



Predicting Limit-Setting Behavior of Gamblers Using Machine Learning Algorithms: A Real-World Study of Norwegian Gamblers Using Account Data

Michael Auer¹ · Mark D. Griffiths² 

Published online: 02 December 2019

© The Author(s) 2019

Abstract

Player protection and harm minimization have become increasingly important in the gambling industry along with the promotion of responsible gambling (RG). Among the most widespread RG tools that gaming operators provide are limit-setting tools that help players limit the amount of time and/or money they spend gambling. Research suggests that limit-setting significantly reduces the amount of money that players spend. If limit-setting is to be encouraged as a way of facilitating responsible gambling, it is important to know what variables are important in getting individuals to set and change limits in the first place. In the present study, 33 variables assessing the player behavior among *Norsk Tipping* clientele ($N = 70,789$) from January to March 2017 were computed. The 33 variables which reflect the players' behavior were then used to predict the likelihood of gamblers changing their monetary limit between April and June 2017. The 70,789 players were randomly split into a training dataset of 56,532 and an evaluation set of 14,157 players (corresponding to an 80/20 split). The results demonstrated that it is possible to predict future limit-setting based on player behavior. The random forest algorithm appeared to predict limit-changing behavior much better than the other algorithms. However, on the independent test data, the random forest algorithm's accuracy dropped significantly. The best performance on the test data along with a small decrease in accuracy in comparison to the training data was delivered by the gradient boost machine learning algorithm. The most important variables predicting future limit-setting using the gradient boost machine algorithm were players receiving feedback that they had reached 80% of their personal monthly global loss limit, personal monthly loss limit, the amount bet, theoretical loss, and whether the players had increased their limits in the past. With the help of predictive analytics, players with a high likelihood of changing their limits can be proactively approached.

Keywords Gambling · Problem gambling · Responsible gambling tools · Limit-setting · Gambling algorithms

For most individuals, gambling is an enjoyable socially acceptable form of leisure activity (Stucki and Rihs-Middel 2007). However, research has consistently demonstrated that a small minority of individuals experience severe negative consequences as a result of excessive

gambling (Meyer et al. 2009) and that problem gambling has become an important public health concern (Shaffer and Korn 2002; Williams et al. 2012). Despite the many different methods of assessment, a recent comprehensive review of 69 gambling prevalence surveys reported that the lifetime prevalence of problem and pathological gambling worldwide ranged from 0.7% (in Denmark) to 6.5% (in Estonia) Calado & Griffiths, 2016.

There have been a number of national gambling prevalence surveys in Norway (where the present study was carried out; Bakken et al. 2009; Götestam and Johansson 2003; Jonsson 2006; Lund and Nordlund 2003). The most recent nationally representative study by Pallesen et al. (2016) reported that 0.9% of the Norwegian population were problem gamblers ($N = 5485$ participants) using the Problem Gambling Severity Index (Ferris & Wynne, 2001).

Limit-Setting Tools in Gambling Environments

Player protection and harm minimization have become increasingly important in the gambling industry and have led to many gambling operators providing responsible gambling tools to their clientele (Griffiths 2012; Harris and Griffiths 2017). One of the most popular responsible gambling applications is the provision of pre-commitment limit-setting tools where players can pre-determine the amount of money and/or time they spend gambling within a specified time period (typically per day, per week, and/or per calendar month) (Wood and Griffiths 2010). Research has indicated that such limit-setting tools are perceived favorably by players (Griffiths et al. 2009). Most operators who utilize limit-setting tools offer one or more of the following: (i) *deposit limits* (i.e., the maximum amount of money that a player can deposit into their account at any one time); (ii) *bet limits* (i.e., the maximum amount of money that a player can gamble with on a single or concurrent games); (iii) *loss limits* (i.e., the maximum amount of money that a player can lose in any single session or sessions); and (iv) *play limits* (i.e., the maximum amount of time that a player can gamble for at any given time) (Wood and Griffiths 2010). In addition to this, Walker et al. (2015) proposed limits on the amount of money a player can win (so-called win limits). Using a simulated slot machine, they tested “win limit” features and reported that self-enforced win limits resulted in players gambling more responsibly.

The setting of limits by players can be either mandatory or voluntary (Wood and Griffiths 2010) whereby players are forced to set limits if they want to access the games on specific gambling websites (i.e., mandatory) or players can make their own choice as to whether to use limit-setting tools on offer (i.e., voluntary). Some gambling operators have introduced global loss limits. For instance, in 2016, the Norwegian government’s gaming provider *Norsk Tipping* introduced a monthly global loss limit of NOK 20,000 (approximately €2100 or \$2450 [US]). Attitudes by players towards this global loss limit were recently evaluated by Auer et al. (2018a). They surveyed 2352 *Norsk Tipping* customers and found that the majority of players (79%) found the mandatory spending limit both useful and helpful. However, Auer et al. also noted that approximately one third of high-risk gamblers were less favorable towards the global loss limit. Some of the gamblers in this high-risk group may have felt that the global loss limit impeded their usual playing activity in some way. However, two thirds of the high-risk gamblers still held positive views about the mandatory loss limit. This in contrast to an earlier Canadian study by Bernhard et al. (2006) who reported that gamblers in their study were strongly opposed to mandatory limits.

A few studies have examined the extent to which online gambling operators provide limit-setting tools on their gambling website. Smeaton and Griffiths (2004) evaluated social responsibility practices of 30 British online gaming companies and reported that the minimum bet size among the 30 companies was £1, whereas the maximum bet size (of those gaming operators that set upper limits) was £20,000. They also reported that many of the gambling websites they visited typically had £10–£25 minimum bets and £250–£1000 maximum bets. However, this study was carried out at a time when issues such as player protection and harm minimization were only just emerging. A similar type of study by Kazhaal et al. (2011) visited 74 online poker sites and reported that less than half offered any limit-setting tools.

In the past couple of years, four studies have described and/or evaluated RG practices of online gambling websites. Bonello and Griffiths (2017) visited 50 of the world's most well-known online gambling sites and reported that 45 of them (90%) offered gamblers monetary limit-setting tools (with money spending limits and deposit limits being the most common types). In a French study, Marionneau and Järvinen-Tassopoulos (2017) reviewed consumer protection practices among all 18 licensed online operators in France. They reported that betting limits as well as deposit limits were offered by all the gaming operators visited. In an Italian study, Calvosa (2017) visited ten Italian regulated online gambling websites and all of them had a mandatory requirement for gamblers to choose a deposit limit prior to gambling. These latter studies were carried out in countries where limit-setting is mandatory explaining why limit-setting options were found in all gambling websites visited.

Finally, Cooney et al. (2018) described the RG practices of 39 Irish online gambling websites using the exact same protocol of Bonello and Griffiths (2017). Of these 39 gambling websites, 35 of them offered at least one limit-setting option. More specifically, 33 sites offered limit-setting for financial deposits, 14 offered limit-setting for financial losses, 12 offered limit-setting for financial spending, and seven offered limit-setting for spending time (i.e., within-session time limits). The study also reported that of the 39 Irish gambling websites, 12 offered one limit-setting option, seven offered two limit-setting options, nine offered three limit-setting options, and two offered four limit-setting options.

Other studies have examined limit-setting in different ways. In an experimental study, Wohl et al. (2013) reported that participants who watched an animated video prior to a gambling session more often stayed within their preset limits than players who did not watch the video. In another experimental study, Stewart and Wohl (2013) investigated the effect of a pop-up reminder concerning monetary limits. They found that individuals were significantly more likely to stick to their limits while gambling if they received a pop-up reminder which informed them that they reached their preset spending limit compared to those that did not. In a virtual reality casino study, Kim et al. (2014) found that participants who were explicitly asked to consider setting a time limit on their EGM play were significantly more likely to do so and spent less time gambling than those who were not given such instructions. Laboratory studies have shown that the majority of players adhere to preset limits. These studies have also shown that the design of limit-setting tools influences their effectiveness. However, most experimental studies have small sample sizes, typically include student samples, and comprise gambling simulations in non-ecologically valid settings (i.e., laboratories rather than real gambling venues).

A few studies have investigated player behavior following the setting of monetary limits. For instance, among Nova Scotian video lottery players, a Canadian study found that RG features (including limit-setting tools) generally reduced the overall levels of player expenditure (Focal Research 2007). However, as Wood and Griffiths (2010) noted, the specific impact

of monetary limit-setting was unknown because these features were not separated out from the other RG tools provided.

The remaining studies on limit-setting tools in gambling environments have been carried out using behavioral tracking account data provided by gambling operators. Using data provided by *bwin Interactive Entertainment*, Broda et al. (2008) investigated the effects of player deposit limits among 47,000 sports bettors over a 2-year period. They compared the gambling behavior of those who tried to exceed their deposit limit to players that did not. At the time of the study, it was mandatory for those gambling with *bwin* to set a deposit limit, and gamblers could not set a deposit limit of more than €1000 a day (or €5000 a month). Gamblers could also set their own deposit limits below the mandatory requirement. Broda et al. (2008) reported that only 0.3% of the gamblers attempted to exceed their deposit limit. In fact, most players (95%) never deposited more than €1050 per month (i.e., approximately one fifth of the monthly maximum €5000). Given the large daily and mandatory limits, it is perhaps unsurprising that so few gamblers reached their financial limits (Wood and Griffiths 2010).

Only one study has ever evaluated the efficacy of voluntary limit-setting using a real-world dataset comprising actual gamblers. Auer and Griffiths (2013) were given access to a random sample of 100,000 online gamblers on the *win2day* gambling website. Their final study sample comprised 5000 online gamblers who set their own limits on the *win2day* website (where deposits were limited to €800 per week). The results demonstrated that voluntary limit-setting had a statistically significant effect on the highest-intensity gamblers (i.e., they significantly decreased their gambling compared to similar players who did not choose a limit). The authors concluded that voluntary limit-setting had the desired effect that the most gaming-intense players appeared to get the most benefit. The results also showed that different players got different benefits based on the types of limit they set. More specifically, results demonstrated that casino and lottery players got the most benefit from the setting of money limits whereas poker players got the most benefit from the setting of time limits.

Auer et al. (2018b) investigated the effect of a loss-limit reminder on subsequent gambling. Players received pop-up messages when they had reached 80% of their personal monthly loss limit. Using a matched-pairs design, the study showed that among the players overall, receiving the message significantly reduced their gambling expenditure 3 months later. Ivanova et al. (2019) compared the gambling intensity between users of an online gambling service prompted to set a deposit limit and non-prompted customers. They found that prompting online gamblers to set a voluntary deposit limit of optional size did not affect subsequent net loss compared to unprompted customers, who set deposit limits without being prompted. They concluded that setting a deposit limit without a prompt or increasing/removing a previously set limit may be a marker of gambling problems and may be used to identify customers in need of help.

Although the empirical evidence is sparse, research suggests that gamblers find limit-setting a useful RG tool and the study by Auer and Griffiths (2013) strongly suggests that limit-setting works for gamblers who most need it (and that players who use limit-setting tools significantly decrease the amount of money they spend gambling on the sites where they use the tools). Consequently, if limit-setting is to be encouraged as a way of facilitating responsible gambling, it is important to know what variables are important in getting individuals to set limits in the first place and to be potentially able to predict which players could be encouraged to change their limits. The present authors of this study are not aware of any other study that has attempted to predict the use of limit-setting tools in a real-world environment.

However, supervised techniques have been applied to predict problem gambling. For instance, Percy et al. (2016) predicted self-exclusion with a sample of 845 online gamblers. They compared different statistical methods and found that the random forest technique performed best (see “Methods” section for a description of this technique and others). Philander (2014) compared different statistical techniques for the prediction of disordered internet sports gamblers. Again, the random forest technique was the most accurate technique. Other techniques used were support vector machines, neural networks, linear models, and logistic regression. Schellinck and Schrans (2011) tested supervised learning algorithms to predict the risk of problem gambling. However, little detail was given concerning the techniques and variables used. Machine learning is a popular technique in marketing and customer relationship management (CRM). Coussemant and De Bock (2013) have used machine learning to model customer retention and found that ensemble techniques (which mix different statistical models) are more accurate than single models.

The present study’s goal was to understand player behavior that is predictive of limit-setting. If it is possible, operators could monitor the behavioral indicators and then actively prompt players who might be likely to set limits. This could then motivate more players to change limits and help nudge players who are close to altering their limits but who do not do it automatically. There are no hypotheses regarding the behavior prior to limit-setting because this is the first study of this kind and by definition exploratory. Consequently, the study attempted to assess gambling behavior with as many variables as possible.

Methods

Participants: The authors were given access to a 20% random sample of all active *Norsk Tipping* players from January to June 2017. From this sample, 70,789 players had played at least one game during January to March 2017 and had a personal global monthly loss limit. The inclusion criterion was thus that players had to have played at least one game across any channel (online casino, VLT [video lottery terminal] sports betting, lottery, etc.) during January to March 2017. A maximum monthly loss limit of NOK 20,000 was introduced in September 2016. The average age of the 70,789 participants was 41.04 years (SD = 13.38), and the percentage of females was 28%. Out of the 70,789 players, 6.7% had changed their personal global monthly loss limit at least once during April to June 2017. Players could increase or decrease their limit during April to June 2017, and 6.3% increased their limit and 0.7% decreased their limit. In the present study, 33 variables assessing the players’ behavior from January to March 2017 were computed (see [Appendix](#)). All variables were computed from January to March 2017. Among these variables were the players’ bet and loss in total and for the different game types (online casino, lottery, VLT, sports betting), age and gender, the players’ personal global monthly loss limit, and whether players received feedback that they had reached 80% of their personal global loss limit. The number of different playing days was also computed. The authors also used information computed by *PlayScan*, a player-tracking tool that identifies potentially problematic gambling. The 33 variables which reflect the players’ behavior from January to March 2017 were then used to predict the likelihood of changing the limit during April to June 2017. The 70,789 players were randomly split into a training dataset of 56,532 and an evaluation set of 14,157 players. This corresponds to an 80%/20% split.

It is important to test mathematical models which have been derived from a set of training data on independent test data. This is especially true of machine learning models such as

random forest which are prone to overfitting because they can potentially use many parameters to describe the data. However, the patterns discovered in training data might not be generalizable and for that reason the reliability of a mathematical model should be tested on independent data which have not been used for the statistical modeling. The open source statistical software *R* was used to perform preparation, statistical analysis, and visualization. The *caret* package was used for machine learning. *Caret* is a package which is also free to download under the GPL V2 open source license. The goal of the present study was to predict a binary variable (limit change yes/no). Various statistical methods can be used. In the present study the following algorithms were used:

1. *Logistic regression*: This is the most classical among the chosen methods and quantifies linear relationships between the independent variables and a binary dependent variable. This method was first described by Cox (1958) and is still frequently used today. Logistic regression builds on linear regression and manages to apply a derived mathematical procedure to a binary outcome variable.
2. *Linear discriminant analysis (LDA)*: LDA (Riffenburgh and Clunies-Ross 1960) quantifies linear relationships between a multinomial dependent variable and a number of independent variables. LDA builds a classifier for each category which can then be used to predict the outcome on future data. In case of binary data, LDA is mathematically very similar to logistic regression.
3. *Random forest (RF)*: RF (Liaw and Wiener 2002) is a popular machine learning method which fits the data with numerous decision trees which are then aggregated into a final model. RF can detect linear as well as non-linear patterns.
4. *Gradient boost machine learning (GBML)*: GBML (Friedman 2001) is a method which fits the data with numerous models which are then aggregated to a final model. GBML can detect linear as well as non-linear patterns.
5. *Naïve Bayes*: Naïve Bayes (John and Langley 1995) is a family of probabilistic classifiers based on Bayes' theorem. It requires the explanatory variables to be independent and the final model is a linear combination.

Except for logistic regression, all models were trained with a fourfold cross-validation. The dataset is randomly split up into four subsets. Removing one of the sets, 75% of the data is used to train the model which is then tested on the remaining 25%. The models' performance was visually evaluated via their respective receiver operating characteristic (ROC) curves (Hanley and McNeil 1982) and numerically via the *area under the curve* (Bradley 1997). The present authors were aware that many of the parameter variables were highly correlated. For instance, loss is total bet minus total won and thus correlated with both total bet and total won. However, a player who bets a lot of money does not necessarily lose a lot of money because chance is a major factor determining a players' loss. For that reason, all variables were submitted to the five machine learning models. The authors also chose a stepwise logistic regression which sequentially selects variables that increase the model accuracy (Lemeshow and Hosmer Jr 1982). This procedure does not select a variable which is highly correlated with another variable that is already in the model because it would not increase the model accuracy. Decision trees—which are the basis for random forest and gradient boost machine—are not affected by multicollinearity among the predictor variables because the final model only contains variables which are in the final model.

Results

Table 1 displays descriptive statistics of the training dataset. From January to March 2017, the 56,632 players on average bet NOK 11,638 and lost NOK 2399. On average they played on 16 days and the average personal monthly loss limit was NOK 5100. A total of 10% of the players had chosen the maximum possible monthly loss limit (i.e., NOK 20,000), and 5% of the players were classified as being problem gamblers according to the behavioral tracking tool *PlayScan* (Forsström et al. 2016). *Norsk Tipping* informs players when they have reached 80% of their chosen monthly loss limit either via SMS or an email, depending on the players' mode of play. A total of 14% of the players received such feedback at least once during January to March 2017, and 6.7% of the players changed their limit at least once during April to June 2017. The latter behavior represents the target variable which the present study wanted to predict.

Figure 1 displays the receiver operating curves (ROC) for the various statistical models. Every statistical method results in a score between "0" and "1" and is assigned to each of the training records (in this case players). However, the present analysis only distinguished between two conditions (changing loss limit or not changing loss limit). In order to classify each player based on the score, a cutoff value has to be determined. If one is very liberal and decides that every player who has a score greater than 0.00001 is predicted to have changed the limit, it will correctly classify all the players who actually changed their limit. This means that in Table 2 the percentage of correctly classified players will be "1." This corresponds to the *true-positive* rate. On the other hand, a very liberal cutoff value will also classify all the players who did not change their limits into the group of players who changed their limit. This means that the *false-positive* rate will be "1" as all the players who did not change their limit will be classified wrongly.

If a very conservative cutoff value is chosen somewhere close to "1," very few or no players will be predicted to have changed the limit. This means the *false-negative* rate will be very high because all the players who changed their limit will be classified as players who did not change the limit. On the other all of the players who did not change the limit will be correctly classified. This means that the *true-negative* rate will be close to "1." In medicine, when correct classification is extremely important and, on the other hand, diagnostic tests are very expensive, the correct cutoff value is of paramount importance. All the potentially ill people should be correctly classified, but few healthy people would also receive a positive diagnosis.

Table 1 Descriptive statistics of the training dataset ($N = 56,632$)

	Mean	SD
Amount bet	11,638	54,286
Amount lost	2399	31,812
Number of days	16	17
Global monthly limit	5100	5909
Maximum monthly limit	10%	30%
<i>PlayScan</i> red	5%	22%
Age	41	13
Female	28%	45%
Limit feedback	14%	35%
Limit changed	6.7%	25.1%

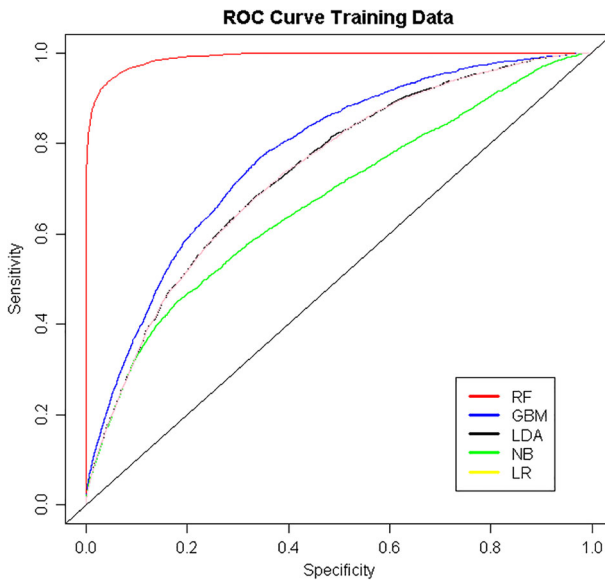


Fig. 1 Receiver operating curve of the training data ($N = 56,632$)

In the current case, the gambling operator would have to make two decisions: (i) the number of true limit changers it wants to detect and (ii) the number of players who do not plan to change their limit it would be willing to wrongly classify as limit changers.

In order to visualize the performance of a statistical model, the classification accuracy can be visualized via plotting the true-negative rate (*specificity*) against the true-positive rate (*sensitivity*). Figure 1 displays the receiver operating curve for the different statistical models across all possible cutoff values. The diagonal line represents the *worst* model, namely, the classification accuracy if players are classified randomly into limit changers and non-limit changers. The closer the line to the upper left corner, the better the predictive model. However, before those decisions are made, it is important to evaluate the performance of the models across the whole data spectrum.

The results demonstrate that the random forest algorithm appears to predict limit changing behavior much better than the other algorithms. With a very low true-negative rate (i.e., high percentage of correctly classified non-limit changers), it also achieves a very high true-positive rate (i.e., high percentage of correctly classified limit changers). The next best model was the gradient boost machine, followed by the linear discriminant analysis, logistic regression, and the naïve Bayes classifier. In order to quantify the predictive power, the area under the curve can be calculated. The minimum value is 50% and the maximum value is 100%. Table 3 reports that the random forest had the highest value, followed by the gradient boost machine.

Table 2 Classification matrix

True Condition		Positive	Negative
Predicted condition	Positive	TP	FP
	Negative	FN	TN

Safari et al. (2016) defined a value between 70 and 80% as fair, and a value lower than 70% as poor.

Machine learning algorithms such as random forest and gradient boost machine potentially use many parameters and are thus prone to overfitting. This means that they explain the training data very well but fail on independent test data. Figure 2 displays the ROC (Receiver Operator Curve) curve when the derived models are applied to the independent test data. It is clearly visible that the high predictive power of the random forest model collapses to almost the worst model when applied to the test data. This means that the random forest model was overfitting the training data. The other models' predictive power did not shrink as much as the random forest's. power. The ranking between the models, except for random forest, remains the same. Ignoring random forest, naïve Bayes is the worst model for the training as well as test data. Table 3 reports the area under the curve for the training and test data. The gradient boosting machine demonstrated the best performance on the test data, and the area under the curve value only dropped from 78 to 76%.

In order to interpret the results, Table 4 lists the most important variables for each method. The list is similar across the five methods. For instance, gradient boost machine contains the variable which measures the feedback players receive when they reach 80% of their personal monthly global loss limit, the personal monthly limit, the amount bet, theoretical loss, and whether the players increased their limit in the past. In order to get a better understanding, the present authors chose to additionally apply the *ctree* decision tree algorithm from the *party* package (Hothorn et al. 2006). Figure 3 displays the most important nodes of the tree. A decision tree hierarchically splits up the dataset according to those variables which separate best between the classification groups (changing or not changing limits).

According to this algorithm, the most important variable is whether players received feedback that they had reached 80% of their personal global monthly loss limit. Approximately one fifth of the players who received this feedback during January to March 2017 (18%) changed their limit during April to June 2017. Overall, 6.7% of the players changed their limit. Another group with a higher percentage of limit changes were the ones who had chosen a very low limit. One tenth of the players who had a limit lower than NOK 900 changed their limit (11%). Approximately one seventh of the players whose limit was greater than NOK 900 but lower than the maximum limit (14%) played on more than 14 days and were classified as problem gamblers by *PlayScan* changed their limit. The latter was the smallest group with only 836 players out of 56,632. In these types of analyses, researchers often use the results of a decision tree, namely the node assigned to the players, as an input for a logistic regression model. In contrast to a logistic regression, a decision tree can detect non-linear relationships and this combined procedure can enhance the performance of a logistic regression. Table 3

Table 3 Area under the curve of the respective ROC curves for the training and test data ($N = 56,632$)

	Training data	Test data
Random forest	99%	73%
Logistic regression	74%	73%
Linear discriminant	74%	73%
Naive Bayes	67%	64%
Gradient boosting machine	78%	76%
Tree + logistic regression	76%	75%




Fig. 2 Receiver operating curve of the test data ($N = 14,157$)

reports the area under the curve for this hybrid approach, and it is only 1% lower than the gradient boost machine's area under the curve value.

Discussion

The present study is the first ever that has attempted to predict limit changing using account data variables in a sample of 70,789 players who played with the Norwegian gaming operator *Norsk Tipping* from January to June 2017. The results demonstrated that it is possible to predict future limit-setting based on player behavior and that the gradient boost machine algorithm appeared to predict limit-changing behavior much better than the other algorithms. Limit-setting is a pre-commitment strategy which has previously been shown to help players reduce their monetary losses from gambling (i.e., Auer and Griffiths 2013) and for that reason it could be potentially helpful for operators to try to predict which players who gamble on their platform are more likely to utilize their limit-setting tools more often.

Table 4 Most important variables per method

RF	GBM	LDA	NB	LR
Limit	Feedback	Feedback	Feedback	Theoretical loss
Theoretical loss	Limit	Limit increases	Limit increases	Number of days
Loss	Amount bet	Limit increases	Limit	Amount bet VLT
Amount bet	Theoretical loss	Limit decreases	Limit decreases	Amount bet casino
Amount won	Limit increases	Number of days	Number of days	Casino playing

RF random forest, *GBM* gradient boost machine, *LDA* linear discriminant analysis, *NB* naïve Bayes, *LR* logistic regression




Fig. 3 Excerpt of a decision tree algorithm on the training data ($N = 56,632$). Notes: Limit, the chosen loss limit between January and March 2017; Number of playing days, number of active days between January and March 2017; red, *Playscan* status high risk between January and March 2017; Feedback, whether players received feedback that they had reached 80% of their loss limit between January and March 2017

For large gambling operators like *Norsk Tipping* with hundreds of thousands of active customers, it is useful to be able to identify customers who are about to display a significant behavioral change or who are in need of specific communication. With the help of predictive analytics, customers with a high likelihood of changing their limits could be proactively approached. In the current dataset, the vast majority of players increased their limits and only 0.7% of those who changed their limit decreased the personal loss limit. An increased limit is very likely to be succeeded by increased losses. The operator could therefore use the prediction model to intervene and persuade players to keep their existing limit instead of increasing it. Alternatively, without predictive analytical support, the gambling operator would have to make guesses about appropriate actions to be taken with their clientele.

The most important variables predicting future limit-setting in the gradient boost machine algorithm were players receiving feedback that they had reached 80% of their personal monthly global loss limit, personal monthly limit, the amount bet, theoretical loss, and whether the players had increased their limits in the past. Given a specific RG initiative like limit-setting, ultimately, a predictive analytics approach will help to identify the right message for the right customers at the right point of time. The findings here came from one gambling operator only (and from Norwegian citizens only), and it is unclear whether the same results would be replicated among other operators and gamblers from other nationalities. Therefore, other operators should be encouraged to utilize predictive analytics on their own data and see if the variables found as important in the present study are as important among other groups of gamblers.

In the marketing of non-gambling products, the prediction of future behavior is commonplace. For example, in the telecommunication industry, customer data is used to try and prevent customers canceling their subscription using predictive marketing techniques (Óskarsdóttir et al. 2017). In order to predict future limit-changing behavior in the present study, the research team applied five algorithms. Among those used was the random forest

technique. This algorithm performed best in two previous gambling-related studies that attempted to predict voluntary self-exclusion (i.e., Percy et al. 2016; Philander 2014). Logistic regression was chosen because it is a classical and well-proven statistical method. Furthermore, in a meta-study of clinical prediction modeling, Jie et al. (2019) found no evidence of superior performance of machine learning algorithms over logistic regression.

Similar to Philander (2014), the dataset in the present study ($N = 70,789$ players) was split into a training set and a test set. Machine learning algorithms which use many behavioral and demographic parameters tend to explain the training data very well, but can fail to explain test data (Lever et al. 2016). For large datasets like the one in the present study, a split into training and test data is commonplace whereas small datasets can be modeled via cross-validation. In the present study, random forest technique provided the best explanation of the training data with the highest accuracy (as shown in Fig. 1). This was followed by gradient boost machine learning across the whole spectrum of *sensitivity* (i.e., correct classification of players who changed limits) and *specificity* (i.e., correct classification of players who did not change limits).

However, on the independent test data, the random forest algorithm's accuracy dropped significantly which means that the algorithm used too many parameters and therefore explained patterns which were specific to the training data. The present authors tried different configurations from the standard configuration in the *caret* package. However, none of these managed to keep the accuracy on the test data. Philander (2014) also reported significant overfitting for this algorithm. The random forest technique comprises a series of trees, and a tree algorithm's overfitting depends on the minimum size of the number of records allowed in a leaf. If the minimum size is 1, then basically every record can end up in one specific leaf based on its specific data. However, it is very unlikely that the test data's records are exactly the same. Pattern recognition is all about identifying generalizable patterns which apply not only to the parameters used to learn the patterns. As measured via the area under the curve (AUC) in Table 3, the best performance on the test data was delivered by the gradient boost machine learning algorithm. Compared to the training data, it drops from 78 to 76%. Safari et al. (2016) defined a value between 0.7 and 0.8 as fair in terms of accuracy.

It is also imperative to gain a deeper understanding into the reasons for limit changing by gamblers. Machine learning algorithms like random forest or gradient boost machine are more complex than logistic regression which reports one odds ratio for each variable. The variable importance (*varImp*) function is part of the R *caret* library and analyzes the importance of each variable in a statistical model. Gradient boost machine learning chose whether players received feedback that they had reached 80% of their loss limit, the personal loss limit, the amount bet, the theoretical loss, and whether players had increased their monetary limit in the past. However, it is difficult to comprehend how machine learning algorithms explain a target variable and for that reason the authors chose a simple classification tree to gain insight into the behavior that precedes limit changing.

However, the mathematical model chosen by (for instance) gradient boost machine learning can be entirely different than the one depicted in Fig. 3. This figure shows the most important branches and leaves of a classification tree used to explain the 56,632 training data records. In the whole sample, 6.7% of players voluntarily changed their monetary limit (and the vast majority of these increased their limit). Among players who received feedback that they had reached 80% of their loss limit, 18% changed their limit. Players in the dataset who received feedback that they had reached 80% of their loss limit tended to reach their personal limit and consequently increased their loss limit. Another important variable is the actual limit value. Just over one tenth of players whose limit was lower than NOK 900 and who did not receive any personalized feedback

changed their monetary limit (11%). The smallest group with an above-average limit change percentage was among those characterized by a red risk rating on the behavioral tracking tool *PlayScan* (i.e., a player displaying characteristics of problem gambling).

The present study was not without its limitations. The present study was carried out with players from one gambling operator in one jurisdiction, and the results might thus not be generalizable to other operators. Furthermore, the players' gambling activity on the *Norsk Tipping* platform might be just part of the players' overall gambling activity (i.e., some players may be gambling with other online operators). Additionally, it is mandatory for *Norsk Tipping's* players to choose a limit which is not common with all other operators (although some do require this). It should also be noted that *Norsk Tipping* players' maximum loss across players' total games portfolio is limited to NOK 20,000 which most other gaming operators do not employ. Out of the 6.7% of players who changed their limit, most increased it. This might be due to the fact that *Norsk Tipping* introduced a new limit-setting system shortly before the data for this study were collected. This means that the machine learning models predict players who are going to increase their loss limits. For this reason, it is important to replicate the present study in other countries and with other operators.

This present study shows that responsible gambling tool usage can potentially be predicted based on a player's behavioral tracking data. Given that the sparse empirical evidence suggests that limit-setting works, and that those players who voluntarily set their own limits reduce the money and/or time spent gambling when using such tools (Auer and Griffiths 2013), gaming operators should be encouraging their clientele to use such tools in order to predict which players are going to increase or decrease their loss limits. Gambling operators can consequently personalize communication about the utilization of pre-commitment (limit-setting) tools. Theoretically, this is not restricted to limit-setting tool usage and future studies should try to evaluate whether the use of other responsible gaming tools can also be predicted using player-tracking behavior.

Authors' Contribution Both authors contributed to the preparation of this manuscript.

Funding Information This research was funded by *Norsk Tipping*, the Norwegian Government's national lottery operator.

Compliance with Ethical Standards

Conflict of Interest The second author's university received funding from *Norsk Tipping* (the gambling operator owned by the Norwegian Government) for this work. The first author was sub-contracted by Nottingham Trent University. The second author has received funding for a number of research projects in the area of gambling education for young people, social responsibility in gambling, and gambling treatment from Gamble Aware (formerly the Responsibility in Gambling Trust), a charitable body which funds its research program based on donations from the gambling industry. Both authors undertake consultancy for various gaming companies in the area of social responsibility in gambling.

Ethical Approval All procedures performed in this study involving human participants were in accordance with the ethical standards of the research team's ethics committee and with the 1975 Helsinki Declaration.

Informed Consent Informed consent was obtained from all participants.

Appendix

Table 5 List of variables used in the algorithms

No.	Variable	Mean	SD
1	Amount bet	11,638	54,286
2	Theoretical loss	9812	50,567
3	Amount won	9240	60,562
4	Loss	2399	31,812
5	Number of playing days	16	17
6	Amount bet other games	194	1278
7	Amount bet Belago VLT	485	6634
8	Amount bet Lottery	1440	2254
9	Amount bet online gambling	6555	47,593
10	Amount bet pre-match betting	875	6674
11	Amount bet live betting	773	10,701
12	Amount bet Multix VLT	1099	8677
13	Amount bet pool sports betting	216	1696
14	Amount bet online casino	6047	47,041
15	Played other games	33%	47%
16	Played Belago VLT	2%	13%
17	Played lottery	90%	29%
18	Played online gambling	49%	50%
19	Played pre-match betting	34%	48%
20	Played live betting	14%	35%
21	Played Multix VLT	6%	23%
22	Played pool sports betting	19%	39%
23	Played online casino	16%	37%
24	Loss limit	5100	5909
25	Number loss limit increases	0.20	0.49
26	Number loss limit decreases	0.02	0.17
27	<i>PlayScan</i> red	5%	22%
28	Number of times feedback	0.25	0.78
29	Feedback	14%	35%
30	Maximum limit	0.10	0.30
31	Years since customer registration	9	4
32	Age	41	13
33	Gender (female)	28%	45%

Table 6 Variable descriptions

Variable	Description
Amount bet	Total amount bet
Theoretical loss	Theoretical loss = amount bet multiplied by the house advantage across all games
Amount won	Total amount won
Loss	Total amount bet minus total amount won
Number of playing days	Number of days with at least one game played
Amount bet other games	Amount bet in games not assigned to one specific category
Amount bet Belago VLT	Amount bet on the VLT category Belago
Amount bet Lottery	Amount bet on Lottery Games
Amount bet online gambling	Amount bet online across all games
Amount bet pre-match betting	Amount bet in pre-match betting
Amount bet live betting	Amount bet during live in-play betting
Amount bet Multix VLT	Amount bet on the VLT category Multix
Amount bet pool sports betting	Amount bet in pool sports betting
Amount bet online casino	Amount bet on online casino games
Played other games	Played at least one game in games not assigned to one specific category
Played Belago VLT	Played at least one game in the VLT category Belago
Played lottery	Played at least one lottery game
Played online gambling	Played at least one game online across all games
Played pre-match betting	Played at least one game in pre-match betting
Played live betting	Played at least one game in live betting
Played Multix VLT	Played at least one game on the VLT category Multix
Played pool sports betting	Played at least one game in pool sports betting
Played online casino	Played at least one game on online casino games
Loss limit	Personal global loss limit
Number loss limit increases	Number of increases of the personal global loss limit
Number loss limit decreases	Number of decreases of the personal global loss limit
<i>PlayScan</i> red	If player is classified as problematic by the player-tracking tool <i>PlayScan</i>
Number of times feedback	Number of times the player received feedback that 80% of the personal global loss limit was reached
Feedback	If the player received feedback that 80% of the personal global loss limit was reached
Maximum limit	If the player had chosen the maximum personal global loss limit (NOK 20,000)
Years since customer registration	Number of years since the player has registered with <i>Norsk Tipping</i>
Age	Age in years
Gender (female)	If the player is female

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Auer, M., & Griffiths, M. D. (2013). Voluntary limit setting and player choice in most intense online gamblers: An empirical study of gambling behaviour. *Journal of Gambling Studies*, 29, 647–660.
- Auer, M., Reiestad, S. H., & Griffiths, M. D. (2018a). Global limit setting as a responsible gambling tool: what do players think? *International Journal of Mental Health and Addiction*. Epub ahead of print. <https://doi.org/10.1007/s11469-018-9892-x>.

- Auer, M., Hopfgartner, N., & Griffiths, M. D. (2018b). The effect of loss-limit reminders on gambling behavior: A real-world study of Norwegian gamblers. *Journal of Behavioral Addictions, 7*, 1056–1067.
- Bakken, I. J., Gotestam, K. G., Grawe, R. W., Wenzel, H. G., & Øren, A. (2009). Gambling behavior and gambling problems in Norway 2007. *Scandinavian Journal of Psychology, 50*, 333–339.
- Bernhard, B. J., Lucas, A. F., & Jang, D. (2006). *Responsible gaming device research report*. Reno, NV: University of Nevada, Las Vegas International Gaming Institute.
- Bonello, M., & Griffiths, M. D. (2017). Analyzing consumer protection for gamblers across different online gambling operators: A descriptive study. *Gaming Law Review and Economics, 21*, 278–285.
- Bradley, A. P. (1997). The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition, 30*(7), 1145–1159.
- Broda, A., LaPlante, D. A., Nelson, S. E., LaBrie, R. A., Bosworth, L. B., & Shaffer, H. J. (2008). Virtual harm reduction efforts for Internet gambling: effects of deposit limits on actual Internet sports gambling behaviour. *Harm Reduction Journal, 5*, 27.
- Calado, F., & Griffiths, M. D. (2016). Problem gambling worldwide: an update and systematic review of empirical research (2000–2015). *Journal of Behavioral Addictions, 5*, 592–613.
- Calvosa, P. (2017). Responsible gambling strategies for internet gambling: An empirical investigation into the Italian gambling market. *International Journal of Business Management, 12*(7), 17–34.
- Cooney, C., Columb, D., Costas, J., Griffiths, M. D., & O’Gara, C. (2018). An analysis of consumer protection for gamblers across different online gambling operators in Ireland: A descriptive study. *International Journal of Mental Health and Addiction*. <https://doi.org/10.1007/s11469-018-9968-7>.
- Coussement, K., & De Bock, K. W. (2013). Customer churn prediction in the online gambling industry: The beneficial effect of ensemble learning. *Journal of Business Research, 66*(9), 1629–1636.
- Cox, D. R. (1958). The regression analysis of binary sequences. *Journal of the Royal Statistical Society. Series B (Methodological), 20*(2), 215–242.
- Ferris, J., & Wynne, H. (2001). *The Canadian Problem Gambling Index: Final Report*. Ottawa, Ontario: Canadian Centre on Substance Abuse.
- Focal Research Consultants. (2007). *Assessment of the behavioural Impact of the Responsible Gaming Device (RGD) features: analysis of Nova Scotia player-card data*. The Windsor trial. Report prepared for the Nova Scotia Gaming Corporation.
- Forsström, D., Hesser, H., & Carlbring, P. (2016). Usage of a responsible gambling tool: A descriptive analysis of latent class analysis of user behavior. *Journal of Gambling Studies, 32*, 889–904.
- Friedman, J. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics, 29*(5), 1189–1232.
- Götestam, K. G., & Johansson, A. (2003). Characteristics of gambling and problematic gambling in the Norwegian context: A DSM-IV based telephone interview study. *Addictive Behaviors, 28*, 189–197.
- Griffiths, M. D. (2012). Internet gambling, player protection and social responsibility. In R. Williams, R. Wood, & J. Parke (Eds.), *Routledge handbook of internet gambling* (pp. 227–249). London: Routledge.
- Griffiths, M. D., Wood, R. T. A., & Parke, J. (2009). Social responsibility tools in online gambling: A survey of attitudes and behaviour among Internet gamblers. *CyberPsychology and Behavior, 12*, 413–421.
- Hanley, J. A., & McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology, 143*(1), 29–36.
- Harris, A., & Griffiths, M. D. (2017). A critical review of the harm-minimisation tools available for electronic gambling. *Journal of Gambling Studies, 33*, 187–221.
- Hothorn, T., Hornik, K., & Zeileis, A. (2006). Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical Statistics, 15*(3), 651–674.
- Ivanova, E. N., Magnusson, K., & Carlbring, P. (2019). Deposit limit prompt in online gambling for reducing gambling intensity: A randomized controlled trial. *Frontiers in Psychology, 10*, 639.
- Jie, M. A., Collins, G. S., Steyerberg, E. W., Verbakel, J. Y., & van Calster, B. (2019). A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of Clinical Epidemiology, 110*, 12–22.
- John, G. H., & Langley, P. (1995). Estimating continuous distributions in Bayesian classifiers. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* (pp. 1–8). San Mateo, CA: Morgan Kaufman Publishers, San Mateo.

- Jonsson, J. (2006). An overview of prevalence surveys of problem and pathological gambling in the Nordic countries. *Journal of Gambling Issues*, 18, 31–38.
- Kazhaal, Y., Chatton, A., Bouvard, A., Khiari, H., Achab, A., & Zullino, D. (2011). Internet poker websites and pathological gambling prevention policy. *Journal of Gambling Studies*, 29, 51–59.
- Kim, H. S., Wohl, M. J., Stewart, M. K., Sztainert, T., & Gainsbury, S. M. (2014). Limit your time, gamble responsibly: Setting a time limit (via pop-up message) on an electronic gaming machine reduces time on device. *International Gambling Studies*, 14, 266–278.
- Lemeshow, S., & Hosmer Jr., D. W. (1982). A review of goodness of fit statistics for use in the development of logistic regression models. *American Journal of Epidemiology*, 115(1), 92–106.
- Lever, J., Krzywinski, M., & Altman, N. (2016). Points of significance: Model selection and overfitting. *Nature Methods*, 13(9), 703–704.
- Liaw, A., & Wiener, M. (2002). Classification and regression by random forest. *R News*, 2(3), 18–22.
- Lund, I., & Nordlund, S. (2003). *Pengespill og pengeproblemer i Norge [Gambling and monetary issues in Norway] (Rapport nr. 2/2000)*. Oslo: Statens Institutt for Rusmiddelforskning.
- Marionneau, V., & Järvinen-Tassopoulos, J. (2017). Consumer protection in licensed online gambling markets in France: The role of responsible gambling tools. *Addiction Research and Theory*, 6, 436–443.
- Meyer, G., Hayer, T., & Griffiths, M. (2009). *Problem gambling in Europe: Challenges, prevention, and interventions*. New York: Springer.
- Óskarsdóttir, M., Bravo, C., Verbeke, W., Sarraute, C., Baesens, B., & Vanthienen, J. (2017). Social network analytics for churn prediction in Telco: Model building, evaluation and network architecture. *Expert Systems with Applications*, 85, 204–220.
- Pallesen, S., Molde, H., Mentzoni, R. A., Hanss, D., & Morken, A. M. (2016). *Omfang av penge-og dataspillproblemer i Norge 2015*. Bergen: University of Bergen.
- Percy, C., França, M., Dragičević, S., & D'Avila Garcez, A. (2016). Predicting online gambling self-exclusion: An analysis of the performance of supervised machine learning models. *International Gambling Studies*, 16(2), 193–210.
- Philander, K. S. (2014). Identifying high-risk online gamblers: A comparison of data mining procedures. *International Gambling Studies*, 14(1), 53–63.
- Riffenburgh, R. H., & Clunies-Ross, C. W. (1960). Linear discriminant analysis. *Pacific Science*, 14(3), 251–256.
- Safari, S., Baraloo, A., Elfil, M., & Negida, A. (2016). Evidence based emergency medicine: Part 5 Receiver operating curve and area under the curve. *Emergency*, 4(2), 111–113.
- Schellinck, T., & Schrans, T. (2011). Intelligent design: How to model gambler risk assessment by using loyalty tracking data. *Journal of Gambling Issues*, 26, 51–68.
- Shaffer, H. J., & Kom, D. A. (2002). Gambling and related mental disorders: A public health analysis. *Annual Review of Public Health*, 23, 171–212.
- Smeaton, M., & Griffiths, M. D. (2004). Internet gambling and social responsibility: An exploratory study. *CyberPsychology and Behavior*, 7, 49–57.
- Stewart, M. J., & Wohl, M. J. A. (2013). Pop-up messages, dissociation, and craving: How monetary limit reminders facilitate adherence in a session of slot machine gambling. *Psychology of Addictive Behaviors*, 27, 268–273.
- Stucki, S., & Rihs-Middel, M. (2007). Prevalence of adult problem and pathological gambling between 2000 and 2005: An update. *Journal of Gambling Studies*, 23, 245–257.
- Walker, D. M., Litvin, S. W., Sobel, R. S., & St-Pierre, R. A. (2015). Setting win-limits: An alternative approach to “responsible gambling”. *Journal of Gambling Studies*, 31, 965–986.
- Williams, R. J., Volberg, R. A. & Stevens, R. M. G. (2012). *The population prevalence of problem gambling: Methodological influences, standardized rates, jurisdictional differences, and worldwide trends*. Report prepared for the Ontario Problem Gambling Research Centre and the Ontario Ministry of Health and Long Term Care.
- Wohl, M. J., Gainsbury, S., Stewart, M. J., & Sztainert, T. (2013). Facilitating responsible gambling: The relative effectiveness of education-based animation and monetary limit setting pop-up messages among electronic gaming machine players. *Journal of Gambling Studies*, 29, 703–717.
- Wood, R. T. A., & Griffiths, M. D. (2010). Social responsibility in online gambling: Voluntary limit setting. *World Online Gambling Law Report*, 9(11), 10–11.

Affiliations

Michael Auer¹ · Mark D. Griffiths²

✉ Mark D. Griffiths
mark.griffiths@ntu.ac.uk

Michael Auer
m.auer@neccton.com

¹ neccton Gmbh, Davidgasse 5, 7052 Muellendorf, Austria

² International Gaming Research Unit, Psychology Department, Nottingham Trent University, 50 Shakespeare Street, Nottingham NG1 4FQ, UK