores were extremely rare: 41 of the 545,232 ingredient scores in the final corpus (0.008%) were missing, affecting 18 of the 11,359 releases.

2.1.4 Validation

Temporal trends derived from LLM annotations can fail in two ways. The first is a false negative: undirected scoring noise (e.g., hallucination) may attenuate the association with release year, causing a genuine trend to be missed. Although increasing the size of the corpus does not eliminate this measurement error, it reduces the minimum annotation quality required to detect

the effect. The second is a false positive: annotation errors that vary systematically with release year may produce a spurious temporal trend. Increasing the corpus size does not reduce such bias. Instead, it narrows the confidence interval around the biased estimate, potentially making even a small systematic error appear conclusive (Meng, 2018).

To evaluate these two risks, we followed the design-specific validation procedure developed by Dubourg and Thouzeau (2026). Rather than applying a generic agreement threshold, this procedure derives the level of annotation quality and systematic bias that the specific design can tolerate. One author independently hand-coded 99 games released between 1994 and 2016 on all 21 ingredients composing the two indices (12 prosocial, 9 competitive), using the same 0-6 scale and definitions provided to the model. The coder did not have access to the model's scores. Because the relevant validation thresholds depend on the effect size and sample size obtained in the study, we report the complete validation analysis after presenting the main results. In addition, we used the audit to correct the corpus-level estimate for annotation error using prediction-powered inference (Angelopoulos et al., 2023).

2.1.5 Analyses

We represented each game as a relative composition rather than as a set of absolute motivational scores. This choice was motivated by a general increase in motivational richness over time: the mean summed across all 48 ingredients rose from 54.3 in 1993 to 78.6 in 2016, corresponding to an estimated increase of 10.1 points per decade ($p < .001$). As a result, increases in the raw scores of individual ingredients may reflect a general expansion of motivational scope rather than a specific shift towards those ingredients. For illustration, *Tetris 2* (1993) received a summed score of 37, whereas *The Witcher 3: Wild Hunt* (2015) received a score of 159. *Tetris 2* is an abstract puzzle game that engages relatively few motivational systems beyond skill and mastery. By contrast, *The Witcher 3* is an open-world RPG whose narrative and mechanics engage a broad range of motivations, including romance, kinship, moral choice, status, fear, and revenge. This contrast illustrates the broader expansion of motivational scope associated with more narratively complex and systems-rich game design.

An increase in the absolute prosocial score would therefore be ambiguous: it could reflect either a specific increase in prosocial content or the general expansion of motivational richness observed across games. To distinguish between these possibilities, we normalized the 48 scores for each game so that they summed to one:

$$s_{ij} = x_{ij} / (\sum_k x_{ik})$$

where x_{ij} is the raw score of game i on ingredient j , and s_{ij} is the corresponding share of that game's total motivational content. Because these shares sum to one, an increase in the share of one ingredient must be offset by a decrease in the combined share of others. The resulting measure therefore captures changes in the relative composition of motivational content rather than changes in its overall level.

For the central analysis, we grouped a subset of the 48 ingredients into two functional indices defined on theoretical grounds. Ingredients were classified not according to their affective valence but according to their primary social function. The prosocial index (P) summed the

normalized shares of 12 ingredients associated with cooperation, attachment, and moral care: Warmth, Moral praise, Compassion, Affiliation, Friendship, Parental protection, Parental investment, Familial love, Moral pride, Guilt, Forgiveness, and Moral blame. The competitive index (C) summed the normalized shares of 9 ingredients associated with rank, dominance, and conflict: Fear of aggressors, Revenge, Status-seeking, Envy, Hatred, Coalitional affiliation, Anger, Pride, and Shame. The remaining 27 ingredients (e.g., Hunger, Curiosity, Skill acquisition) were assigned to neither pole and were analyzed individually. Both indices are computed per game and averaged across the games released each year, rather than pooled within years, which would implicitly weight years by motivational richness ($r = .58$ with the prosocial index).

Because our primary aim was to estimate historical change, the main analyses were conducted at the level of release year. Each index is reported on two scales: as a share of the game's summed 48-ingredient content, and as the mean per ingredient (dividing by 12 and 9), which places the two poles on a common scale. For each index taken separately, these are the same quantity expressed differently. We regressed the annual mean of each index on release year and scaled year so that the coefficients represent percentage-point changes per decade. To summarize the relative balance between the two poles in a single quantity, we also computed the prosocial share, Q , from the per-ingredient values of P and C :

$$Q = P / (P + C)$$

Unlike P and C taken separately, Q changes depending on whether the indices are divided by their number of ingredients, because the two poles contain 12 and 9 ingredients. Throughout the article, "index" refers to P or C , and "prosocial share" refers only to Q . We regressed (Q) on release year. We predicted that the prosocial index would increase while the competitive index decreased, producing an increase in Q .

Because analyses based on annual means may obscure heterogeneity among the games released within each year, we complemented them with three models estimated at the level of the 11,359 individual releases. First, we regressed each index on release year using ordinary least squares to determine whether the trend held at the level of individual games and to quantify the variance it explained without the smoothing introduced by annual averaging. Second, we added genre, number of platforms, and log-transformed global sales as covariates, thereby adjusting the year coefficient for temporal changes in the composition of the corpus along these dimensions. Third, we stacked the observations for the two indices and estimated a pole-by-year interaction, with standard errors clustered by title. This interaction provides a direct estimate of the difference between the temporal trends of the prosocial and competitive poles.

For descriptive purposes, we also regressed the share of each of the 48 ingredients separately on release year, applying the Benjamini-Hochberg correction for multiple testing. These analyses identify which motivations contributed to the overall change, but the primary weight rests on the two composite indices rather than on any individual ingredient. We then assessed the scope and robustness of the result using a series of alternative specifications, detailed in the Robustness section. We tested whether the temporal reallocation held within genres, within single-player and multiplayer titles, within franchises, and among best-selling games. We also examined sales-weighted estimates, changes across the full distribution of games, robustness to removing

each ingredient in turn, and every admissible alternative definition of the two indices. Finally, we assessed whether the ingredients' co-occurrence structure changed over time. All analyses used a fixed random seed and can be fully reproduced from the annotated corpus.

2.2 Results

2.2.1 *A directional reallocation across the 48 ingredients*

When examined separately as a function of release year, the 48 ingredient shares showed a structured, directional pattern of change rather than uniform drift. After Benjamini-Hochberg correction, 29 shares increased significantly, 14 decreased significantly, and 5 showed no reliable temporal trend. The largest increases concerned affiliative and caring motivations: Friendship (+1.04 points of relative share per decade), Warmth (+0.92), Compassion (+0.68) and Affiliation (+0.66), all $q < .001$. The largest declines concerned skill-focused, strategic and agonistic motivations: Skill acquisition (-1.69), Pride (-1.36), Anger (-1.12), Fear of aggressors (-1.01), Mental time travel (-0.89) and Status-seeking (-0.79), all $q < .001$ (Figure 1, panel A). Not every increasing ingredient was prosocial or moral. Romantic affiliation (+0.50) and Desire (+0.44), for example, also gained relative share, both $q < .001$. We therefore interpret the overall pattern as a reallocation away from a predominantly skill-focused and agonistic profile towards more affiliative and caring content. Because several non-prosocial ingredients also increased, however, our central inference rests on the specific contrast between the prosocial and competitive indices rather than on a general moralization of games.

2.2.2 *The convergence of the two indices*

The two functional indices quantify the reallocation. At the level of individual games, the two poles were strongly opposed ($r = -.68$, 95% CI [-.69, -.67], $N = 11,359$; Figure 1, panel B). In the primary year-level regressions, the prosocial index increased by an estimated 5.44 percentage points per decade, or 0.45 points per ingredient (R-squared = .91, $p < .001$), whereas the competitive index decreased by 6.34 points per decade, or 0.70 points per ingredient (R-squared = .90, $p < .001$). Averaged over the first three years of the corpus (1993-1995, $N = 379$), the competitive index accounted for 38.60% of the average game's total motivational content, compared with 12.65% for the prosocial index, a gap of about 26.0 percentage points. Over the final three years (2014-2016, $N = 864$), the competitive index had fallen to 27.36%, and the prosocial index had risen to 22.98%, narrowing the gap to about 4.4 percentage points. On the per-ingredient scale, the same quantities were 4.29% and 1.05% at the start of the period, and 3.04% and 1.92% at its end (Figure 1, panel C). Read at the endpoints of the observation window, the fitted values placed the prosocial index at 11.03% in 1993 and 23.55% in 2016, and the competitive index at 40.53% and 25.94%, reducing the fitted gap from 29.5 to 2.4 percentage points. Games therefore moved from a marked predominance of competitive over prosocial content towards a substantially reduced gap between the two poles. They did not reverse order within the observed period, however: even in 2016, the average competitive index remained slightly larger than the prosocial index.

Using the game-level data, we tested this convergence directly in a single model by stacking the two poles and regressing index value on release year, pole, and their interaction, with standard errors clustered by title. The pole-by-year interaction was +1.31 percentage points per decade

(95% CI [1.24, 1.38], $p < .001$), indicating that the temporal slope of the prosocial index was about 1.3 points more positive than that of the competitive index. This interaction is one of the quantities for which the two scales do not coincide: on the summed scale the same model returned +13.26 points per decade (95% CI [12.56, 13.97], $p < .001$). It provides the formal test of the convergence shown in Figure 1, panel C.




Figure 1. Temporal evolution of motivational ingredient shares (1993-2016). (a) The annual mean relative share of each motivational ingredient, grouped according to whether its temporal trend was significantly increasing (left), significantly decreasing (middle), or non-significant after Benjamini-Hochberg correction (right). Labeled trajectories correspond to the five ingredients with the largest temporal trends within each group; Romantic affiliation is highlighted instead of Moral praise (rank 5) to illustrate that not all increasing ingredients were prosocial. Note that the y-axes differ across panels. (b) The 11,359 games projected into the space defined by the two indices, with the linear fit in black and a selection of well-known titles annotated. (c) Annual mean prosocial and competitive indices, as mean intensity per constituent ingredient, with fitted trends and 95% confidence intervals.

2.2.3 The trend holds at the level of individual games

The temporal pattern also held at the level of individual game releases. In unadjusted models estimated across all 11,359 releases, the prosocial index increased by 6.0 percentage points per decade (0.50 per ingredient) and the competitive index decreased by 7.3 percentage points per decade (0.81 per ingredient). These slopes were slightly larger in magnitude than the year-level estimates, but the game-level observations showed substantially greater variability: release year explained approximately 8% (prosocial) and 10% (competitive) of the variance at the game level, compared with approximately 90% in the corresponding year-level models. The high year-level R-squared therefore reflects the smoothing produced by annual aggregation rather than a

stronger relationship at the level of individual games. Consistent with this interpretation, a variance decomposition indicated that only about 9% of the variance in the prosocial index and 11% in the competitive index occurred between release years, the remainder occurring between games released within the same year. Adjusting for genre, number of platforms, and log-transformed global sales attenuated but did not eliminate the temporal associations. The adjusted prosocial index increased by 3.9 percentage points per decade (0.32 per ingredient; $R^2 = .42$, $p < .001$), whereas the competitive index decreased by 4.5 percentage points per decade (0.49 per ingredient; $R^2 = .52$, $p < .001$). Finally, adding a quadratic term for release year provided no evidence of a departure from linearity for either pole (prosocial, $p = .21$; competitive, $p = .45$). Thus, within the period examined, we found no evidence that the reallocation was driven by a recent acceleration.

2.2.4 Annotation validation

As described in the Methods, one author, blind to the model's scores, hand-coded 99 games released between 1994 and 2016 on all 21 ingredients used to construct the two indices. The sample was drawn at random within each release year, with the number of entries selected from each year proportional to the number of titles released in that year. This procedure yielded 100 entries. Two entries represented the same game, *Backyard NBA Basketball*, listed separately for two platforms. We therefore coded the game only once, yielding a final sample of 99 distinct games. This audit served two distinct purposes: first, to establish whether annotation accuracy was sufficient to prevent random error from obscuring a genuine trend (false-negative risk); and second, to estimate how much of the observed trend could be attributable to annotation error systematically correlated with release year (false-positive risk). The thresholds for sufficient accuracy and tolerable systematic bias are design-specific and depend on the observed effect size and sample size (Dubourg & Thouzeau, 2026).

The present design is a game-level OLS with release year as the sole predictor, a standardized beta of .31 for the prosocial share, and $N = 11,359$ releases. For this design, the minimum true annotation quality required to retain 80% power was $\rho^* = .09$ (i.e., the lowest correlation between LLM annotations and the true construct at which the effect would still be detected with 80% power). With 99 paired human-model annotations in the audit, this requirement corresponded to an observed correlation of $\rho = .25$. The required annotation quality decreases as corpus size increases: for an effect of the same magnitude, ρ would be .65 at $N = 200$, .41 at $N = 500$, .29 at $N = 1,000$ and .13 at $N = 5,000$. Thus, the large size of our corpus makes the design highly tolerant of random annotation noise.

Observed agreement exceeded the audit threshold by a wide margin. Human and model annotations correlated at $\rho = .82$ for the prosocial index, .76 for the competitive index, and .87 for the prosocial share. Pooling the 99-by-21 ingredient scores yielded $\rho = .73$. This value partly reflects differences in mean score across ingredients (e.g. games generally scored low on Envy and high on Pride), but after centring scores within each ingredient, the pooled within-ingredient, between-game correlation remained $r = .65$. Across the 21 individual ingredients, median agreement was $\rho = .65$, and only four correlations fell below .50: Forgiveness (.36), Pride (.43), Shame (.48), Affiliation (.49). Dropping these four ingredients left the estimated temporal increase in the prosocial share within the audit sample virtually

unchanged and closely matched across annotation methods (human-coded, +13.8 points per decade, $p = .001$; model-coded, +13.4, $p = .004$). The observed trend therefore does not depend on the ingredients with the weakest human-model agreement.

To test for systematic temporal bias, we regressed the model-minus-human discrepancy on release year separately for each index. The estimated drift was +0.09 percentage points per decade for the prosocial index (95% CI [-0.14, +0.31], $p = .44$), -0.05 for the competitive index ([-0.27, +0.18], $p = .69$), and +1.15 for the prosocial share ([-2.6, +4.9], $p = .55$). None was distinguishable from zero. At $N = 11,359$, a very small bias ($\gamma = .018$ standardized units, or 0.69 points of prosocial share per decade) would be sufficient to produce a significant slope under a true null, compared with the observed slope of 11.59 points per decade. Detecting a bias this small with 80% power would take approximately 5,600 hand-coded games, and the present 99-game audit had a power of only .065 against it. What the audit can rule out is a much larger bias accounting for approximately half or more of the reported effect: its minimum detectable drift was 5.4 points per decade, with simulations yielding a comparable threshold of approximately 46% of the observed effect. Thus, the audit excludes large systematic discrepancies but cannot rule out the small temporal bias that could achieve statistical significance in the full corpus. No realistically sized hand-coding audit could provide that degree of resolution. We therefore conducted further tests that do not depend on the audit's ability to detect such a small bias: subsampling the model-annotated corpus commensurate with the audit's detection limit, testing whether the temporal reallocation could be recovered from the human annotations alone and correcting the corpus-level estimate for annotation error.

The audit constrains the year-correlated component of annotation error. However, the magnitude of bias capable of affecting inference at $N = 11,359$ is too small to be detected reliably from 99 hand-coded games. A complementary strategy operates on the other term of the same inequality (Dubourg & Thouzeau, 2026): reducing the size of the analyzed corpus rather than increasing the size of the audit. The minimum bias required to produce a significant temporal slope under a true null is not fixed. As the corpus shrinks, the standard error of the year coefficient increases, and a larger undetected bias is therefore needed to produce a false-positive result. In the full corpus, this critical bias is only 0.69 points of prosocial share per decade, which is below the audit's detection limit. At $N = 194$, however, it rises to 5.42 points per decade, exceeding the upper bound of the audit's 95% confidence interval for the observed drift (4.94). Thus, at $N = 194$, any year-correlated bias still compatible with the audit is too small, by itself, to produce a significant temporal trend under a true null.

The subsampling strategy has a lower as well as an upper bound, because shrinking the corpus also increases the annotation quality required to avoid a false negative. Given the agreement observed across the 99 audit pairs ($\rho = .87$), the design continued to satisfy this requirement at $N = 130$ but not at $N = 100$, where an agreement of $\rho = .93$ would have been required. Combined with the false-positive ceiling derived above, this defined an admissible range of 130 to 194 games. We used $N = 193$, near the upper end. We then re-estimated the regression of prosocial share on release year in 2,000 random subsamples of 193 releases drawn from the full corpus. The median slope was +11.61 percentage points per decade, with a central 95% empirical interval of [6.58, 16.45], closely matching the full-corpus estimate of +11.59. The slope was positive in all 2,000 subsamples and statistically significant at $p < .05$ in 99.5% of them.

Results were similar at the lower bound of the admissible range, $N = 130$: the median slope was +11.63 points per decade, with a central 95% empirical interval of [6.39, 17.15], and 98.7% of subsamples produced a significant slope. Subsampling does not itself reduce annotation bias: each subsample inherits any systematic bias present in the model's annotation. What changes is the magnitude of bias required to produce a significant result under a true null. At sample sizes between 130 and 194, that critical magnitude exceeds the upper bound on year-correlated bias permitted by the audit. The persistence of the trend within this range therefore cannot be explained solely by a systematic drift that the validation failed to detect.

Regressing annotation discrepancies on release year provides an indirect test of systematic model bias. A more direct test removes the model annotations from the loop and asks whether the same temporal reallocation is present in the human-coded data. The bias of concern is one that would make a true null association appear significant. The human annotations provide direct evidence against that null. Across the 99 audited games, the human-coded share increased by 11.9 percentage points per decade (95% CI [+4.8, +18.9], $p = .001$), compared with 13.0 points for the model annotations of the same games ($p < .001$). Thus, a coder who was blind to all model outputs recovered the temporal reallocation independently, with an estimated slope closely matching that produced by the model.

The separate indices showed a less symmetrical pattern in this smaller sample. The prosocial index increased under both annotation methods (human-coded: +0.57 scale points per decade, $p = .001$; model-coded: +0.66, $p < .001$), whereas the competitive index showed no reliable temporal change under either method (human-coded: +0.01, $p = .96$; model-coded, -0.04, $p = .81$). The 99-game audit lacked sufficient power to detect the competitive decline on the raw 0-6 scale; in this subset, the overall reallocation was therefore recovered through the prosocial share measure rather than through a separately significant decline in the competitive index. This limitation is a property of the validation subsample, not of the full corpus, in which the competitive decline is itself significant at $N = 11,359$.

We conducted two additional checks. First, the model's overall scoring differences favored the competitive rather than the prosocial pole. Relative to the human coder, the model assigned lower scores to prosocial ingredients (mean model-minus-human error = -0.22 on the 0-6 scale) and higher scores to competitive ingredients (+0.29). Thus, the model portrayed games as less prosocial and more competitive overall than did the human coder. These overall scoring differences cannot by themselves explain the observed temporal shift toward the prosocial pole; they could do so only if they varied systematically with release year. Second, we tested for precisely such temporal variation. When the audit sample was split in 2005, pooled human-model agreement was virtually identical in the earlier and later periods ($\rho = .725$ vs $.724$), and the mean signed difference remained close to zero in both (-0.04 vs +0.02). We therefore found no indication that the estimated reallocation resulted from an era-dependent decline in agreement or a systematic temporal drift in model scores.

Finally, the audit can be used not only to test for bias but to correct for it. Prediction-powered inference (Angelopoulos et al., 2023) takes the slope estimated on the full corpus and subtracts the model's drift measured in the audit (i.e., how much faster the model's scores rise over time than the human scores). The confidence interval is widened accordingly, to reflect the

uncertainty of this drift. Because the audited games were drawn at random from the corpus, the corrected estimate is valid whether or not the model is biased, even if the bias is too small for the audit to detect. The corrected slope of the prosocial share was +10.44 percentage points per decade (95% CI [+6.64, +14.24], $p < .001$), close to the uncorrected +11.59, and its interval was about 1.8 times narrower than with human annotations alone. The temporal reallocation toward prosocial content therefore holds once annotation error is corrected.

2.3 Robustness

All robustness analyses address the same central claim: over the period examined, games reallocated their motivational content away from competitive and towards prosocial ingredients. Rather than treating each analysis as a separate finding, we used them to test whether this reallocation persisted under a series of plausible alternative explanations. Specifically, we examined whether the pattern could be attributed to changes in genre composition, the balance between single-player and multiplayer titles, the franchises represented in the corpus, differential commercial success, a small subset of unusually prosocial games, the particular ingredients assigned to the two indices, or changes in how motivations combine (i.e. the ingredients' co-occurrence structure). The reallocation remained evident across all these tests (Table 1).

Check	Specification	Result	Alternative ruled out
Within genres	index $\sim$ year, per genre and with genre fixed effects	+3.7 / -4.2 pts per decade; significant in 11/12 and 10/12 genres	temporal changes in genre composition
Single-player vs multiplayer	index $\sim$ year by mode, and year-by-mode interaction	present in both modes (+4.5 / -4.8 solo; +5.2 / -6.1 multiplayer)	the increasing prevalence of multiplayer games
Within franchises	index $\sim$ year + franchise fixed effects	+2.5 / -2.2 pts per decade (773 franchises, 5,673 releases)	changes in which franchises are represented over time
Best-sellers and sales weighting	top-N titles per year; observations weighted by sales	prosocial rise significant at every Top-N; weighted +4.0 / -2.5	a trend driven primarily by low-selling or niche titles
Across the distribution	quantile regression of each index on year	10th +3.6, 50th +7.9, 90th +3.8 pts per decade	a shift confined to one segment of the game distribution
Any single ingredient	leave-one-ingredient-out on each index	every slope keeps sign and significance (min +4.4)	a result driven by a single ingredient
Index definition	all 642,560 admissible pole pairings	share rises significantly in 100%. Range +6.4 to +17.0 pts per decade.	dependence on the particular assignment of ingredients to the two poles
Co-occurrence structure	Mantel test, early vs late correlation matrix	$r = .93$, $p < .001$	a reorganization of the motivational architecture (i.e. the ingredients' co-occurrence structure)

Table 1. Overview of the robustness analyses. All coefficients are expressed in percentage points of total motivational content per decade. The per-ingredient values reported alongside them in the Results follow by division by 12 (prosocial) and 9 (competitive).

2.3.1 Within genres

The temporal reallocation occurred within genre rather than resulting from changes in the genre composition of the corpus. When we re-estimated the two index trends separately within each of the 12 genres, the prosocial index increased significantly in 11 of the 12 genres and the competitive index decreased significantly in 10. The prosocial increase was nonsignificant only for Strategy games, whereas the competitive decrease was nonsignificant only for Strategy and Role-Playing games. In a pooled model with genre fixed effects, in which the year coefficient was identified only from temporal variation within genres, the prosocial index increased by 3.7 percentage points per decade (0.31 per ingredient) and the competitive index decreased by 4.2 points per decade (0.46 per ingredient), both $p < .001$. Because each slope is estimated within a single genre, the pattern cannot be attributed to changes in the relative prevalence of genres characterized by more prosocial or more competitive content (Figure 2).




Figure 2. Within-genre temporal trends in prosocial and competitive indices (1993-2016). Each panel shows the annual mean prosocial (purple) and competitive (orange) indices, as mean intensity per constituent ingredient, within one of the 12 genres. Lines show OLS fits with 95% confidence bands. Asterisks indicate the significance of the temporal trend (* $p < .05$, *** $p < .001$); NS, non-significant.

2.3.2 Single- and multiplayer titles

The temporal reallocation was evident in both single- and multiplayer games. For each platform version of a game, GPT-5.5 indicated whether local multiplayer (several players on a single device) and online multiplayer (several networked devices) were available at release (yes, no, or

not determinable). A title's local- and online-multiplayer shares are the proportions of its platform versions offering each feature. We classified a title as multiplayer when either its local- or online-multiplayer share was at least .50 and re-estimated the two index trends separately within each mode. Among single-player titles, the prosocial index increased by 4.5 points per decade (0.37 per ingredient) and the competitive index decreased by 4.8 points (0.54 per ingredient); among multiplayer titles, the corresponding changes were +5.2 and -6.1 points per decade (+0.44 and -0.68 per ingredient), all $p < .001$. The year-by-mode interaction indicated that the reallocation was, if anything, slightly steeper among multiplayer titles (prosocial, +1.5 points per decade, $p < .001$; competitive, -1.2 points, $p < .001$). Because the same pattern appeared in games involving human co-players, it cannot be explained solely by the increasing use of algorithmically simulated companionship in single-player games (Figure 3).

We further divided multiplayer titles according to whether their multiplayer content was predominantly local ($n = 4,028$) or online ($n = 2,170$), and compared them with single-player titles ($n = 5,161$). The same reallocation appeared in all three categories. In single-player, local-multiplayer, and online-multiplayer titles, respectively, the prosocial index increased by 4.5, 4.3, and 5.5 percentage points per decade, whereas the competitive index decreased by 4.8, 5.5 and 3.9 (all $p \leq .001$).

Relative to single-player titles, the prosocial increase did not differ among online-multiplayer titles (+0.65 points per decade, $p = .26$), and the competitive decline was weaker (+2.63 points per decade, $p < .001$). The reallocation was therefore not concentrated in the titles most exposed to online play, which argues against reducing it to the spread of online connectivity.




Figure 3. Temporal trends in prosocial and competitive indices for single-player, local-multiplayer, and online-multiplayer titles (1993-2016). Points represent the annual mean prosocial (purple) and competitive (orange) indices, as mean intensity per constituent ingredient, within each play mode. Lines show OLS fits with 95% confidence bands. Asterisks indicate the significance of the temporal trend (** $p < .01$, *** $p < .001$).

2.3.3 Within franchises

The temporal reallocation documented thus far compares games released in different years, leaving open a compositional explanation: successive waves of releases may draw on different franchises, such that the aggregate trend reflects which franchises entered or disappeared from the market rather than changes within the franchises themselves. To address this possibility, we examined whether successive installments of the same franchise shifted towards the prosocial pole.

To identify franchises, we provided the full list of titles in the annotated corpus to a large language model (Claude Opus 4.8) and instructed it to group releases belonging to the same franchise. This procedure identified 832 franchises covering 6,063 releases. The automatic grouping procedure sometimes merged unrelated titles that shared short, generic name prefixes. We therefore first audited the groups most at risk, flagging those whose assigned franchise name either consisted of a single common English word or formed a prefix of another franchise name. These criteria identified 121 candidate groups, which we inspected manually. Fifty-nine proved to combine unrelated titles on the basis of a shared word. For example, the group labeled “Space” included *Space Jam*, *Space Hulk*, and *Space Chimps*. We removed these groups, leaving 773 franchises comprising 5,673 releases. To estimate the residual error rate after this audit, we manually inspected a random sample of 30 retained franchises. Twenty-eight formed coherent series and two did not, corresponding to an estimated residual error rate of 6.7%. Because erroneous groups combine unrelated titles, they should add noise and attenuate the within-franchise estimates towards zero rather than generate a systematic temporal trend, provided that grouping errors are not themselves associated with release year.

For each index, we regressed its value on release year with franchise fixed effects and standard errors clustered by franchise. The year coefficients were therefore identified exclusively from temporal variation within franchises. Later releases within a franchise were, on average, more prosocial and less competitive than earlier releases: the prosocial index increased by 2.5 percentage points per decade (0.21 per ingredient), whereas the competitive index decreased by 2.2 points (0.25 per ingredient), both $p < .001$. These within-franchise slopes were about half the magnitude of the corpus-wide slopes, consistent with stable differences between franchises contributing to the aggregate pattern. Nevertheless, both retained the predicted direction and remained highly significant. Estimated franchise by franchise among the 576 franchises spanning at least four distinct release years, the prosocial index rose in 65.8% and the competitive index fell in 62.7%, with both moving in the predicted direction in 48.6%. *FIFA* illustrates how this reallocation can occur without removing competition from the game. Across its 37 titles, the prosocial share (Q) rose from 11.4% to 21.8%, while the raw scores of key competitive motivations remained stable: Pride was rated 5 in every edition, and Anger and Coalitional affiliation were rated as high in the latest editions as in the 1990s (4 and 5.9 on the 0-6 scale). The model describes team rivalry as “one of the game’s clearest motivational structures” in both *FIFA 17* and *FIFA: Road to World Cup 98*. The motivational repertoire surrounding match play expanded, with the summed score across all ingredients rising from 56 to 74. Resource accumulation showed the largest increase, from 0.1 to 3.5: the model describes early titles as “not built around stockpiling money, materials, or consumables”, whereas later editions introduce career-mode budgets and transfers and, eventually, Ultimate Team, where players earn coins,

collect cards and build squads. Competition thus remained central to match play but came to occupy a smaller share of a broader motivational repertoire. Holding franchise identity constant, later installments therefore devoted a larger proportion of their motivational content to prosocial ingredients and a smaller share to competitive ingredients (Figure 4).




Figure 4. Within-franchise reallocation of motivational content (1993-2016). (a) The distribution of franchise-specific temporal slopes, in points per decade, on the per-ingredient scale, for the prosocial and competitive indices among 576 franchises represented in at least four distinct release years. Dashed vertical lines indicate the corresponding fixed-effects estimate; asterisks indicate its significance (** $p < .001$). (b) The temporal trajectories of three illustrative franchises (Grand Theft Auto, Naruto, FIFA). Points represent annual means within franchises; lines show the OLS fits, with 95% confidence bands.

2.3.4 Best-selling titles and sales weighting

The temporal reallocation was also present among the games with the greatest commercial reach, although its magnitude depended on the weight assigned to the largest sellers. Restricting each release year to its best-selling titles left the prosocial increase significant at every threshold tested, from the top 10 to the top 100 titles per year. Among the annual top 100, the prosocial index increased by 4.4 percentage points per decade (0.37 per ingredient, $p < .001$). The competitive decrease was significant at every threshold from the top 30 onwards but not in the more restrictive top 10 sample, where the slope was indistinguishable from zero (-0.004 points per decade, $p = .99$) and each year contributed only 10 observations.

We also weighted each game in proportion to its global sales, thereby giving greater influence to titles that sold more copies. Sales weighting attenuated both temporal slopes but preserved their direction and statistical significance: the prosocial index increased by 4.0 percentage points per decade (0.33 per ingredient) and the competitive index decreased by 2.5 points (0.27 per

ingredient), compared with unweighted estimates of +6.0 and -7.3 points (+0.50 and -0.81 per ingredient), respectively. This attenuation suggests that lower-selling titles contributed disproportionately to the magnitude of the unweighted reallocation. Nevertheless, the shift remained significant both among the best-selling titles and when all releases were weighted by their commercial reach (Figure 5, panel B).

2.3.5 Across the distribution of games

The reallocation was not confined to a fringe of unusually prosocial games but was evident across the distribution of releases. Quantile regressions of each index on release year yielded positive prosocial slopes at the 10th (+3.6 percentage points per decade, 0.30 per ingredient), 50th (+7.9, 0.66), and 90th (+3.8, 0.32) percentiles (all $p < .001$), with the largest estimated increase occurring at the median. The competitive index declined at all three percentiles as well (-6.0, -8.3, and -6.5 points per decade; -0.67, -0.93, and -0.72 per ingredient; all $p < .001$), again most steeply at the median. If the aggregate had been driven solely by an increasingly prosocial upper tail, little or no change would have been expected at the median or lower quantiles. Instead, all three quantiles increased substantially. The reallocation therefore characterized the typical game released in each year, rather than only a small group of outlying titles.

2.3.6 Independence from any single ingredient

Neither temporal trend depended on any single constituent ingredient. We recomputed each index after excluding its ingredients one at a time, yielding 12 leave-one-out specifications for the prosocial index and 9 for the competitive index. In every specification, the temporal slope retained its original direction and remained significant. Excluding Friendship produced the largest attenuation of the prosocial trend, yet the prosocial index still increased by 4.4 percentage points per decade (0.40 per ingredient); excluding Pride produced the largest attenuation of the competitive trend, which still declined by 5.0 points (0.62 per ingredient) (Figure 5, panel A). A bootstrap over games with 1,000 resamples likewise yielded narrow intervals around both year-level slopes: +5.4 per decade for the prosocial index (95% bootstrap interval [5.0, 5.9]) and -6.3 for the competitive index (95% bootstrap interval [-6.8, -5.9]). The two trends therefore reflect changes distributed across the ingredients composing each pole rather than changes in any single motivation.

2.3.7 Robustness to alternative index definitions

The temporal reallocation was robust to the particular composition of the two indices. We conducted a multiverse analysis (Steegen et al., 2016), recomputing the prosocial share $Q = P/(P+C)$ for every admissible pair of ingredient subsets containing at least 6 of the original 12 prosocial ingredients and at least 5 of the original 9 competitive ingredients, each subset divided by its own size so that specifications of different sizes remain comparable. The procedure yielded 642,560 alternative index definitions. Regressing Q on release year across the 11,359 releases, Q increased significantly in every specification (range +6.4 to +17.0 percentage points per decade; median +11.6, against +11.6 for the reference definition) (Figure 5, panel C). Thus, within this admissible specification space, no choice of which ingredients count as prosocial or competitive removed or reversed the effect.




Figure 5. Robustness checks for the two functional indices. (a) The leave-one-ingredient-out analysis. Each point represents the temporal trend, in points of total motivational content per decade, obtained after removing one ingredient from its respective index; the dashed line marks the estimate from the full definition. All trends retain their expected direction and remain significant. (b) The estimated trends after restricting each release year to its top 10 through top 100 best-selling titles; asterisks indicate the significance of each trend (***) $p < .001$; NS, non-significant. (c) The distribution of temporal trends in the prosocial share (Q) across all 642,560 admissible alternative index definitions. The dashed vertical line marks the null value ($\beta = 0$), and the black line indicates the estimate obtained using the original index definitions.

2.3.8 Stability of the ingredient co-occurrence structure

Finally, we examined whether the temporal reallocation was accompanied by a broader change in how motivational ingredients combine within games. We constructed separate 48-by-48 ingredient co-occurrence matrices for the early (1993-2004, $n = 3,925$) and later (2005-2016, $n = 7,434$) halves of the period. The two matrices showed a very strong correspondence (Mantel $r = .93$, $p < .001$, 9,999 permutations), indicating that the overall pattern of which motivations tended to occur together remained highly stable. The reallocation therefore primarily concerned the relative prevalence of different motivations within a stable co-occurrence structure, rather than a reorganization of that structure. This result does not imply that every pairwise relationship remained unchanged, but it provides no evidence of a large-scale transformation in the motivational architecture of games.

Taken together, Study 1 shows that, between 1993 and 2016, the motivational content of games shifted away from competition and dominance and towards affiliation, warmth, and care. The temporal slope of the prosocial index was approximately 1.3 points per decade more positive, on the per-ingredient scale, than that of the competitive index, substantially narrowing the initial gap in favor of the competitive pole. This shift reflected a compositional reallocation of motivational content rather than a general increase in motivational richness over time. The reallocation was evident at the level of individual game releases, survived adjustment for genre, number of platforms, and global sales, and extended across the distribution of games. It also held within genres, play modes, and franchises, among best-selling and sales-weighted titles, after removing each ingredient in turn, and across every admissible alternative index definition. The prosocial increase held in every one of these specifications; the competitive decline, while robust overall, was attenuated among the highest-selling titles. Finally, the trend was independently recovered from the human annotations, while the broader co-occurrence structure of motivational ingredients remained highly stable.

3. Study 2: Personality predicts players' preferred motivational content

Study 1 established that the balance of motivational content in games shifted over two decades, away from the agonistic pole of status, anger, hatred, revenge and towards the caring pole of affiliation, friendship, warmth and compassion. One implication of a medium that tracks its audience's motivations is that attraction to these forms of content should vary systematically across individuals. Study 2 therefore examines whether players' Big Five personality traits predict the relative balance of prosocial and competitive motivations in their favorite games.

The design also distinguishes two forms of sociality that are often conflated. One concerns the *social context* of play: a preference for playing with – and often competing against – other people. The other concerns the *social content* of play: a preference for games that engage motivations related to care, warmth, and affiliation. These are distinct preferences, which we expected to load on personality traits: preference for a social play context with Extraversion, reflecting sociability and assertiveness, and preference for prosocial content with Agreeableness, reflecting warmth and concern for others.

3.1. Method

Player data were collected through a gamified online platform launched in 2024. After completing the assessment (a personality inventory and a short questionnaire on gaming habits, approximately 15 minutes), users received a personalized player profile and game recommendations. The study was approved by the Research Ethics Committee of Université Paris Cité (IRB No. 00012024-50) on June 11, 2024, and was conducted in accordance with the General Data Protection Regulation (GDPR). All participants provided informed consent. For each participant, we recorded their preferred game genres, up to five favorite games entered as free text, total number of games played, weekly playing time, age, and gender. Personality was assessed using the Big Five

Inventory-2 (BFI-2-S), which yields scores for five broad personality domains and 15 facets, with three facets nested within each domain (Soto & John, 2017).

The initial sample comprised 4,657 players. Gender was self-reported from four options: 3,785 reported male, 713 female, 67 non-binary or third gender, and 92 preferred not to report it. Because the latter two groups were too small to yield stable estimates, we restricted analyses to the 4,498 players who reported male or female (mean age = 32.6 years; 15.9% women). Gender was modeled as a binary variable. Our estimates therefore describe the contrast between men and women, and their applicability to the 159 excluded players remains untested. Every listed game was annotated on the same 48 motivational ingredients as in Study 1, using the same large-language-model procedure (GPT-5.4, identical prompt). Annotation was run separately, on the game names as written by participants. Because some names were ambiguous (e.g., “fifa”), titles present in both corpora did not receive identical scores, but these correlated at $r = .86$. After excluding game entries for which motivational content could not be coded, the dataset contained 21,539 player-game observations.

For each game, we calculated the prosocial (P) and competitive (C) index scores using the same fixed lists of 12 and 9 ingredients, respectively, as in Study 1, each divided by its number of constituent ingredients. We then expressed the game’s position within this two-pole composition as the prosocial share,

$$Q = P / (P+C),$$

with the complementary competitive share given by

$$1-Q = C / (P+C)$$

Thus, Q represents the proportion of the game’s caring-plus-agonistic content allocated to the caring pole ($M = .39$, $SD = .12$).

Because the two shares are exact complements, linear models predicting them are mirror images: each predictor has coefficients of equal magnitude but opposite sign. We therefore report only the model predicting the prosocial share. Positive coefficients indicate a preference for more caring relative to agonistic content, whereas negative coefficients indicate the reverse.

3.2. Analyses

Players completed the BFI-2-S, from which we calculated scores for the five broad personality traits (Extraversion, Agreeableness, Conscientiousness, Neuroticism, Openness) and their 15 constituent facets. All trait and facet scores were standardized before analysis; the outcome Q was left on its original scale.

We fitted a linear mixed-effects model predicting the prosocial share Q of each listed game from the five traits, so that coefficients express the change in Q associated with a one-standard-deviation difference in a trait. The model included a random intercept for players to account for the clustering of multiple games within players, together with age, gender, weekly

playing time, and years of gaming experience as covariates. Because all five domains were included in the same model, each coefficient represents the association with one trait adjusted for the other four. Fixed effects were tested using Wald z tests. Because Q is a property of the game rather than of the player, every nomination of the same title carries the same value. We nonetheless checked that the trait-level results held under standard errors clustered by both player and title.

We then conducted three complementary analyses using the same covariate and random-effects structure. First, to determine whether Agreeableness predicted specifically the relative composition of motivational content rather than overall motivational richness, we fitted a parallel model with the summed raw score across all 48 ingredients as the outcome. Second, we fitted a facet-level model in which the five domain scores were replaced by the 15 BFI-2-S facet scores. Third, we constructed an ingredient-level personality fingerprint by fitting 48 separate models, each with one of the raw ingredient scores as the outcome and the five domain scores entered simultaneously as predictors. The resulting 240 p-values (48 outcomes * 5 traits) were corrected using the Benjamini-Hochberg procedure.

3.3. Results

Personality traits were systematically associated with the relative balance between the two motivational poles. All coefficients reported below come from the model predicting the prosocial share (Q). Positive coefficients therefore indicate a shift towards the caring pole, whereas negative coefficients indicate a shift towards the agonistic pole.

Extraversion predicted a lower share ($b = -0.006$, $z = -5.36$, $p < .001$), whereas Agreeableness predicted a higher share ($b = 0.003$, $z = 3.47$, $p < .001$). Conscientiousness ($b = 0.005$, $z = 4.48$, $p < .001$) was also associated with a greater relative preference for the caring pole, as were Neuroticism ($b = 0.003$, $z = 2.56$, $p = .011$), and, more weakly, Openness ($b = 0.002$, $z = 2.02$, $p = .043$). These associations were modest but statistically reliable: a one-standard-deviation difference in a trait corresponded to a shift of between 0.2 and 0.6 percentage points in Q. Gender showed a substantially larger association. Holding the other predictors constant, women's listed games had a prosocial share (Q) approximately 4.7 percentage points higher than men's ($b = 0.047$, $z = 17.0$, $p < .001$), an effect more than eight times larger than the largest personality-trait coefficient (Figure 6). Under standard errors clustered by both player and title, the associations with Extraversion, Agreeableness, Conscientiousness and Neuroticism, and that of gender, remained significant, whereas the association with Openness did not ($p = .065$). The associations between personality traits and the prosocial share were largely consistent across genders. Estimated separately, Extraversion, Agreeableness and Conscientiousness were significant predictors among men ($N = 3,785$ players), and Extraversion, Agreeableness, Conscientiousness ($p = .049$) and Neuroticism among women ($N = 713$); Openness reached significance in neither group. None of the five trait-by-gender interactions was significant (all $p \geq .08$), providing no evidence that these associations differed by gender.




Figure 6. Personality predictors of motivational content preferences (Study 2). (a) Fixed-effect coefficients from the linear mixed-effects model predicting the prosocial share (Q); positive coefficients indicate a relative preference for caring content, whereas negative coefficients indicate a relative preference for agonistic content. (b) The coefficients for the complementary competitive share (1-Q); these are exact sign reversals of those in Panel A. Points represent coefficient estimates (b), and bars represent 95% confidence intervals. Big Five traits were z-scored, so coefficients express the change in Q associated with a one-standard-deviation difference in a trait; gender was coded 0 = male and 1 = female. Age, weekly playing time and years of gaming were entered on their original scales. The model included random intercepts for players (not shown).

The association with Agreeableness concerned the relative composition of motivational content rather than its overall richness. Agreeableness predicted a higher prosocial share ($b = 0.003$, $p < .001$) but was unrelated to the summed intensity of the 48 motivational ingredients ($b = -0.022$, $p = .92$). Thus, the favorite games of more agreeable players were not motivationally richer overall; instead, a greater proportion of their motivational content was allocated to the caring rather than the agonistic pole. This pattern argues against the alternative explanation that agreeable players simply prefer games with denser or more motivational content.

The facet-level results sharpened, and partly complicated, the trait-level pattern. Positive associations with the prosocial share (Q) emerged for Aesthetic sensitivity (a facet of Openness, $b = 0.005$) and Compassion (a facet of Agreeableness, $b = 0.004$). By contrast, the strongest negative association was observed for Sociability, the facet of Extraversion most directly related to social engagement ($b = -0.003$), followed by Trust, a facet of Agreeableness ($b = -0.003$, $p = .011$). These four associations survive Benjamini-Hochberg correction across the 15 facets; Anxiety ($b = 0.003$) and Respectfulness ($b = 0.002$) do not.

That Sociability predicted a preference for relatively more agonistic rather than caring content is consistent with the distinction between social context and social content. A disposition towards seeking social interaction does not necessarily imply a preference for games whose content centers on care, warmth and affiliation. These two forms of social preference appear to relate to different aspects of personality. The 48-ingredient profile of Agreeableness provided a more

nuanced account than the simple conclusion that more agreeable players prefer caring content. Its largest raw-score associations were negative and concerned agonistic ingredients: higher Agreeableness was associated with lower scores for Status-seeking ($b = -0.061$), Coalitional affiliation (-0.050), Envy (-0.044), Revenge (-0.033), Fear of aggressors (-0.031), Shame (-0.030) and Anger (-0.027). Among the caring ingredients, Warmth was the only one showing a significant positive association ($b = 0.026$). Affiliation was instead negatively associated with Agreeableness ($b = -0.032$), whereas Compassion, Friendship, and Familial love showed no reliable associations. The same contrast appeared in the ingredient-level profile of Extraversion, which mirrored that of Agreeableness: positive associations for Status-seeking ($b = 0.078$), Coalitional affiliation (0.052), Shame (0.047), Fear of aggressors (0.045), Envy (0.042) and Anger (0.039), and negative ones for Warmth, Compassion, Friendship and Familial love. All ingredient-level models entered the five domains simultaneously and were Benjamini-Hochberg corrected across the 240 tests.

The positive association between Agreeableness and the prosocial share therefore did not reflect a broad preference for caring motivations. It arose substantially from lower scores on agonistic ingredients: because $Q=P/(P+C)$, a reduction in competitive content necessarily increases the relative prosocial share even if prosocial content does not increase. Together with the specific positive association with Warmth, this pattern suggests that the Agreeableness effect reflects a relative avoidance of agonistic content more than a general attraction to care.

4. General Discussion

Across two studies, we document a historical reallocation in the motivational content of video games and show that preferences along the same dimension covary with individual differences in personality. In Study 1, the prosocial index, the share of games' motivational content devoted to prosocial motivations – including affiliation, friendship, warmth, and compassion – increased between 1993 and 2016, whereas the competitive index devoted to agonistic motivations – including status, anger and revenge – decreased. The temporal slope of the prosocial index was approximately 1.3 points per decade more positive, on the per-ingredient scale, than that of the competitive index, substantially narrowing the initial gap between the two poles. Because the analysis captured relative composition rather than absolute motivational intensity, this pattern reflects a reallocation of motivational content rather than merely the general increase in motivational richness over time. We refer to this historical shift as a “prosocial turn.”

Study 2 complements this historical analysis by asking whether the same motivational axis also explains variation in players' preferences. Players higher in Agreeableness listed favorite games with a larger prosocial share – an association driven primarily by lower agonistic content – whereas players higher in Extraversion showed the reverse pattern. Together, the two studies connect a historical trend with a corresponding dimension of individual variation: the shift occurred in the direction of the preferences associated with higher Agreeableness and lower Extraversion. This correspondence does not, however, establish that historical changes in the composition of the gaming audience caused the shift.

We interpret this convergence between the historical shift and contemporary player preferences through the entertainment technology hypothesis (Dubourg & Baumard, 2022). This framework

represents entertainment products as bundles of motivational ingredients that engage evolved preferences, with each ingredient defined by the fundamental motive it activates (Dubourg, Thouzeau, et al., 2026). A game's appeal should therefore depend partly on the match between its motivational composition and the preferences of its audience, while popular culture may more broadly reflect the preferences to which it caters (Saad, 2011). From a cultural-evolutionary perspective, the distribution of content within a medium should change when the environment created by audience preferences changes. New cultural variants are continually generated, and those that fit available cognitive niches are more likely to be adopted and retained. Through repeated cycles of production, selection, and transmission, a cultural repertoire may therefore shift towards cultural attractors grounded in consumers' minds (Claidière & Sperber, 2007; Claidière et al., 2014; Acerbi & Mesoudi, 2015; Mesoudi, 2016). The prosocial turn can be interpreted as the cultural signature of a changing alignment between the motivational content offered by games and the preferences of their audience.

At least three non-exclusive mechanisms could have produced this realignment, and our data cannot distinguish between them.




Figure 7. Three non-exclusive accounts of the prosocial turn. Each account could produce the observed increase in the prosocial share (Q) of game content (bottom), and the present data cannot distinguish between them. Circles represent the player population before and after the shift; grey boxes summarise the mechanism at work in each panel. (a) Audience-composition: new players with different preferences enter the market, altering demand for motivational content, while existing players' preferences remain unchanged. The larger circle reflects the expansion of the player population. (b) Psychological-ecological: existing players become more prosocial as part of a broader shift in the population (e.g., because of rising material security). The values expressed in cultural products shift accordingly. (c) Design-drift: the player population is assumed to be predominantly prosocial and remains unchanged in composition and preferences. New technological and design affordances, such as cooperative mechanics and bonds with characters, expand developers' capacity to represent and sustain caring content. This allows game content to catch up with players' existing preferences.

The first possibility is an “audience-composition” account (Figure 7, panel A): as gaming expanded beyond its early audience, the distribution of player preferences may also have shifted. If the groups entering the market were, on average, less attracted to status and conflict and more attracted to care and affiliation, the changing composition of the player population could have altered demand for these different forms of motivational content. This possibility is consistent with the substantial gender difference observed in Study 2 and with previous evidence of gendered game preferences (Lucas & Sherry, 2004; Lange et al., 2021). Nevertheless, the personality associations were similar in direction among men and women, with no significant trait-by-gender interactions. This suggests that dispositional differences in preferences along the caring-agonistic axis were present in both groups. Long-running ESA and ISFE reports also document an increasing representation of women and older players since the late 1990s. These demographic series are external to our corpus, however, and cannot establish that changes in audience composition caused the historical shift in game content.

A second possibility operates at a broader level than the first. Under a “psychological-ecological” account (Figure 7, panel B), the shift need not reflect changes specific to the composition of the gaming audience. Instead, the wider population may have shifted, on average, towards greater cooperation, tolerance, and concern for others. This account treats such dispositions as responsive to local ecological conditions. When access to basic resources such as food, health care, and education is secure, individuals are hypothesized to adopt slower, more future-oriented life strategies, which in turn favor patience, trust, and concern for others (Dubourg et al., 2026; Baumard et al., 2015; Baumard & Chevallier, 2015; Sheskin et al., 2014). Because the decades our corpus spans were, in most of the markets it covers, decades of rising affluence, this account predicts a broader shift towards prosocial and caring dispositions, accompanied by a corresponding change in the values expressed in cultural products (Baumard & André, 2025). Related mechanisms have been invoked to explain the historical rise of prosociality in Western fiction (Martins & Baumard, 2020) and the postmaterialist shift in values documented by decades of survey research (Inglehart et al., 2017). On this interpretation, the prosocial turn in video games may be another cultural trace of the same long-term realignment, with games serving as a relatively recent medium in which it can be observed. Of the three accounts, this one connects our findings most directly to the broader literature on how material security may reshape human motivation (Boon-Falleur et al., 2025).

A third possibility is a “design-drift” account (Figure 7, panel C), under which neither the composition nor the preferences of the player population need have changed. Instead, new technological and design affordances (including persistent worlds, cooperative mechanics, parasocial relationships with characters, and online connectivity) may have expanded developers’ capacity to represent and sustain caring forms of motivational content. Study 1 speaks against the strongest version of this account, in which the prosocial turn results primarily from the emergence of online or multiplayer gaming. The reallocation was also present in single-player titles and within long-running franchises. Although it was slightly steeper among multiplayer titles considered as a single group, the prosocial increase did not differ reliably across single-player, local multiplayer and online multiplayer titles when the three categories were compared. It therefore cannot be reduced to the arrival of multiplayer or internet connectivity. A broader design-drift account nevertheless remains plausible, because many relevant affordances

also transformed single-player games. Distinguishing changes in who plays from changes in what designers can build would require longitudinal data linking audience composition and preferences to evolving design capabilities and game content. We regard this as a central question for future research.

The present work has several limitations. First, the corpus is restricted to console and PC titles tracked through unit sales and excludes mobile and free-to-play games, which constitute a substantial segment of contemporary gaming. Our estimates should therefore be interpreted as characterizing the traditional console and PC market rather than video games as a whole. Although the accessibility and social orientation of many mobile and free-to-play games suggest that their inclusion might strengthen the observed trend, this possibility remains to be tested. Second, Study 2 is cross-sectional. It establishes that personality traits covary with preferences for motivational content among players, but it cannot show that the trait composition of the gaming audience changed over the period examined in Study 1. Longitudinal evidence of such a change would be required to test the audience-composition account directly. Finally, the personality associations were modest, consistent with those typically observed between broad traits and specific behaviors (Vize et al., 2023; Mammadov, 2022; Entringer et al., 2021), while gender showed a substantially larger association with motivational-content preference. We therefore interpret the personality findings as evidence that stable dispositions are associated with players' content preferences, not as evidence that changes in those dispositions drove the historical shift.

5. Acknowledgements

We thank Paco Maunoir-Guérin and Valentin Thouzeau for their help during the conceptualization of the project.

6. Funding

This work was supported by the EUR FrontCog grant [ANR-17-EURE-0017]; PSL [ANR-10-IDEX-0001-02]; the ANR MAGIC grant [ANR-25-CE28-5813-01]; and the Major Research Program of PSL Research University "CultureLab" [ANR-10-IDEX-0001]. The funding sources had no involvement in study design; in the collection, analysis and interpretation of data; in the writing of the report; or in the decision to submit the article for publication.

7. CRediT

Alexandre Javelle: Data curation, Formal analysis, Software, Visualization, Validation, Writing – review and editing. Benjamin Dahan: Data curation, Formal analysis, Software, Writing – review and editing. Toni Zian Alabdeen Ali: Conceptualization, Data curation, Software, Writing – review and editing. Valérian Chambon: Conceptualization, Funding acquisition, Project administration, Supervision, Validation, Writing – review and editing. Edgar Dubourg: Conceptualization, Funding acquisition, Project administration, Supervision, Validation, Writing – review and editing, Writing – original draft.

8. AI Statement

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process. During the preparation of this work the authors used Claude Opus 5.5 in order to improve the readability and language of the manuscript and to assist with analysis code. They also used GPT-5.4 for LLM annotations of the video games in Studies 1 and 2, GPT-5.5 to annotate multiplayer features, and Claude Opus 4.8 to group titles into franchises (see Method). After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

9. Data availability

The annotated corpus, the human validation data and all analysis scripts are available on OSF [https://osf.io/sj6ch/overview?view_only=615f1e71542744b8bc0cd6184388a0c3]. The full analysis pipeline was reproduced independently in R and in Python.

10. Competing interests

The authors declare no competing interests.

11. References

- Acerbi, A., & Mesoudi, A. (2015). If we are all cultural Darwinians what's the fuss about? Clarifying recent disagreements in the field of cultural evolution. *Biology & Philosophy*, 30(4), 481–503. <https://doi.org/10.1007/s10539-015-9490-2>
- Anderson, C. A., Shibuya, A., Ihori, N., Swing, E. L., Bushman, B. J., Sakamoto, A., Rothstein, H. R., & Saleem, M. (2010). Violent video game effects on aggression, empathy, and prosocial behavior in Eastern and Western countries: A meta-analytic review. *Psychological Bulletin*, 136(2), 151–173. <https://doi.org/10.1037/a0018251>
- Angelopoulos, A. N., Bates, S., Fannjiang, C., Jordan, M. I., & Zrnic, T. (2023). Prediction-powered inference. *Science*, 382(6671), 669–674. <https://doi.org/10.1126/science.adi6000>
- Baumard, N., & André, J.-B. (2025). The ecological approach to culture. *Evolution and Human Behavior*, 46(3), Article 106686. <https://doi.org/10.1016/j.evolhumbehav.2025.106686>
- Baumard, N., & Chevallier, C. (2015). The nature and dynamics of world religions: A life-history approach. *Proceedings of the Royal Society B: Biological Sciences*, 282(1818), Article 20151593. <https://doi.org/10.1098/rspb.2015.1593>
- Baumard, N., Hyafil, A., Morris, I., & Boyer, P. (2015). Increased affluence explains the emergence of ascetic wisdoms and moralizing religions. *Current Biology*, 25(1), 10–15. <https://doi.org/10.1016/j.cub.2014.10.063>
- Baumard, N., Safra, L., Martins, M., & Chevallier, C. (2024). Cognitive fossils: Using cultural artifacts to reconstruct psychological changes throughout history. *Trends in Cognitive Sciences*, 28(2), 172–186. <https://doi.org/10.1016/j.tics.2023.10.001>

- Bongini, P., Becattini, F., & Del Bimbo, A. (2023). Is GPT-3 all you need for visual question answering in cultural heritage? In L. Karlinsky, T. Michaeli, & K. Nishino (Eds.), *Computer vision – ECCV 2022 workshops: Proceedings, Part I* (Lecture Notes in Computer Science, Vol. 13801, pp. 268–281). Springer. https://doi.org/10.1007/978-3-031-25056-9_18
- Boon-Falleur, M., Baumard, N., & André, J.-B. (2025). The effect of income and wealth on behavioral strategies, personality traits, and preferences. *Perspectives on Psychological Science*, 20(4), 669–690. <https://doi.org/10.1177/17456916231201512>
- Claidière, N., Scott-Phillips, T. C., & Sperber, D. (2014). How Darwinian is cultural evolution? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1642), Article 20130368. <https://doi.org/10.1098/rstb.2013.0368>
- Claidière, N., & Sperber, D. (2007). The role of attraction in cultural evolution. *Journal of Cognition and Culture*, 7(1–2), 89–111. <https://doi.org/10.1163/156853707X171829>
- Cole, H., & Griffiths, M. D. (2007). Social interactions in massively multiplayer online role-playing gamers. *CyberPsychology & Behavior*, 10(4), 575–583. <https://doi.org/10.1089/cpb.2007.9988>
- Domahidi, E., Festl, R., & Quandt, T. (2014). To dwell among gamers: Investigating the relationship between social online game use and gaming-related friendships. *Computers in Human Behavior*, 35, 107–115. <https://doi.org/10.1016/j.chb.2014.02.023>
- Dubourg, E., & Baumard, N. (2022). Why and how did narrative fictions evolve? Fictions as entertainment technologies. *Frontiers in Psychology*, 13, Article 786770. <https://doi.org/10.3389/fpsyg.2022.786770>
- Dubourg, E., & Chambon, V. (2025). DEEP: A model of gaming preferences informed by the hierarchical nature of goal-oriented cognition. *Entertainment Computing*, 53, Article 100930. <https://doi.org/10.1016/j.entcom.2025.100930>
- Dubourg, E., Jach, H. K., Beuchot, T., Smillie, L. D., & Baumard, N. (2026). Openness to Experience: From ecology to culture. *Trends in Cognitive Sciences*. Advance online publication. <https://doi.org/10.1016/j.tics.2026.03.006>
- Dubourg, E., & Thouzeau, V. (2026). *The double edge of AI annotation in cultural sciences* [Manuscript submitted for publication]. Institut Jean Nicod, Département d'études cognitives, École normale supérieure, Université PSL.
- Dubourg, E., Thouzeau, V., & Baumard, N. (2024). A step-by-step method for cultural annotation by LLMs. *Frontiers in Artificial Intelligence*, 7, Article 1365508. <https://doi.org/10.3389/frai.2024.1365508>
- Dubourg, E., Thouzeau, V., Beuchot, T., Bonard, C., Boyer, P., Chevallier, C., Clasen, M., Boon-Falleur, M., Fiorio, G., Fitouchi, L., Fisher, M. L., Gantman, A., Grant, A., Hye-Knudsen, M., Katiyar, T., Kjeldgaard-Christiansen, J., Mercier, M., Mercier, H., Morin, O., . . . Baumard, N. (2026). Carving stories at their natural joints. *Psychological Review*. Advance online publication. <https://doi.org/10.1037/rev0000648>
- Entringer, T. M., Gebauer, J. E., Eck, J., Bleidorn, W., Rentfrow, P. J., Potter, J., & Gosling, S. D. (2021). Big Five facets and religiosity: Three large-scale, cross-cultural, theory-driven, and

- process-attentive tests. *Journal of Personality and Social Psychology*, 120(6), 1662–1695.
<https://doi.org/10.1037/pspp0000364>
- Ferguson, C. J. (2015). Do Angry Birds make for angry children? A meta-analysis of video game influences on children's and adolescents' aggression, mental health, prosocial behavior, and academic performance. *Perspectives on Psychological Science*, 10(5), 646–666.
<https://doi.org/10.1177/1745691615592234>
- Ferguson, C. J., & Olson, C. K. (2013). Friends, fun, frustration and fantasy: Child motivations for video game play. *Motivation and Emotion*, 37(1), 154–164.
<https://doi.org/10.1007/s11031-012-9284-7>
- Gentile, D. A., Anderson, C. A., Yukawa, S., Ihori, N., Saleem, M., Lim, K. M., Shibuya, A., Liau, A. K., Khoo, A., Bushman, B. J., Huesmann, L. R., & Sakamoto, A. (2009). The effects of prosocial video games on prosocial behaviors: International evidence from correlational, longitudinal, and experimental studies. *Personality and Social Psychology Bulletin*, 35(6), 752–763. <https://doi.org/10.1177/0146167209333045>
- Gilardi, F., Alizadeh, M., & Kubli, M. (2023). ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30), Article e2305016120. <https://doi.org/10.1073/pnas.2305016120>
- Granic, I., Lobel, A., & Engels, R. C. M. E. (2014). The benefits of playing video games. *American Psychologist*, 69(1), 66–78. <https://doi.org/10.1037/a0034857>
- Greitemeyer, T., & Osswald, S. (2010). Effects of prosocial video games on prosocial behavior. *Journal of Personality and Social Psychology*, 98(2), 211–221.
<https://doi.org/10.1037/a0016997>
- Inglehart, R. F., Ponarin, E., & Inglehart, R. C. (2017). Cultural change, slow and fast: The distinctive trajectory of norms governing gender equality and sexual orientation. *Social Forces*, 95(4), 1313–1340. <https://doi.org/10.1093/sf/sox008>
- Kenrick, D. T., Griskevicius, V., Neuberg, S. L., & Schaller, M. (2010). Renovating the pyramid of needs: Contemporary extensions built upon ancient foundations. *Perspectives on Psychological Science*, 5(3), 292–314. <https://doi.org/10.1177/1745691610369469>
- Kowert, R., Domahidi, E., Festl, R., & Quandt, T. (2014). Social gaming, lonely life? The impact of digital game play on adolescents' social circles. *Computers in Human Behavior*, 36, 385–390. <https://doi.org/10.1016/j.chb.2014.04.003>
- Lange, B. P., Wühr, P., & Schwarz, S. (2021). Of Time Gals and Mega Men: Empirical findings on gender differences in digital game genre preferences and the accuracy of respective gender stereotypes. *Frontiers in Psychology*, 12, Article 657430.
<https://doi.org/10.3389/fpsyg.2021.657430>
- Lucas, K., & Sherry, J. L. (2004). Sex differences in video game play: A communication-based explanation. *Communication Research*, 31(5), 499–523.
<https://doi.org/10.1177/0093650204267930>
- Mammadov, S. (2022). Big Five personality traits and academic performance: A meta-analysis. *Journal of Personality*, 90(2), 222–255. <https://doi.org/10.1111/jopy.12663>

- Martins, M. de J. D., & Baumard, N. (2020). The rise of prosociality in fiction preceded democratic revolutions in Early Modern Europe. *Proceedings of the National Academy of Sciences*, 117(46), 28684–28691. <https://doi.org/10.1073/pnas.2009571117>
- Meng, X.-L. (2018). Statistical paradises and paradoxes in big data (I): Law of large populations, big data paradox, and the 2016 US presidential election. *The Annals of Applied Statistics*, 12(2), 685–726. <https://doi.org/10.1214/18-AOAS1161SF>
- Mesoudi, A. (2016). Cultural evolution: Integrating psychology, evolution and culture. *Current Opinion in Psychology*, 7, 17–22. <https://doi.org/10.1016/j.copsyc.2015.07.001>
- Neel, R., Kenrick, D. T., White, A. E., & Neuberg, S. L. (2016). Individual differences in fundamental social motives. *Journal of Personality and Social Psychology*, 110(6), 887–907. <https://doi.org/10.1037/pspp0000068>
- Pei, X., Li, Y., & Xu, C. (2023). *GPT self-supervision for a better data annotator* (arXiv:2306.04349). arXiv. <https://doi.org/10.48550/arXiv.2306.04349>
- Prot, S., Gentile, D. A., Anderson, C. A., Suzuki, K., Swing, E., Lim, K. M., Horiuchi, Y., Jelic, M., Krahé, B., Liuqing, W., Liao, A. K., Khoo, A., Petrescu, P. D., Sakamoto, A., Tajima, S., Toma, R. A., Warburton, W., Zhang, X., & Lam, B. C. P. (2014). Long-term relations among prosocial-media use, empathy, and prosocial behavior. *Psychological Science*, 25(2), 358–368. <https://doi.org/10.1177/0956797613503854>
- Raith, L., Bignill, J., Stavropoulos, V., Millea, P., Allen, A., Stallman, H. M., Mason, J., De Regt, T., Wood, A., & Kannis-Dymand, L. (2021). Massively multiplayer online games and well-being: A systematic literature review. *Frontiers in Psychology*, 12, Article 698799. <https://doi.org/10.3389/fpsyg.2021.698799>
- Rathje, S., Mirea, D.-M., Sucholutsky, I., Marjich, R., Robertson, C. E., & Van Bavel, J. J. (2024). GPT is an effective tool for multilingual psychological text analysis. *Proceedings of the National Academy of Sciences*, 121(34), Article e2308950121. <https://doi.org/10.1073/pnas.2308950121>
- Saad, G. (2011). *The consuming instinct: What juicy burgers, Ferraris, pornography, and gift giving reveal about human nature*. Prometheus Books.
- Safra, L., Chevallier, C., Grèzes, J., & Baumard, N. (2020). Tracking historical changes in perceived trustworthiness in Western Europe using machine learning analyses of facial cues in paintings. *Nature Communications*, 11(1), Article 4728. <https://doi.org/10.1038/s41467-020-18566-7>
- Schaller, M., Kenrick, D. T., Neel, R., & Neuberg, S. L. (2017). Evolution and human motivation: A fundamental motives framework. *Social and Personality Psychology Compass*, 11(6), Article e12319. <https://doi.org/10.1111/spc3.12319>
- Sheskin, M., Chevallier, C., Lambert, S., & Baumard, N. (2014). Life-history theory explains childhood moral development. *Trends in Cognitive Sciences*, 18(12), 613–615. <https://doi.org/10.1016/j.tics.2014.08.004>
- Soto, C. J., & John, O. P. (2017a). The next Big Five Inventory (BFI-2): Developing and assessing a hierarchical model with 15 facets to enhance bandwidth, fidelity, and predictive power.

Journal of Personality and Social Psychology, 113(1), 117–143.

<https://doi.org/10.1037/pspp0000096>

Soto, C. J., & John, O. P. (2017b). Short and extra-short forms of the Big Five Inventory–2: The BFI-2-S and BFI-2-XS. *Journal of Research in Personality*, 68, 69–81.

<https://doi.org/10.1016/j.jrp.2017.02.004>

Steege, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science*, 11(5), 702–712.

<https://doi.org/10.1177/1745691616658637>

Trepte, S., Reinecke, L., & Juechems, K. (2012). The social side of gaming: How playing online computer games creates online and offline social support. *Computers in Human Behavior*, 28(3), 832–839. <https://doi.org/10.1016/j.chb.2011.12.003>

Vize, C. E., Sharpe, B. M., Miller, J. D., Lynam, D. R., & Soto, C. J. (2023). Do the Big Five personality traits interact to predict life outcomes? Systematically testing the prevalence, nature, and effect size of trait-by-trait moderation. *European Journal of Personality*, 37(5), 605–625. <https://doi.org/10.1177/08902070221111857>

Yee, N. (2006). Motivations for play in online games. *CyberPsychology & Behavior*, 9(6), 772–775.

<https://doi.org/10.1089/cpb.2006.9.772>

Yee, N., Ducheneaut, N., & Nelson, L. (2012). Online gaming motivations scale: Development and validation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2803–2806). Association for Computing Machinery.

<https://doi.org/10.1145/2207676.2208681>