No credible evidence that UK safer gambling messages reduce gambling

Philip W. S. Newall¹

Leonardo Weiss-Cohen²

Henrik Singmann^{3,4}

Lukasz Walasek³

Elliot A. Ludvig³

1 Experimental Gambling Research Laboratory, School of Health, Medical and Applied Sciences, CQUniversity, 400 Kent St, Sydney, NSW 2000, Australia

2 Leeds University Business School, Maurice Keyworth Building, University of Leeds, Leeds, LS2 9JT, United Kingdom

3 Department of Psychology, University of Warwick, Coventry, CV4 7AL, United Kingdom

4 Experimental Psychology, University College London, 26 Bedford Way, London, WC1H, 0AP

Funding statement: This research was funded by a BA/Leverhume Small Research Grant no. SRG19\191299 awarded to Leonardo Weiss-Cohen, a Research Development Fund awarded to Lukasz Walasek from the University of Warwick, and by a Swiss National Science Foundation (SNSF) grant 100014_179121 awarded to Henrik Singmann.

Acknowledgement: Thanks to David Zendle for his useful comments and to Tyson Hayes for help in developing the online tasks.

Abstract

Safer gambling messages are a common freedom-preserving method of protecting individuals from gambling-related harm. Yet, there is a striking lack of independent and rigorous evidence to attest to the effectiveness of safer gambling messages. This study presents results from three large ($N \approx 3,000$), preregistered, and incentivized experimental tests of the UK's commonly used "When the fun stops, stop" gambling message. Variants of this message were tested on gamblers' propensity to accept bets on soccer events (Experiment 1), behavior on a realistic online roulette table (Experiment 2), and betting patterns on an online soccer betting platform (Experiment 3). All three experiments find either no beneficial effect or a small backfire effect on gamblers' behaviour. We conclude that independent evaluations should inform policy-makers' decisions on how to best implement improved safer gambling messages.

Keywords: public health, gambling policy, gambling regulation, safer gambling messages, independent testing, behavioral science

The UK has the world's largest regulated online gambling market (1). Gambling can be a normal pastime for some, but can potentially affect others (2,3), with an estimated 168,000 problem gamblers in England, and a further estimated 1,590,000 also at risk of experiencing gambling-related harm (4). Consequently, a key issue for gambling regulators is the trade-off between making gambling products freely available and protecting consumers' welfare, especially those who are vulnerable to harm. In 2019, the UK's Gambling Commission announced a new National Strategy to Reduce Gambling Harms (5). With its focus on prevention and education, the strategy set out a requirement that operators must use gambling messages to promote safer gambling behaviors, both at point of sale, and as part of broader, population-based campaigns. In addition, the strategy stresses the importance of rigorous empirical evaluation of safer gambling intervention, with a preference toward independent appraisal (i.e., not sponsored by gambling operators). In a similar vein, the UK government has recently issued a call for evidence to review "what, if any, is the evidence on the effectiveness of mandatory safer gambling messages in adverts in preventing harm?" (6). Yet, despite the urgent need to address the growing prevalence of gambling related harm, there are no studies that address these questions.

At the time of writing, "when the fun stops, stop" is the most commonly adopted UK gambling message (see Figure 1). This type of message can be encountered on television, in print, online, and in bookmakers' shop windows, and is said to be recognized by at least 54% of all UK adults (7,8). Similar messages are also used in other jurisdictions: The Nevada Council on Problem Gambling has trademarked the slogan "When the fun stops" (9), and the state monopoly gambling provider in Ontario, Canada, has a "PlaySmart" campaign which tells gamblers to "keep the fun in the game." Despite their prominence across different jurisdictions, to best of our knowledge, there is currently no independent evidence that these gambling messages influence gambling behaviours at all.

Figure 1. Typical gambling messages used in the UK and experimental stimuli used in the present experiments. (a) Yellow message used in the Experiment 1; (b) yellow message used in Experiment 2 and 3; (c) black-and-white message with de-emphasized “fun” also used in Experiment 2.

Outside the domain of gambling, the effectiveness of warning labels/messages has been tested in the areas of alcohol, tobacco, and financial decision-making (10-12). For example, research shows that consumers devote minimal attention to UK alcohol product warning labels (13), whereas financial market warnings, such as “past performance does not guarantee future returns”, do not appear to meaningfully shape investors’ choices (14). On the other hand, tobacco warnings are generally considered to be effective (15), especially when explicit messages designed to generate fear and disgust are paired with other interventions, such as plain packaging (16,17).

Despite the growing pressure to consider gambling a public health issue, little justification for the use of the “when the fun stops, stop” message exists. At the same time, this particular slogan has been criticized by political leaders (18), and the Gambling Commission’s CEO (19). One common observation is the emphasis that most versions of the message place on the word “fun” (See Figure 1a and 1b). A non-independent evaluation in 2017 argued that the message had led to 5.7 million adults, “...warn[ing] other people about

their gambling, if only jokingly” (7). This result, however, is based on a self-report measure and not a behavioral outcome.

To better assess the impact of this safer gambling message, we conducted three preregistered, incentivized online experiments on large samples of UK gamblers, comparing effects of multiple versions of the commonly used “when the fun stops, stop” message. In each experiment, gamblers were given a cash bonus and allowed to keep that bonus or to wager it on games of chance, which shared many features with those available on online gambling platforms. More specifically, gamblers could place bets on either real soccer games (Experiment 1 and 3) or on online roulette (Experiment 2). Safer gambling message presence was randomized between-participants. Experiment 1 was posted as a preprint in August 2019 (20). Consequently, at least one large gambling operator has used a new black-and-white message with a de-emphasized word “fun” written in the same size and font as the rest of the text (Figure 1c), an action that was attributed to being in response to Experiment 1’s reception (21). Therefore, the subsequent study (Experiment 2), also experimentally tested the effect of this novel black-and-white safer gambling message on gamblers’ behaviour.

In Experiment 1, 506 UK gamblers were presented with nine bets of £0.10 (each) on the upcoming FA cup final, the highest-profile regularly scheduled match in the domestic UK soccer calendar. Participants were shown either a bet description accompanied by the safer gambling message (Figure 1a) or no message at all. Their task was simply to indicate whether, for each bet, they would accept it (Bet) or not (Don’t bet). Bets differed in how specific (or precise) the prediction was concerning soccer events (low, medium, and high), with more specific ones offering higher returns. Our primary research question was whether the presence of the message would make gamblers less likely to accept a bet.

In Experiments 2 and 3, participants were given the option to gamble with larger endowments (£3), which they earned themselves by successfully completing an unrelated effortful task. The goal of this step was to assure that participants earn and feel ownership over their endowments (22). The task involved correctly transcribing seven or more out of 10 captcha codes (similar to when confirming that one is not a computer when visiting certain websites). The codes were not easy to complete, and so served a double purpose of screening-out inattentive participants (see Supplementary Information, Section 2.1): 317 of 1819 initial participants failed the task in Experiment 2, and 310 of 1318 initial participants failed the task in Experiment 3.

Participants who passed the captcha task proceeded to the next stage of each experiment, in which they could gamble their £3 endowment. Experiment 2 was based on a commercially-available online roulette game (Figure 2). This presentation format closely resembles both the physical roulette tables found in casinos and digital tables available to online gamblers. By comparison, many experimental tasks in gambling research use simplified games or graphical representations (23-27). Participants could place many different types of bets on the roulette table, using different values of tokens, and play as many spins of the roulette as they wanted. Experiment 3 also used a realistic soccer betting platform. In Experiment 3, participants could browse and bet on the odds of 921 unique bets distributed across the upcoming round of 10 English Premier League soccer matches (Figure 3), collected from a large online UK betting operator's website. Participants could browse these bets and their odds across 10 unique pages, calculate potential payouts of bets from different sizes, and confirm their selection before checking out, just like when betting on soccer on gambling operators' websites. Participants were free to bet as much of their endowment as they wanted and keep the rest, or they could keep it all without placing any bets. All payments were made based on participants' choices, and any bets placed were paid

according to random outcomes of the roulette or actual soccer match results. In all experiments (1-3), gambling messages were displayed prominently with the relevant logos. All experiments included a control condition with no message. Experiment 1 included one treatment group with the “when the fun stops, stop” message (Figure 1a), Experiment 2 used two treatment conditions (Figure 1b/1c), and Experiment 3 used one treatment condition (Figure 1b).

Figure 2. Screenshot of the roulette task (yellow message condition). The top screenshot illustrates various potential bets: one £0.10 chip on all black numbers, one £0.50 chip on the first street (numbers 1,2,3) and a simple bet of 2 x £0.20 chips on the number 36, for a total bet amount of £1, the maximum allowed per spin. In the lower screenshot, the roulette wheel is shown spinning before settling on a randomly selected number after the participant clicked the "spin" button.

Figure 3. Screenshot of the soccer betting platform (yellow message condition). The screenshot illustrates a selection of possible bets from a single match, here between Everton and Liverpool (more bets were visible when a participant scrolled down this page). Participants could click on any event, specifying the amount of money they would like to bet. Selected bets would be added to the Bet Slip total and subtracted from the Remaining Credit (top right). At any point a participant could continue (bottom right), either if they were satisfied with their bets or if they did not want to bet at all.

Results

Hypotheses and analyses for all experiments were pre-registered and can be found at <https://osf.io/fwunh/>

Experiment 1

The primary dependent variable in Experiment 1 was participants' decisions to accept (or reject) a series of soccer bets varying in their specificity (and payoffs).

Participants chose to place a bet more often in the condition where the message was present (41.3%) than in the control condition (37.8%). Following our pre-registered analysis

plan, we report results from the best generalized linear mixed model (GLMM) with binomial family and logistic link function, with by-participant random-intercepts and no further random-effects terms. Results of pre-registered models with different fixed- and random-effects structures can be found on the Open Science Framework (OSF) repository, and all show the same pattern of significant results with only very small numerical differences to the results reported here.

The mild increase in willingness to place a bet following the gambling message label was not significant, $\chi^2(1) = 2.10$, $p = .15$. This corresponds to an effect size (odds ratio) of $OR = 1.22$ [95% CI: 0.93, 1.61]. Following our a priori power analysis this suggests that any effect of the message is likely smaller than 6%, independent of the direction (which in the present case was opposite to the intention of the message).

The specificity (and thus odds) of the gamble showed a significant effect on gambling behaviour, $\chi^2(2) = 34.05$, $p < .0001$. This effect was non-monotonic: the predicted probability to gamble was .37, [95% CI: .33, .41] for low-specificity gambles, .31, [.27, .34] for medium-specificity gambles, and .42, [.38, .46] for high-specificity gambles. The interaction between message presence and specificity was not significant, $\chi^2(1) = 1.71$, $p = .41$.

Exploratory results of models that include effects of problem gambling severity, previous year gambling frequency, and gender, did not change the findings reported above (see OSF repository for details). Thus overall, we found no evidence that the presence of a gambling message influenced people's propensity to gamble on sporting events.

Experiment 2

In Experiment 2, our dependent variable was the proportion of money bet by each participant on the roulette game. This dependent variable shows a clearly multimodal distribution with three peaks, reflecting that a) some participants do not gamble at all, b)

some bet all their endowment, and c) others bet a proportion of their endowment on roulette bets. Given this distribution, we applied a Bayesian zero-one-inflated beta regression (ZOIBR) model across the three message conditions. This model therefore explores whether the message has an impact on participants' propensity to gamble at all, the propensity to bet everything, and the mean proportion of money bet (if they bet at all but do not bet everything).

Figure 4 (top row) shows the posterior distributions of the parameter estimates from the preregistered ZOIBR model across the three message conditions, which represent the uncertainties for the message-condition-specific parameter estimates given the observed data. Due to the large sample size ($N=1,500$), the precision of the parameter estimates is high (i.e., uncertainties are low). Nevertheless, the differences between the three messages seem comparatively small. To facilitate comparison across conditions, the bottom row of Figure 4 shows the difference distributions between the two message conditions from the no-message control condition.

Figure 4. ZOIBR model estimates for data in Experiment 2. Top row panels show the posterior distributions of condition-specific parameter estimates. Bottom row panels show posterior difference distribution for comparison of the two message conditions from the no-message control condition. If a difference distribution does not include 0 (which is indicated by a grey line), this would constitute evidence for a difference between the corresponding message condition and the no-message condition. The grey area shows the full posterior distribution, the black dot underneath shows the median, the thick line shows the 66% credibility interval, and the thin line shows the 95% credibility interval.

Figure 4 (left column) shows how the presence of the messages impacted the probability to gamble at all. The visual impression suggests that the messages could have a slight protective effect. Participants who see the messages are on average around 2% less likely to gamble at all. As shown in the lower row of Figure 4, however, the uncertainty of the difference is substantially larger than 2% and thus a statistically meaningful difference is not supported by the model. There might be a slight protective effect, but, if true, the effect is so small that it would require even larger sample sizes to be detected reliably.

Figure 4 (middle column) shows how the messages clearly had no effect on the conditional probability of gambling everything. The difference between the message and no-message conditions is on average only around 1%. In addition, the difference distributions peak close to zero, indicating that the data provides evidence for no difference between the conditions for any reasonable prior distribution.

Figure 4 (right column) depicts how the messages affected the mean proportion bet. Both visual impression and statistical results provide evidence for a backfire effect with the yellow message. The mean of the difference distribution between the no-message condition and the yellow message is around 4% and the 95% credibility interval of the difference distribution just excludes zero (the lower bound is 0.003%). In contrast, the mean of the difference distribution between the no-message condition and the white message is lower at 2% and the 95% credibility interval of the difference distribution does include zero. Thus, the data provides evidence that, when presented with the yellow message, participants' proportion of money bet is larger compared to the no-message condition.

The confirmatory analysis did not include Bayesian priors due to the difficulty of specifying what a good prior width is, and as default priors do not exist for bespoke models such as the ZOIBR model employed here. However, we considered the evidence provided by the posterior distribution of the confirmatory model in light of different possible priors. The analysis showed that even with extremely tight difference priors ($SD = 0.05$), there is evidence for the null of no-difference in all but one case, the backfire effect of the yellow message on the proportion bet (See Supplementary Information, Section 2.2).

Exploratory analyses were also performed on participants' individual difference scores on the Problem Gambling Severity Index (PGSI) and the Enhancement Motives subscale of the Gambling Motives Questionnaire (ENH). As both of those scales were

collected after the roulette task and the experimental manipulation, the possibility exists that participants' scores might be affected by the manipulation. To test for this, two Bayesian ANOVAs using default priors were estimated (28) with the two questionnaire scores as dependent variables and message condition as independent variable. The Bayesian ANOVAs provided strong evidence that the manipulation did not affect the scores, with Bayes factors for the null above 30 for both (see Supplementary Information, Table S1 in Section 2.3).

Exploratory analysis was therefore performed to see if the main results hold when adjusting (or controlling) for participants' two questionnaire scores. Thus, the main analysis reported above was repeated twice. In one analysis, the mean-centred scores of both of the questionnaires were included as additive effects, while another analysis also added interactions of the questionnaires with the condition effect (Supplementary Information Table S2 in Section 2.3). Both of those models showed the same pattern as reported earlier, after controlling for their individual gambling-related scores. The backfire effect of the yellow message remained, with the respective 95% confidence intervals still excluding zero in both additional models.

Finally, different bets in roulette are associated with different levels of risk, as indicated by their likelihood of being paid out, and the magnitude of the payoff. For example, if £1 is bet, then this can provide a total payoff between £2 (e.g., bets on red or black) and £36 (bets on a single number). A preregistered exploratory analysis investigated if the gambling messages would affect the riskiness of bets chosen. A beta regression model showed evidence of no differences in the riskiness of bets chosen across the three experimental conditions (see Supplementary Materials Section 2.4).

Experiment 3

Experiment 3 examined the impact of the safer gambling messages on gambling on an online football betting platform. As in Experiment 2, we applied the ZOIBR model to determine whether the presence of a message influenced people's propensity to gamble at all, the propensity to bet everything, and the mean proportion of money bet (if betting at all but not betting everything).

Figure 5 shows the results of the ZOIBR model in Experiment 3. Despite shifting from the roulette task to an online soccer betting platform, the results are strikingly similar between the two studies (and conceptually similar to Experiment 1).

Figure 5. ZOIBR model estimates for data in Experiment 3. Top row panels show the posterior distributions of condition-specific parameter estimates. Bottom row panels show posterior difference distribution for comparison of the yellow message condition from the

no-message control condition. If a difference distribution did not include 0 (which is indicated by a grey line), this would constitute evidence for a difference between the corresponding message condition and the no-message condition. The grey area shows the full posterior distribution, the black dot underneath shows the median, the thick line shows the 66% credibility interval, and the thin line shows the 95% credibility interval.

Figure 5 (top left) shows how the protective effect of the message on the probability to gamble at all is even smaller than that found in Experiment 2. In the bottom panel of Figure 5, we can see that the uncertainty around the 1% difference clearly crosses over 0, indicating no reliable difference between the conditions. Figure 5 (middle column) shows how the largest difference appears when we assess whether the message affects the conditional probability to gamble everything, but it is in the opposite direction to the desired effect of the safer gambling message. Participants who saw the “when the fun stops, stop” message were 6% more likely to gamble all of their money if they gambled at all. Despite the effect size, the 95% credibility interval does include 0. Finally, does the message affect the mean proportion bet? The answer is no, as it is clearly visible in Figure 5 (right column) that the difference between the condition is very close to 0 (-0.4% difference).

As in Experiment 2, we explored the evidence provided by the ZOIBR model in light of different possible priors for the difference between the two message conditions. The results replicate Experiment 2 that even for very narrow priors we either find evidence for the null or slight evidence for a backfire effect (see Supplementary Information Table S4, in Section 3.2). Likewise, including the two gambling scales as covariates does not alter the conclusions (Supplementary Information Table S5 in Section 3.3). The final result concerns the relationship between the treatment and level of riskiness among people who placed at least one bet. Beta regressions showed no effect of the manipulation, consistent with Experiment 2 (see Supplement 3.4).

Discussion

The UK's gambling regulator stresses the importance of independently evaluating safer gambling messages (5). The results from our three incentivized experiments on the relevant population of interest provided no credible evidence that gambling messages based on the phrase "when the fun stops, stop" prompt safer concurrent gambling behaviour. Any potential improvements on gamblers' decisions to gamble at all were too small to be reliably detected here and can be contrasted with the largest effect in each experiment being a backfire effect. The results might be considered relevant to other jurisdictions, where similar messages can be observed. For example, "when the fun stops" used in Nevada (US) or "keep the fun in the game" used in Ontario (Canada), might be similarly ineffective. Meanwhile, in Australia the phrase "gamble responsibly" is frequently placed on marketing materials, but without attracting much attention from gamblers, suggesting that it may also have little effect on gambling behaviour (29).

The present results are limited to a group of UK gamblers and to immediate behavioral outcomes. It might be argued that UK gamblers have seen "when the fun stops, stop" messages so often that improvements in safer gambling had already taken effect before the experimental manipulation. Although participants in all experiments were experienced gamblers, no direct measurements of gambling frequency were taken. PGSI is, however, a proxy for gambling frequency, given that problem gamblers tend to gamble much more than recreational gamblers (30). Although the results showed no differential effect of gambling messages across PGSI score, a replication study using gamblers from a different jurisdiction could perhaps best test this potential explanation. Contrastingly, repeated exposure to an identical message may create psychological reactance that could lead to backfire effects (31). This is a potential explanation that appears more consistent with the observed results. Furthermore, these messages may have weak immediate effects, but stronger effects with repeated exposure over longer timescales. A longitudinal study could test this explanation.

The strength of our contribution is in the use of three different gambling tasks in which gamblers could choose how much (if at all) they would like to bet on various online gambling games. Despite diversity across tasks and messages used, the results were consistent in showing no effect of the messages. Whereas our objective was to maximize the ecological validity of evaluating gambling behavior, this could only be done within the constraints of ethical considerations. As a result, we could not test whether the messages might have had an effect in the most hazardous situations, such as when a person continues betting more than they can afford, having incurred significant losses (i.e., chasing losses). Such a limitation merely further highlights the need for evidence to inform policy in cases where individuals may engage in risky behaviours with severe consequences for their health and wellbeing.

Nevertheless, despite these limitations, there would seem to be sufficient evidence for the Gambling Commission and other actors to consider trialling different safer gambling messages. Although the gambling industry could assist in trialling different safer gambling messages, evidence from tobacco and alcohol suggests that industry actors are more likely to slow down rather than facilitate the widespread adoption of effective messages (32,33). Alternative messages could, for example, highlight that there are an estimated current 168,000 problem gamblers in England, with a further estimated 1,590,000 also at risk of experiencing gambling-related harm (4). Safer gambling messages could also highlight the diffuse pathways through which gambling-related harm can operate (34,35), or how gambling can harm affected-others too (36).

Independent evaluation of safer gambling messages can provide an important input to evidence-based gambling policy (5,37,38). The present paper makes such a contribution, offering no credible evidence for the effectiveness of “when the fun stops, stop” safer gambling messages.

Methods

Each study was pre-registered. Anonymized data, analysis output, materials and the preregistration document can be accessed from the OSF repository available at <https://osf.io/fwunh/>.

Participants were paid a baseline fee of £0.50, £2.50, and £1.50 in Experiment 1, 2, and 3, respectively. In addition, participants received additional money that could be used to spend on gambling. In Experiment 1, all participants were endowed with £0.90, whereas in Experiments 2 and 3, participants could earn £3 by completing an unrelated task. In Experiments 2 and 3, only those who correctly completed the transcription of seven or more captcha codes earned the additional funds. No participant was required to gamble in any study. Each study received ethical approval from the University of Warwick's Human Research Ethics Committee (HSSREC).

Design

In each experiment, participants were randomly allocated to groups that completed the gambling task with or without a safer gambling message. The treatment group in Experiment 1 included the yellow message shown in Figure 1a. Experiments 2 and 3 used a slightly different version of the yellow message, shown in Figure 1b, based on feedback from an industry actor (39). Finally, Experiment 2 also included a third treatment group in which participants were presented with the black-and-white message with the de-emphasized word “fun” (Figure 1c). In Experiment 1, we also varied the specificity of the described soccer bets (low, medium, high; see the OSF repository for details).

Experiment 2 and 3's main (pre-registered) outcome measure was proportion of available funds bet, which was defined as the total amount of money bet by a participant out of the total that could have been bet:

$$\text{Proportion bet} = \text{total amount bet} / (\text{initial } \pounds 3 \text{ endowment} + \text{all winnings})$$

Proportion bet has a minimum of 0, for a participant who bet nothing, and a maximum of 1, for a participant who bet everything they possibly could (including in Experiment 2 any money they had won in addition to the initial endowment) and lost it all. This dependent variable therefore takes participants' luck into account.

Participants

Participants were recruited and paid via Prolific Academic. All were UK nationals, and had, in the case of Experiments 1 and 3, indicated prior experience in online sports betting. Participants in Experiment 1 also previously reported an interest in UK Premier League soccer on Prolific Academic (this filter was not used for Experiment 3 as it would have overly restricted the available pool given the larger sample size of $N=1000$ targeted in that experiment). In Experiment 2, eligible participants indicated prior experience in playing online roulette.

For Experiment 1, 506 unique participants took part (Age: $M=36$; Female: $N=168$), with an average completion time of 4.4 minutes, and an average bet payout of $\pounds 0.68$, which added to the fixed payment of $\pounds 0.50$, resulted in an equivalent pay rate of $\pounds 9.36/\text{hour}$. For Experiment 2, 1,819 participants took part in the real-effort captcha task, but 317 transcribed six or fewer captchas correctly and did not proceed to the roulette stage. The data from two participants were excluded because of missing data, resulting in a final sample size of 1,500 (Age: $M=32$; Female: $N=578$). The average completion time was 7.2 minutes and the average

final roulette payout was £2.80 (range: [£0, £30]), which added to the fixed payment of £2.50 resulted in an equivalent pay rate of £44/hour. For Experiment 3, 1,318 participants took part in the real-effort captcha task, but 310 transcribed six or fewer captchas correctly and did not proceed to the football betting stage. The data from four participants had to be excluded because they subsequently abandoned the task before checking out their bets and completing the final gambling questionnaire, and one participant had to be excluded due to missing captchas data, resulting in a final sample size of 1,003 (Age: $M=35$; Female: $N=351$). The average completion time was 5.2 minutes, and the average final bonus payout was £2.20 (range: [£0, £33]), which added to the fixed payment of £1.50 resulted in an equivalent pay rate of £43/hour.

Experiment 1's sample size was justified via a simulation-based power analysis, using a generalized linear mixed model (GLMM) with a binomial family, indicating that the experiment had a minimum power of 80% to detect a difference in probability of accepting the bet of roughly 6% (the further away the average probability of accepting the bet is from 50%, the larger the power to detect an effect of 6% difference). It was not feasible to run a power analysis based on the necessary statistical model for Experiments 2 and 3, so the sample size was set to exceed the 322 participants per group required to detect a small effect size with power = 0.8 using a traditional ANOVA model. As our funds allowed us to collect around 500 participants per condition we opted to do so.

Procedure

Experiment 1

In Experiment 1 participants were presented with a random order of nine mock gambling adverts (three at low, medium, and high level of specificity), which in the UK can be seen on TV displaying the odds on specific bets, e.g., "Manchester City to win 4-1. Win

£1.70” (40). All bets were based on the upcoming 2019 FA Cup final between Manchester City and Watford, the highest-profile match of the domestic soccer calendar. Payoff sizes were based on the total payoff based on a bet of £0.10 receiving currently available betting odds from a major gambling operator’s website, and participants received all bonuses the day after the match. The text of each bet and the corresponding potential payoff were taken from a major gambling operator’s website. Safer gambling message presence was manipulated between participants. Eleven percent of participants declined to place any bets, thereby keeping their entire £0.90 bonus. After the nine trials, participants self-reported demographics and completed the Problem Gambling Severity Index (41).

Experiment 2

The main task consisted of six individual pages (see Supplementary Information, Section 2.1 for more information). Upon correct transcription of at least seven captchas, participants were informed that they had earned a further £3 and were given details of the roulette game. Fourteen percent of participants left the roulette without placing any bets, thereby keeping their entire £3 bonus. This page also included information on how to place bets. The roulette game was shown on the next page. Gambling message presence was manipulated between participants, with messages shown both on the instructions page and overlaid in the top-right corner of the roulette game (see Figure 2).

The roulette game was configured to resize based on screen resolution, so that it would always be clearly visible in its entirety. Participants could play as many spins of the roulette as they wanted, by betting between £0.10 and £1 per spin, using one or more betting chips of different values. The game was a realistic version of online roulette, programmed in JavaScript. On average, participants who gambled played for 9 spins (range: [1, 204]). The winning number was randomly generated by the server at each spin. Participants were

allowed to play roulette for as long as they wanted, with the exception that any participant who managed to reach a balance of zero (lost everything) or more than £90 would have been moved on from the roulette game (no one reached the upper limit; the maximum balance reached at any point was £33).

After the roulette game, participants completed the Problem Gambling Severity Index and the Enhancement Motives subscale of the Gambling Motives Questionnaire (42).

Experiment 3

The procedure closely resembled that of Experiment 2, with the exception that participants could place bets on a selection of soccer matches rather than roulette (see Figure 3). Detailed description of the soccer gambling platform is provided in Supplementary Information Section 3.1. After successfully transcribing captchas, participants were informed that they can use their £3.00 bonus to place bets on the upcoming round of the English Premier League matches. The betting platform was programmed in Python using oTree. The message in the treatment group was displayed on top of each screen of the betting platform.

The betting platform started with pop-up instructions on how to place bets. After reading the instructions, participants could place as many individual bets as they wanted, as long as the total amount bet was below £3, or they could proceed without placing any bets. They could bet on 10 different matches, and within each match they could choose between an average of 40 unique bets, across three subcategories: winning team, exact score, and first goal-scorer. Thirty-nine percent of participants did not place any bets. Participants who decided to bet placed an average of 2.2 bets. The payments for all bets were calculated based on the actual outcomes of the matches. All bonuses were paid after all the matches had been played, including any amounts that were not bet, to ensure there was no time-discounting benefit to not placing a bet.

To ensure all odds were ecologically valid, the odds for each individual bet were extracted from a leading high-street bookmaker's website on the morning before data collection started. Full description of the bets and their odds can be found on the OSF repository. Bet amounts started at £0.01 and could change in that increment. Potential payouts were calculated automatically and shown to participants according to the odds and the amount they wanted to bet. Participants could review and edit their bet slip, until they pressed on the "Continue" button to save and submit their bets. Upon submitting their bet slip, participants completed the Problem Gambling Severity Index measure and the Enhancement Motives subscale of the Gambling Motives Questionnaire.

Data analysis

As per the preregistration, the analysis of data from Experiment 1 involved running a series of generalized mixed models (GLMMs) with binomial family and logistic link function. These models included fixed effects of condition and specificity, as well as by-participant random-intercepts. Results of different preregistered variants of the models can be accessed on the OSF, but these all yielded similar results. A Bayesian analysis similar to the ones preregistered for Experiments 2 and 3 (reported on OSF) also showed no effect of the message condition.

In Experiments 2 and 3, the statistical analysis was performed using a preregistered zero-one-inflated beta regression (ZOIBR) model implemented in a Bayesian statistical framework using the *brms* (43) package in R (44). The ZOIBR model was parameterized as follows:

g : probability to gamble at all (i.e., probability of proportion bet being greater than 0)

e : conditional probability to gamble everything (i.e., conditional probability of proportion bet being equal to one, if one gambles)

μ (mu): mean proportion gambled if participants gamble, but do not gamble everything (i.e., mean of beta regression model)

ϕ (phi): precision of the beta regression model

The beta regression with parameters μ and ϕ only matters for the case when the response is neither a zero nor a one (i.e., with probability $g \times (1 - e)$). The three main parameters (i.e., g , e , and μ) were allowed to differ across experimental conditions. This allowed a direct test for the following three patterns for each of the two messages.

RQ1: Does the message affect the probability to gamble at all (i.e., test of g)?

RQ2: Does the message affect the conditional probability to gamble everything, if one gambles (i.e., test of e)?

RQ3: Does the message affect the conditional mean proportion bet, if participants gamble, but do not bet everything (i.e., the mean of the beta distribution; test of μ)?

References

- (1) Gambling Commission. Review of online gambling. 2018; Available at: <http://www.gamblingcommission.gov.uk/PDF/Online-review-March-2018.pdf>, 2020.
- (2) Muggleton N, Parpart P, Newall P, Leake D, Gathergood J, Stewart N. The association between gambling and financial, social, and health outcomes in big financial data. *Nature Human Behaviour* 2021.
- (3) Browne M, Langham E, Rawat V, Greer N, Li E, Rose J, et al. Assessing gambling-related harm in Victoria: a public health perspective. Melbourne: Victorian Responsible Gambling Foundation; 2016.
- (4) NHS Digital. Health Survey for England 2018: Supplementary analysis on gambling. 2019; Available at: <https://digital.nhs.uk/data-and-information/publications/statistical/health-survey-for-england/2018/health-survey-for-england-2018-supplementary-analysis-on-gambling>, 2020.
- (5) Gambling Commission. National Strategy to Reduce Gambling Harms. 2019; Available at: <http://www.reducinggamblingharms.org/asset-library/national-strategy-to-reduce-gambling-harms.pdf>, 2019.
- (6) DCMS. Review of the Gambling Act 2005 Terms of Reference and Call for Evidence. 2020; Available at: <https://www.gov.uk/government/publications/review-of-the-gambling-act-2005-terms-of-reference-and-call-for-evidence/review-of-the-gambling-act-2005-terms-of-reference-and-call-for-evidence>, 2021.
- (7) The Senet Group. Report on campaign evaluation for The Senet Group. October 2017. 2017; Available at: <https://senetgroup.org.uk/wp-content/uploads/2017/11/senet-group-report-sept-2017-v2.pdf>, 2019.
- (8) Critchlow N, Moodie C, Stead M, Morgan A, Newall PW, Dobbie F. Visibility of age restriction warnings, harm reduction messages, and terms and conditions: a content analysis of paid-for gambling advertising in the United Kingdom. *Public Health* 2020.
- (9) Nevada Council on Problem Gambling. WHEN THE FUN STOPS - Trademark Details. 2016; Available at: <https://trademarks.justia.com/867/14/when-the-fun-86714704.html>, 2020.
- (10) Argo JJ, Main KJ. Meta-analyses of the effectiveness of warning labels. *Journal of Public Policy & Marketing* 2004;23(2):193-208.
- (11) Rogers WA, Lamson N, Rousseau GK. Warning research: An integrative perspective. *Hum Factors* 2000;42(1):102-139.
- (12) Stewart DW, Martin IM. Intended and unintended consequences of warning messages: A review and synthesis of empirical research. *Journal of Public Policy & Marketing* 1994;13(1):1-19.

- (13) Kersbergen I, Field M. Alcohol consumers' attention to warning labels and brand information on alcohol packaging: Findings from cross-sectional and experimental studies. *BMC Public Health* 2017;17(1):123.
- (14) Mercer M, Palmiter AR, Taha AE. Worthless warnings? Testing the effectiveness of disclaimers in mutual fund advertisements. *Journal of Empirical Legal Studies* 2010;7(3):429-459.
- (15) Hammond D. Health warning messages on tobacco products: a review. *Tob Control* 2011 Sep;20(5):327-337.
- (16) Maynard OM, Munafò MR, Leonards U. Visual attention to health warnings on plain tobacco packaging in adolescent smokers and non-smokers. *Addiction* 2013;108(2):413-419.
- (17) Munafò MR, Roberts N, Bauld L, Leonards U. Plain packaging increases visual attention to health warnings on cigarette packs in non-smokers and weekly smokers but not daily smokers. *Addiction* 2011;106(8):1505-1510.
- (18) Gaming Intelligence. Lib Dems call for 'when the fun stops, stop' slogan to be scrapped. 2019; Available at: <https://www.gamingintelligence.com/marketing/advertising/47077-lib-dems-call-for-when-the-fun-stops-stop-slogan-to-be-scrapped/>, 2021.
- (19) McArthur N. Gambling Commission makes online gambling safer. 2018; Available at: <https://www.gamblingcommission.gov.uk/news-action-and-statistics/News/gambling-commission-makes-online-gambling-safer>, 2019.
- (20) Newall PWS, Walasek L, Singmann H, Ludvig EA. Testing a gambling warning label's effect on behavior. 2019; Available at: <https://psyarxiv.com/dxfkj/>.
- (21) Glenday J. Warning slogan in gambling ads has little effect on efforts to influence behaviour. 2019; Available at: <https://web.archive.org/web/20190805163105/https://www.thedrum.com/news/2019/08/05/warning-slogan-gambling-ads-has-little-effect-efforts-influence-behaviour>, 2020.
- (22) Erkal N, Gangadharan L, Nikiforakis N. Relative earnings and giving in a real-effort experiment. *Am Econ Rev* 2011;101(7):3330-3348.
- (23) Cherkasova MV, Clark L, Barton JJS, Schulzer M, Shafiee M, Kingstone A, et al. Win-Concurrent Sensory Cues Can Promote Riskier Choice. *J Neurosci* 2018 Nov 28;38(48):10362-10370.
- (24) Eben C, Chen Z, Vermeylen L, Billieux J, Verbruggen F. A direct and conceptual replication of post-loss speeding when gambling. *Royal Society Open Science* 2020;7(5):200090.
- (25) McGivern P, Hussain Z, Lipka S, Stupple E. The impact of pop-up warning messages of losses on expenditure in a simulated game of online roulette: a pilot study. *BMC Public Health* 2019;19(1):822.

- (26) Spetch ML, Madan CR, Liu YS, Ludvig EA. Effects of winning cues and relative payout on choice between simulated slot machines. *Addiction* 2020.
- (27) Tobias-Webb J, Limbrick-Oldfield EH, Vearncombe S, Duka T, Clark L. The effects of alcohol on sequential decision-making biases during gambling. *Psychopharmacology (Berl)* 2020;237(2):395-407.
- (28) Rouder JN, Morey RD, Speckman PL, Province JM. Default Bayes factors for ANOVA designs. *J Math Psychol* 2012;56(5):356-374.
- (29) Lole L, Li E, Russell AM, Greer N, Thorne H, Hing N. Are sports bettors looking at responsible gambling messages? An eye-tracking study on wagering advertisements. *Journal of behavioral addictions* 2019;8(3):499-507.
- (30) Rockloff MJ. Validation of the consumption screen for problem gambling (CSPG). *Journal of Gambling Studies* 2012;28(2):207-216.
- (31) Cacioppo JT, Petty RE. Effects of message repetition and position on cognitive response, recall, and persuasion. *J Pers Soc Psychol* 1979;37(1):97-109.
- (32) Hiilamo H, Crosbie E, Glantz SA. The evolution of health warning labels on cigarette packs: the role of precedents, and tobacco industry strategies to block diffusion. *Tob Control* 2014 Jan;23(1):e2.
- (33) Stockwell T, Solomon R, O'Brien P, Vallance K, Hobin E. Cancer warning labels on alcohol containers: A consumer's right to know, a government's responsibility to inform, and an industry's power to thwart. *Journal of Studies on Alcohol and Drugs* 2020;81(2):284-292.
- (34) Langham E, Thorne H, Browne M, Donaldson P, Rose J, Rockloff M. Understanding gambling related harm: A proposed definition, conceptual framework, and taxonomy of harms. *BMC Public Health* 2016;16(1):80.
- (35) Markham F, Doran B, Young M. The relationship between electronic gaming machine accessibility and police-recorded domestic violence: A spatio-temporal analysis of 654 postcodes in Victoria, Australia, 2005–2014. *Soc Sci Med* 2016;162:106-114.
- (36) Goodwin BC, Browne M, Rockloff M, Rose J. A typical problem gambler affects six others. *International Gambling Studies* 2017;17(2):276-289.
- (37) van Schalkwyk MCI, Cassidy R, McKee M, Petticrew M. Gambling control: in support of a public health response to gambling. *Lancet* 2019 Apr 27;393(10182):1680-1681.
- (38) Wardle H, Reith G, Langham E, Rogers RD. Gambling and public health: we need policy action to prevent harm. *Br Med J* 2019;365.
- (39) Gilmore O. Our message works, whatever you might have read. 2019; Available at: <https://www.campaignlive.co.uk/article/message-works-whatever-read/1593157>, 2020.
- (40) Newall PWS, Thobhani A, Walasek L, Meyer C. Live-odds gambling advertising and consumer protection. *PLOS One* 2019.

- (41) Ferris J, Wynne HJ. The Canadian problem gambling index: Final report. Ottawa, ON: Canadian Centre on Substance Abuse; 2001.
- (42) Stewart SH, Zack M. Development and psychometric evaluation of a three-dimensional Gambling Motives Questionnaire. *Addiction* 2008;103(7):1110-1117.
- (43) Bürkner P. brms: An R package for Bayesian multilevel models using Stan. *Journal of statistical software* 2017;80(1):1-28.
- (44) R Core Team. R: A language and environment for statistical computing. 2016; Available at: <https://www.R-project.org/>.

Supplementary information

Introduction

This supplementary information material contains more details on Experiments 2 and 3 of the manuscript “No credible evidence that UK safer gambling messages reduce gambling.” The document is subdivided into sections, one for each experiment. Because these two experiments were similar and the analyses conducted for both followed the same preregistered plan, in-depth details are only given when they are first mentioned in Experiment 2, and only differences and specific results are provided for Experiment 3. All raw data, analysis code, pre-registration, and more information can be found online on the project’s OSF repository available at <https://osf.io/fwunh/>.

Experiment 2

Experimental design

Experiment 2 consisted of six individual pages. Figure S1 shows the flow of Experiment 2 through the pages, with the average time spent on each page. Screen captures of all pages are available online (<https://osf.io/fwunh/>).

Figure S1. Flow of Experiment 2. Each box represents an individual page of the experiment. All participants had to complete the consent form and type in the captchas. Only participants who completed seven or more correct captchas proceeded into the roulette game part of the experiment. Two participants who correctly filled in the captchas were excluded from the analyses due to missing data. The average time spent on each page is shown in brackets.

A total of N=1,819 participants started the task by reading the consent form in which they were informed that high performance on the captcha task would yield “a bonus that can then be used to play an online roulette game.” An initial sample of N=500 was discarded because participants were not explicitly told about the roulette aspect of the bonus. After giving consent, participants then proceeded to the captchas page, which consisted of ten captchas that had to be transcribed by the participant. Capitalization was ignored. The captchas shown to each participant were randomly selected from a full set of 21 captchas. The captchas were created in Python using the “captcha” library (1), and each was 160 pixels wide by 60 pixels high. The full set is available online (<https://osf.io/fwunh/>), and examples are shown in Figure S2. Participants (N=317) who typed in six or fewer correct captchas were immediately moved to the final thank-you page which ended the task and paid £2.50. Participants (N=1,502) who typed in seven or more correct captchas were awarded a bonus of £3.00, which they could play in a roulette game, in the next phase of the experiment. Two participants were excluded from further analyses due to missing data (one missing bet and one missing questionnaire data), for a total sample of N=1,500 used for analysis.

Figure S2. Example of captchas used in Experiments 2 and 3. A total of 21 captchas were created in Python, and ten were randomly selected for each participant.

The roulette phase started with a page with instructions on how to play the roulette game. The full instructions are available online (<https://osf.io/fwunh/>). Participants were told that they could play as often or as little as they liked. The instructions page also included information on how to place bets and operate the game. Gambling message presence was manipulated between participants, with messages shown both on the instructions page and overlaid in the top-right corner of the roulette game. After reading the instructions, participants could then proceed to the roulette game itself on the next page.

The roulette game was configured to resize based on screen resolution, so that it would always be clearly visible in its entirety. Participants with small screens (smaller than 500 pixels available for the browser in either height or width) were not allowed to participate in the experiment, which excludes most mobile phones and small tablets. Participants with larger screens but smaller browser windows were told to increase the size of the browser window above 500 pixels in both dimensions before continuing. Participants were not required to gamble. Fourteen percent of participants left the roulette without placing any bets, thereby keeping their entire £3 bonus. Participants could play as many spins of the roulette as they wanted, by betting between £0.10 and £1 per spin, using one or more betting chips of different values (£0.1, £0.2, £0.5, £1). The game was a realistic version of roulette, allowing individuals to place diverse types of bets, such as bets on a single number, red/black,

odds/evens, highs/lows, and more complicated bets such as splits (two numbers), streets (three numbers), corners (four numbers), columns, and others. The winning number was randomly generated by the server at each spin.

Participants were allowed to play roulette for as long as they wanted, with the exception that any participant who managed to reach a balance of zero (lost everything) or more than £90 would be moved on from the roulette game (no one reached the upper limit; the maximum balance reached at any point was £33). Participants could proceed by clicking the "close roulette" button. Participants who closed the roulette by accident were given a chance to reopen the roulette; otherwise, after closing the roulette they could continue to the questionnaire page.

The questionnaire page contained 13 questions, split across the five questions of the Enhancement Motives scale (ENH) of the Gambling Motives Questionnaire, followed by the eight questions of the Problem Gambling Severity Index (PGSI). Once all the questions had been answered, participants were shown the thank-you page which ended the task.

Participants who passed the captchas and played the roulette game were paid £2.50 plus whatever bonus they had when they closed the roulette (average bonus: £2.80, range: [£0, £30]).

Exploratory analysis: Bayes factors supplement

The confirmatory analysis did not include Bayesian priors due to the difficulty of specifying a good prior width. Default priors do not exist for bespoke models such as the ZOIBR model employed here. Without default priors, it is questionable whether reasonable Bayes factors can be calculated. Thus, calculating Bayes factors for the ZOIBR model is beyond the scope of the present work. However, the density estimates for each hypothesis of the confirmatory

analysis (gamble at all, gamble everything, and proportion gambled, when comparing the two messages against each other and against the no-message condition) were plotted against a set of possible prior distributions with different widths (Figure S3; note that we did not refit the model with these priors so the possible priors are only an illustration). The normal distribution was chosen as the possible prior distribution as it has a higher peak at 0 compared to the Cauchy and t distributions (which have fatter tails than the normal distribution). Three different Standard Deviations (SDs) were chosen to explore different possibilities: 0.05, 0.1 and 0.25.

A prior with SD of 0.05 means that with 95% probability the largest expected difference between conditions observed under the null model (i.e., under the model assuming no effect) is around 4% on the respective response scale. The reason for this is that 95% probability corresponds approximately to the area within -1.96 and 1.96 quantiles of the standard normal distribution (i.e., on the linear scale before applying the link function). Here, we are only interested in the magnitude of the difference and not the sign, so we can focus on the absolute value 1.96. Thus, with 95% probability all differences of a prior with SD of 0.05 are $1.96 * 0.05 = 0.098$ on the linear scale. Because the probit link function is non-linear, an absolute difference of 0.098 on the linear scale does not always correspond to the same difference on the response scale. What matters is where this difference occurs. However, because of the s-shape and symmetry of the link function, the largest possible difference on the probability scale happens at the midpoint where the linear scale is 0 and the probability scale is 0.5. Thus, the maximum difference on the probability scale for a difference Δ on the linear scale corresponds to $\Phi\left(\frac{\Delta}{2}\right) - \Phi\left(-\frac{\Delta}{2}\right)$, where $\Phi(\cdot)$ is the inverse of the probit link function (i.e., the cumulative distribution function of the normal distribution). In our case this results in

$\Phi\left(\frac{0.098}{2}\right) - \Phi\left(-\frac{0.098}{2}\right) \approx 0.04$. For the prior with SD = 0.1 the maximum difference

expected on the response scale is around 8% and for a prior with SD = 0.25 it is around 20%.

It should be noted that these priors are comparatively narrow. For example, (2) also use a model with a probit link function and define a narrow prior as a normal distribution with SD = 0.52 (see their Appendix E). This means that the width of their narrowest prior is more than double the width of our widest prior.

Figure S3. Comparison of density estimates of the difference distributions (in black) against three possible prior distributions with different prior widths (red: SD = 0.05; green: SD = 0.1; blue: SD = 0.25). Each row shows results for one set of difference distribution (e.g., first row shows difference between white and yellow label), and each column shows results for one hypothesis.

Having specified several possible priors, this allows us to judge evidence for or against the null hypothesis by looking at the ratio of the densities at zero, using the Savage-Dickey procedure (3-5).

This analysis shows that even with extremely tight priors (a normal distribution with $SD = 0.05$), there is evidence for no difference (i.e., black density estimate above the prior at 0, the grey vertical line) in all but one of the cases: the lower right plot (proportion bet, yellow message minus no-message). This interpretation is consistent with the results reported in the main text of no evidence for the effect of safer gambling messages apart from the backfire effect observed with increased betting proportion for the yellow message in comparison to the control condition. With somewhat more reasonable priors (SD of 0.1) there is relatively clear evidence for the null for both gamble at all and gamble everything when comparing message and no-message conditions. And with a wider prior width of $SD = 0.25$, which corresponds to a density of 1.60 at 0, the evidence for the null becomes quite high, as in each case the observed density for the null is at least 10 (i.e., the Bayes factor for the null is 6.25 or larger).

Exploratory analysis: Covariates

After the task, participants completed two metrics to assess their relationship with gambling: the Problem Gambling Severity Index (PGSI) and the Enhancement Motives subscale (ENH) of the Gambling Motives Questionnaire (GMQ). Those who score high on the PGSI are considered to be problem gamblers ($PGSI > 7$), and those who score high on the ENH gamble for ‘fun’ and to ‘feel good.’ Bayesian ANOVAs comparing the PGSI and the ENH provided evidence for no difference in scores between the three experimental conditions, with Bayes factors above 30 for the null hypothesis for both scales (Table S1).

Individual difference scale	Control: No message	Yellow Message	White Message	Bayes factor for the null
Problem Gambling Severity Index (PGSI)	3.03 (4.83)	2.70 (4.40)	2.66 (4.30)	47.4
Enhancement Motives subscale of the Gambling Motives Questionnaire (ENH)	6.01 (3.91)	5.86 (3.89)	6.27 (3.98)	32.3

Table S1. Bayesian ANOVAs for potential effect of manipulation on scale scores. Columns 2 to 4 show means and SDs (in parentheses). The last column shows the default Bayes factors for the null model, comparing the three means.

Therefore, the same model as the main confirmatory analysis was run, but this time including the PGSI and ENH as covariates. This approach allowed us to assess if the main findings would be replicated when adjusting for individual differences in the PGSI or ENH scores. This analysis was conducted twice, with and without interactions between the messages and the two indices. In the first analysis, only the main effects of both indices were included (PGSI and ENH, both mean-centred), in addition to the experimental conditions. The second analysis also included the two-way interaction between the indices and the messages (but not between the two indices themselves). Both models showed the same patterns for the influence of the safer gambling messages as reported in the main analysis, without the covariates (Table S2): there was no significant effect of the messages on gambling behaviour, with the exception of the backfire effect of the increased proportion bet with the yellow message (which can be seen on the table as “Message=Yellow” under “proportion bet”).

There was a positive relationship (in the no-message condition) with the ENH on the probability to gamble (0.15, 95% CI [0.10, 0.19]) and the mean proportion bet (0.02, [0.00, 0.04]), as well as positive relationships of the PGSI with the probability to gamble everything (0.07, [0.04, 0.09]). Additionally, there was evidence for one interaction in relation to the mean proportion bet: a positive interaction term for PGSI and the white message condition (0.04, [0.00, 0.08]) Therefore, when the white message was present, individuals with high PGSI gambled a larger proportion of their bonus compared to participants in the no-message condition— a result similar to the effect found with the yellow message in the main analysis, but amplified by the individual PGSI scores. There were no changes to the overall effects of

experimental conditions (no message vs. yellow message vs. white message) after adjusting for the effects of PGSI and ENH.

Proportion bet	Main model: No covariates		Model with covariates: Main effects only		Model with covariates: Main effects + interactions	
Predictors	Estimates (SD)	CI (95%)	Estimates (SD)	CI (95%)	Estimates (SD)	CI (95%)
Coefficients for “probability to gamble at all” (g)						
Intercept	1.96* (0.14)	[1.69, 2.23]	2.08* (0.14)	[1.80, 2.36]	2.16* (0.16)	[1.85, 2.49]
Message=White	-0.19 (0.19)	[-0.55, 0.18]	-0.23 (0.19)	[-0.60, 0.14]	-0.35 (0.21)	[-0.76, 0.05]
Message=Yellow	-0.17 (0.19)	[-0.53, 0.20]	-0.16 (0.19)	[-0.53, 0.21]	-0.17 (0.22)	[-0.61, 0.26]
PGSI			-0.01 (0.02)	[-0.05, 0.03]	-0.03 (0.03)	[-0.09, 0.03]
ENH			0.15* (0.02)	[0.10, 0.19]	0.19* (0.05)	[0.11, 0.29]
Message=White : PGSI					0.04 (0.05)	[-0.06, 0.13]
Message=Yellow : PGSI					0.04 (0.05)	[-0.06, 0.14]
Message=White : ENH					-0.10 (0.06)	[-0.22, 0.01]
Message=Yellow : ENH					-0.02 (0.06)	[-0.14, 0.10]
Coefficients for “conditional probability to gamble everything” (e)						
Intercept	-1.31* (0.12)	[-1.54, - 1.08]	-1.37* (0.12)	[-1.60, - 1.14]	-1.38* (0.12)	[-1.63, - 1.15]
Message=White	-0.07 (0.17)	[-0.40, 0.26]	-0.05 (0.17)	[-0.38, 0.29]	-0.03 (0.18)	[-0.37, 0.31]
Message=Yellow	0.05 (0.17)	[-0.27, 0.38]	0.07 (0.17)	[-0.25, 0.40]	0.08 (0.17)	[-0.25, 0.42]
PGSI			0.07* (0.01)	[0.04, 0.09]	0.05* (0.02)	[0.00, 0.09]
ENH			0.02 (0.02)	[-0.02, 0.05]	0.06 (0.03)	[-0.01, 0.12]
Message=White : PGSI					0.01 (0.04)	[-0.06, 0.08]
Message=Yellow : PGSI					0.05 (0.04)	[-0.02, 0.12]
Message=White : ENH					-0.04 (0.05)	[-0.13, 0.05]
Message=Yellow : ENH					-0.09 (0.05)	[-0.18, 0.00]
Coefficients for beta regression model (μ and ϕ), or “proportion bet”						
Intercept	-0.23* (0.05)	[-0.34, - 0.13]	-0.23* (0.05)	[-0.34, - 0.13]	-0.24* (0.05)	[-0.34, - 0.13]
Message=White	0.08 (0.07)	[-0.06, 0.23]	0.08 (0.07)	[-0.07, 0.23]	0.09 (0.08)	[-0.06, 0.24]
Message=Yellow	0.15* (0.07)	[0.00, 0.29]	0.15* (0.07)	[0.00, 0.29]	0.16* (0.08)	[0.01, 0.30]
PGSI			0.01 (0.01)	[-0.01, 0.02]	-0.01 (0.01)	[-0.04, 0.01]
ENH			0.02* (0.01)	[0.00, 0.04]	0.04* (0.01)	[0.01, 0.07]
Message=White : PGSI					0.04* (0.02)	[0.00, 0.08]
Message=Yellow : PGSI					0.03 (0.02)	[-0.00, 0.07]
Message=White : ENH					-0.02 (0.02)	[-0.06, 0.02]
Message=Yellow : ENH					-0.04 (0.02)	[-0.08, 0.00]
Phi (ϕ)	1.11* (0.04)	[1.04, 1.19]	1.12* (0.04)	[1.04, 1.20]	1.13* (0.04)	[1.05, 1.20]

Table S2. Regression coefficients of the ZOIBR models of proportion bet. Standard deviation of the posterior, the Bayesian equivalent of the standard error, in brackets. Estimates are given on the linear scale (i.e., without applying the link function). CI gives the 95% Bayesian credibility interval. PGSI and ENH are mean-centered. The table is split into three sub tables, each corresponding to the coefficients corresponding to one part of the ZOIBR model. The intercept corresponds to the mean of the no-message condition and all coefficients are in relation to this mean (i.e., the models use treatment coding). For example, “Message=white” is the difference from the no-message condition for the white message condition, and “PGSI” refers to the relationship between PGSI and the corresponding ZOIBR parameter in the no message condition. “:” indicates an interaction. (*) Credibility interval does not include zero, indicating evidence that the coefficient differs from 0.

Exploratory analysis: Riskiness

Different roulette bets are associated with different levels of risk, as indicated by their likelihood of being paid out and the magnitude of the payoff. For example, if £1 is bet on a single number, this bet will pay £36 with a probability of $1/37$. In comparison, if £1 is bet on red, this bet will pay £2 with a probability of $18/37$. This is because there are 37 numbers on the roulette wheel (from 0 to 36), of which 18 are red. These payoff figures, 36 and 2, are the decimal odds for these two bets, respectively, as they indicate the amount won per unit bet. For example, a £0.50 bet with decimal odds of 36 will pay £18 if it wins.

Additionally, gamblers can place multiple bets per spin, (e.g., betting £0.50 on red and £0.50 on 7). Because the number 7 is a red colour, the bet on 7 can only win if the first bet on red also wins. Multiple bets per spin can either be placed in a way that increases risk, as in this example, or in a way that reduces risk (for example, a bet on black added to the bet on 7, or betting on reds and blacks together). All roulette bets have the same expected value ($36/37$), however, a £1 bet on a single number is riskier than £1 bet in all 18 red numbers at once, because the former is more concentrated and pays a higher payoff but only if the winning number is picked, but it is less likely to occur, while the latter will pay a lower payoff, but the probability of a winning number being hit is higher.

To evaluate if the safer gambling messages influenced risk taking, a risk metric was calculated for each spin made by participants. The metric was calculated as the variance of the decimal odds from each individual number on the roulette wheel, according to the bets placed by the participant. The metric measured the concentration of the bet, with more risk represented by more concentrated bets (e.g., betting on a single number), and lower risk represented by more spread bets (e.g., betting on reds or betting on evens). A value of zero was assigned for the numbers in which participants did not bet (since they win zero times the amount bet if the roulette stops on those numbers). The risk variable was calculated to be the variance of the decimal odds for every number on the roulette table, taking into account the

bets the participant has placed, and including zeroes for non-winning numbers. The value of this risk variable can range between 0.03 and 35.03. Higher numbers indicate more risk taking (with the highest, 35.03, associated with concentrating the bet on a single number). Lower numbers indicate lower risk taking (with the lowest, 0.03, associated with spreading the bet across every number excluding the zero, for example by placing a token on odds and another same-value token on evens). The risk variable cannot be zero in the task because participants are not able to bet the same amount across all numbers including the roulette's zero due to the limitations in the values of the available tokens and total bet amounts.

The aim of this exploratory analysis was to test if the gambling messages influenced the level of risk taking. This was evaluated with a beta-regression model. Because beta-regressions are bounded between 0 and 1, the risk metric was first divided by 36 to create a variable ranging from just above zero to just below one. From the model-fitted estimates, it can be seen that the 95% CIs for the two messages both include 0, which means that there was no influence of either message manipulation on the risk-taking variable (Table S3).

Predictors	Estimates (SD)	CI (95%)
Intercept	-2.09* (0.05)	[-2.18, -1.99]
Message=White	-0.03 (0.06)	[-0.16, 0.10]
Message=Yellow	-0.08 (0.06)	[-0.21, 0.04]
Phi	1.77* (0.04)	[1.68, 1.85]

Table S3. Beta regression coefficients of the model on riskiness. Parameterisation of the model was the same as for the beta regression model in Table S2 (i.e., using treatment coding).

Experiment 3

Experimental design

The structure of Experiment 3 was very similar to that of Experiment 2, with the only difference being that participants were allowed to bet on a single round football matches

instead of potentially multiple rounds of roulette. Experiment 3 consisted of five individual pages. Figure S4 shows the flow of Experiment 3 along the pages, with the average time spent on each page. Screen captures of all pages are available online (<https://osf.io/fwunh/>).

A total of $N = 1,318$ participants started the task by reading the consent form in which they were informed that high performance on the captcha task would yield “a bonus that can then be used to bet on the upcoming round of Premier League football fixtures.” After giving consent, participants then proceeded to the captchas page, which followed the same procedure as Experiment 2. Participants ($N = 310$) who typed in six or fewer correct captchas were immediately moved to the final thank-you page which ended the task and paid £1.50. Participants ($N = 1,008$) who typed in seven or more correct captchas were awarded a bonus of £3.00, which they could use to bet on football matches, in the next phase of the experiment. Four participants were excluded from further analyses because they abandoned the task before proceeding to the gambling questionnaires, and therefore did not finish the task. One additional participant was excluded due to missing captchas data, resulting in a total sample of $N=1,003$ for analysis.

Figure S4. Flow of Experiment 3. Each box represents an individual page of the experiment. All participants had to complete the consent form and type in the captchas. Only participants who completed seven or more correct captchas proceeded into the football betting part of the experiment. Four participants who abandoned the task were excluded from the analyses due to missing questionnaire data. The average time spent in each page is shown in brackets.

The football betting phase started with instructions on how to place bets appearing as a pop-up over the football betting page. The full instructions are available online as a screen capture (<https://osf.io/fwunh/>). Participants were told that they could place as many bets as they wanted, as long as the total amount did not exceed £3.00, and that they did not have to place any bets if they did not want to. Gambling message presence was manipulated between participants, with the message overlaid in the top of the football betting page at all times. After reading the instructions, participants closed the pop-up and were presented with a list of available football fixtures on which they could place bets. There was an on-screen help button that recalled the instructions at any point if needed.

To avoid on-screen display and formatting problems, participants with small screens (smaller than 1000 pixels available for the browser in width) were not allowed to participate in the experiment, which should exclude most mobile phones and small tablets. Participants with larger screens but smaller browser windows were told to increase the size of the browser window above 1000 pixels in width before continuing.

Participants were not required to place any bets. Thirty-nine percent of participants continued past the football betting page without placing any bets, thereby keeping their entire £3 bonus. Participants could place as many individual bets as they wanted over many different fixtures, as long as the total bet amount did not exceed £3. Among participants who placed bets, the average number of bets placed was 2.2. The football betting page allowed participants to place bets across 10 different fixtures. These fixtures were the complete upcoming round of head-to-head matches of all the 20 teams in England's Premier League. The games were to be played on Saturday 17th, Sunday 18th and Monday 19th of October 2020, and data collection started on Tuesday 13th and ended on Saturday before the first game started. For each fixture, they could bet on three different types of bets: which team would win the match (e.g., Liverpool to win, betting on a draw was also possible), what would be the correct final score

of the match (e.g., 2-1), and who would be the first goal-scorer in the match (e.g., Thiago Silva to score first, own goals not included, betting on no goal-scorers was also possible). Odds were shown next to each individual bet as fractional odds, the most common in the UK (e.g., 2/1 is equal to a £2 profit if placing a £1 bet that wins, for a total payout of £3). To ensure ecological validity, the odds used were actual odds extracted from a leading high street bookmaker's website on the Tuesday before the match weekend, immediately before the start of data collection. A list of all the available bets and the odds used is available online on the OSF repository.

Upon selecting any individual bet, participants were asked for their bet amount, and the potential earnings for that bet were calculated using the odds and shown to the participant. The minimum bet amount was £0.01 and the maximum was £3, in increments of £0.01. After viewing the calculated potential earning for a bet, participants could then choose to save this bet into their bet slip, change it, or cancel it. Bets saved in their bet slips could be later reviewed and deleted. Once participants were happy with the bets in their bet slip, they had to click the continue button to proceed to the final part of the experiment, which was the same gambling questionnaire as in Experiment 2. Once confirming their bet slip, participants could no longer change their bets.

All participants were paid the fixed participation fee of £1.50 within 24 hours of finishing the task. Bonuses were calculated based on the actual results of the individual matches plus any amount that was not bet. All bonuses, including the amounts that were not bet, were paid after all the individual games had been played and results were known. Therefore, there was no time-discounting advantage to not betting or betting on early matches, as all bonuses were paid at the same time on October 20th. The average bonus was £2.20, and the range was [£0, £33].

Exploratory analysis: Bayes factors supplement

Figure S5. Comparison of density estimates of the posterior of the difference distribution between the two conditions (in black) against three possible prior distributions with different prior widths (red: SD = 0.05; green: SD = 0.1; blue: SD = 0.25) for each hypothesis.

As in Experiment 2, the confirmatory analysis did not include Bayesian priors due to the difficulty of specifying what a good prior width is, as default priors do not exist for the ZOIBR model employed here. However, the density estimates for each hypothesis of the confirmatory analysis (gamble at all, gamble everything, and proportion gambled, when comparing the message against the no-message condition) were plotted against a set of possible prior distributions with different widths (Figure S5). Three different Standard Deviations (SDs) were chosen to explore different possibilities: 0.05, 0.1 and 0.25.

Even with an extremely small prior with SD = 0.05, we still see evidence for no difference (i.e., black density estimate above the prior at 0) for all but “gamble everything” for which there is some evidence for a backfire effect. Note that a prior with SD of 0.05 means that we expect with 95% probability that the largest effect we observe is 2.50% on the response scale.

Exploratory analysis: Covariates

As in Experiment 2, after the task, participants completed two metrics to assess their relationship with gambling: the Problem Gambling Severity Index (PGSI) and the Enhancement Motives subscale (ENH) of the Gambling Motives Questionnaire (GMQ).

Bayesian ANOVAs provided evidence for the null of no difference between the two the two experimental conditions for the PGSI and no clear evidence (pointing towards a null effect) for or against a difference for the ENH (Table S3).

Individual difference scale	Control: No message	Yellow Message	Bayes factor for the null
Problem Gambling Severity Index (PGSI)	2.12 (3.84)	1.82 (3.45)	5.98
Enhancement Motives subscale of the Gambling Motives Questionnaire (ENH)	5.23 (3.66)	4.77 (3.73)	2.14

Table S4. Bayesian ANOVAs for potential effect of manipulation on scale scores. Columns 2 and 3 show means and SDs (in parentheses). The last column shows the default Bayes factors for the null model, comparing the two means.

We proceeded to run the same model as the main confirmatory analysis, but this time including the PGSI and ENH as covariates. This analysis was conducted twice, with and without interactions between the message and the two indices. Both models showed the same patterns for the influence of the message as reported in the main analysis, without the covariates (Table S4): There was no credible effect of the safer gambling message on gambling behavior.

Proportion bet	Main model: No covariates		Model with covariates: Main effects only		Model with covariates: Main effects + interactions	
Predictors	Estimates (SD)	CI (95%)	Estimates (SD)	CI (95%)	Estimates (SD)	CI (95%)
Coefficients for “probability to gamble at all” (g)						
Intercept	0.41* (0.09)	[0.23, 0.59]	0.41* (0.09)	[0.22, 0.59]	0.41* (0.09)	[0.23, 0.59]
Message=Yellow	0.04 (0.13)	[-0.21, 0.30]	0.09 (0.13)	[-0.17, 0.35]	0.10 (0.13)	[-0.16, 0.36]
PGSI			-0.04* (0.02)	[-0.08, - 0.00]	-0.05* (0.03)	[-0.10, - 0.00]
ENH			0.12* (0.02)	[0.08, 0.16]	0.12* (0.03)	[0.06, 0.17]
Message=Yellow : PGSI					0.03 (0.04)	[-0.05, 0.11]
Message=Yellow : ENH					0.01 (0.04)	[-0.07, 0.09]
Coefficients for “conditional probability to gamble everything” (e)						
Intercept	0.67* (0.12)	[0.43, 0.91]	0.64* (0.12)	[0.40, 0.88]	0.66* (0.12)	[0.41, 0.90]
Message=Yellow	0.28 (0.18)	[-0.07, 0.62]	0.30 (0.18)	[-0.04, 0.65]	0.29 (0.18)	[-0.06, 0.64]
PGSI			0.01 (0.03)	[-0.04, 0.06]	0.02 (0.04)	[-0.05, 0.64]
ENH			0.05* (0.03)	[0.00, 0.11]	0.03 (0.04)	[-0.05, 0.10]
Message=Yellow : PGSI					-0.02 (0.06)	[-0.13, 0.09]
Message=Yellow : ENH					0.05 (0.05)	[-0.05, 0.16]
Coefficients for beta regression model (μ and ϕ), or “proportion bet”						
Intercept	-0.39* (0.08)	[-0.53, - 0.24]	-0.41* (0.08)	[-0.56, - 0.26]	-0.40* (0.07)	[-0.55, - 0.26]
Message=Yellow	-0.02 (0.11)	[-0.24, 0.20]	0.02 (0.11)	[-0.20, 0.24]	0.04 (0.11)	[-0.18, 0.25]
PGSI			-0.02 (0.02)	[-0.05, 0.01]	-0.05* (0.02)	[-0.09, - 0.00]
ENH			0.04* (0.02)	[0.01, 0.08]	0.01 (0.02)	[-0.03, 0.05]
Message=Yellow : PGSI					0.03 (0.03)	[-0.03, 0.09]
Message=Yellow : ENH					0.09* (0.03)	[0.02, 0.15]
Phi (ϕ)	1.80* (0.10)	[1.60, 1.98]	1.83* (0.10)	[1.63, 2.02]	1.88* (0.10)	[1.68, 2.07]

Table S5. Regression coefficients of the ZOIBR models of proportion bet. Standard deviation of the posterior, the Bayesian equivalent of the standard error, in brackets. Estimates are given on the linear scale (i.e., without applying the link function). CI gives the 95% Bayesian credibility interval. PGSI and ENH are mean-centred. The table is split into three sub-tables, each corresponding to the coefficients corresponding to one part of the ZOIBR model. The intercept corresponds to the mean of the no-message condition, and all coefficients are in relation to this mean (i.e., the models use treatment coding). (*) Credibility interval does not include zero, indicating evidence that the coefficient differs from 0.

There was a positive relationship with the ENH (for the main effect, which corresponds to the no-message condition) on the probability to gamble at all (0.12, 95% CI [0.08, 0.16]), probability to gamble everything (0.05, 95% CI [0.00, 0.11]), and the mean proportion bet (0.04, [0.01, 0.08]). Furthermore, there was a credible interaction between the message condition and ENH for the mean proportion bet indicating that this relationship was attenuated (by 0.09 [0.02, 0.15]) for participants with the safer gambling message. Therefore, individuals with high ENH gambled more (across all three measurements of gambling

behaviour), and this effect appeared to be attenuated for the safer gambling message and the mean proportion bet.

Somewhat surprisingly, we found negative relationships of the PGSI with the probability to gamble at all (-0.04, [-0.08, -0.00]) as well as with the mean proportion bet (-0.05, [-0.09, -0.00]) for the model with interaction. Taken at face value, this negative relationship suggests that individuals with high PGSI scores seemed to be less likely to gamble at all or gamble less. Further analysis, however, suggests that this is not a justifiable inference when considering the results in more detail. In particular, it helps to only consider a model in which only PGSI is a covariate given the substantial correlation between both covariates ($r = .36, p < .001$).

In the model with only PGSI as covariate as a main effect, the relationship between PGSI and the probability to gamble at all is zero (0.00 [-0.03, 0.04]). This suggests that the negative relationship of PGSI with the probability to gamble at all reported in Table S5 is a consequence of the collinearity between the covariates. In other words, the negative relationship between PGSI and the probability to gamble occurs only after adjusting for ENH scores (given the positive relationship between both such results patterns are notoriously difficult to interpret).

In the model with only PGSI as covariate and interactions, we again find a negative main effect of PGSI with mean proportion bet (-0.04 [-0.09, -0.00]), but also a positive interaction between PGSI and message condition (0.06 [0.01, 0.12]). This indicates that the effect of PGSI in the no-message condition is negative (i.e., the main effect of PGSI, -0.04 [-0.09, -0.00]), but this effect is descriptively positive in the yellow message condition (0.02 [-0.02, 0.06]). Furthermore, the overall effect does not credibly differ from 0 (-0.01 [-0.04, 0.02]). Overall, the evidence for a negative relationship with PGSI with mean proportion bet only

emerges for the no-message condition. Given that this is a unique finding in the overall pattern of results we refrain from discussing this further and would want to see a replication before putting too much weight into it.

Most importantly, there were no changes to the overall effects of experimental conditions (no message vs. yellow message) after adjusting for the effects of PGSI and ENH.

Exploratory analysis: Riskiness

Different football bets are associated with different levels of risk, as indicated by their odds. For example, a bet with a 150/1 fractional odds is riskier than a bet with a 2/1 fractional odds. The former gives a £150 profit on a £1 bet, while the latter gives a £2 profit. It is expected that bets with higher odds are much less likely to occur than bets with lower odds.

Though we presented participants with fractional odds, which are more common in the UK and participants would be better accustomed to, for analysis we used decimal odds. Fractional odds provide the profits made on a bet, excluding the amount staked, while decimal odds provide the total payout, including the stake. Furthermore, fractional odds are usually represented as fractions (150/1, 2/5) while decimal odds are usually represented as decimals (151 and 1.4 for the same two bets, respectively).

Fractional Odds = (Payout – Stake) / Stake

Decimal Odds = Payout / Stake

Decimal Odds = Fractional Odds + 1

Amongst the potential bets available to participants, the riskier bets had fractional odds of 150/1 (decimal odds = 151) and the safest bets had fractional odds of 2/5 (decimal odds = 1.4). The maximum potential payout on a single bet was £450, although no one placed such a

bet. In fact, no one placed any bets at 150/1 odds. The riskiest bets taken were at 125/1, and the safest bets taken were at 2/5. The highest potential payout of any single bet which was actually placed was £273, although this bet was not won. The largest individual bet won was £33.

To evaluate if the gambling messages influenced risk taking, a risk metric was calculated for the bets made by participants. This analysis was only conducted for participants who placed bets. First, we started by converting all odds into decimal odds by adding 1 to them. Then we calculated the average non-weighted decimal odds for each participant across all the individual bets they placed. Finally, we calculated relative risk, taking into account the maximum and minimum risk taken by participants. Because the riskiest bet taken by any participant has decimal odds of 126 and the lowest was 1.4, relative risk was calculated taking this range into account, with the formula

$$\text{Relative_Risk} = (\text{Avg_odds} - 1.4) / (126 - 1.4)$$

The metric could therefore vary between 0 and 1. In order to analyse this data with a beta-regression (because the model does not allow for values at the extremes), two participants who had a relative risk of zero had to be excluded, and no participants had a relative risk of one. From the model-fitted estimates, it can be seen that the 95% CIs for the gambling message includes 0, which means that there is no influence of the message manipulation on the risk-taking variable (Table S6).

Predictors	Estimates (SD)	CI (95%)
Intercept	-3.07* (0.07)	[-3.21, -2.93]
Message=Yellow	-0.10 (0.08)	[-0.26, 0.05]
Phi	2.31* (0.07)	[2.16, 2.44]

Table S6. Beta regression coefficients of the model on riskiness. Parameterisation of the model is the same as for the beta regression model in Table S2 (i.e., using treatment coding).

Including the two participants whose relative risk was zero, we re-ran the analysis using a zero-inflated beta-regression. The results were the same, with no evidence that the safer gambling message affected the riskiness of the bets.

References

- (1) Yang H. A captcha library that generates audio and image CAPTCHAs. 2014; Available at: <https://github.com/lepture/captcha>, 2020.
- (2) Gronau QF, Wagenmakers E, Heck DW, Matzke D. A simple method for comparing complex models: Bayesian model comparison for hierarchical multinomial processing tree models using Warp-III bridge sampling. *Psychometrika* 2019;84(1):261-284.
- (3) Penny WD, Ridgway GR. Efficient posterior probability mapping using Savage-Dickey ratios. *PLoS One* 2013;8(3):e59655.
- (4) Dickey JM. The weighted likelihood ratio, linear hypotheses on normal location parameters. *The Annals of Mathematical Statistics* 1971;42:204-223.
- (5) Kass RE, Raftery AE. Bayes factors. *Journal of the american statistical association* 1995;90(430):773-795.