playtech PROTECT

BEHAVIOURAL MARKERS OF HARM

Insights from BetBuddy machine learning models

Industry Research Brief Vo2. (1) – Markers of Harm
December 2021

Behavioural Markers of Harm:
Insights from BetBuddy machine learning models

What is the topic?

Gambling operators typically have visibility of diverse datapoints for a gambler's behaviour, particularly where their play is registered, i.e. identified against a specific player account number, rather than anonymous play. Behavioural markers of harm seek to identify patterns of activity in such datapoints that suggest a player is at risk.

Such behaviours can include choice of gambling activity, pattern and outcome of wagers, timing and variability of play, financial transactions like deposits or withdrawals, in-session patterns like loss chasing, communication activity (e.g. conversations with customer services), responsible gambling tool usage, etc.

This report analyses which specific markers of harm are the most powerful predictors of player risk in five of BetBuddy's machine learning models that have been deployed in a live environment. These models analyse player behaviours and identify individuals more likely to be at risk of harm

Why is it important?

Behavioural markers of harm are key to responsible gambling techniques that hope to intervene or engage with players in personalised ways. Such personalised interventions are more likely to succeed in raising awareness, reducing risky play patterns, and preventing harmful outcomes.

Fig. 1. Framework for player risk

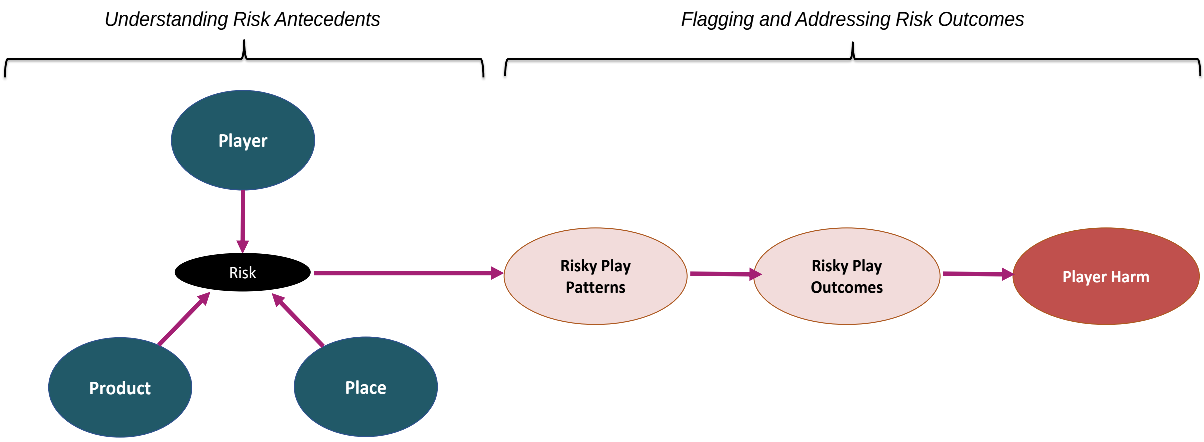

Interventions based on behavioural markers of harm are of interest because they take place closer to the point of harm, i.e. they take place in the pale red sections of the diagram above. Because it is more targeted, this approach enables more proactive, personalised interventions that would be inappropriate or ineffective to apply to the full player base. In a previous industry research brief, we cite research suggesting that personalised targeted messages are more effective at encouraging deposit limit setting than mass communications (vol 1; issue 2).

Other strategies to support responsible gambling might choose to intervene at other points in the figure 1 diagram. For instance, interventions at the level of risk antecedents, such as product-level regulation or changes to the platforms and places where people place bets. Full strategies can also include investments in broad-based awareness campaigns and the provision of support services directly aimed at players experiencing harm.

What did the research do?

Playtech Protect provides a service, BetBuddy, to gambling operators to help them analyse player data to understand which players appear to be at greater risk of harm. The service comprises a series of analytics engines, including machine learning models that are tailored to the player bases of individual brands.

These models each draw on around 40 behavioural features optimised for that player base that summarise each player's recent pattern of behaviour. The included features are similar but not identical across models, reflecting different data availability and which types of features are most effective in that operator's context. A machine learning model called random forest then draws on these features to test how similar each player looks to other players identified by the operator as at risk. We discuss the methodology and motivation for choosing the random forest model in academically published work (see Percy et al, 2016).

This analysis examines which behavioural features contribute most to the decision making of random forest models built in a broadly similar way, with similar algorithmic optimisation techniques and harm outcomes (being players who self-excluded from gambling for at least six months[i]). By examining similar models across different player bases, the research helps understand the consistency of different behavioural markers of harm across different types of game, such as online slots, online bingo, and other online gaming products, and across different player bases from different countries. The models perform well with average test set performance of 95%[ii].

What did the research find?

The five models include over 70 different features between them, with 41 features appearing in at least two models. The top 16 features, in terms of how much they contribute to the models' estimation of player risk, are shown in figure 2 overleaf.

These top 16 features span a wide range of categories:

- Deposits
- Game selection
- Frequency of days played
- Number of bets per day
- Responsible Gambling tool usage
- Session time
- Bet amount
- Declined deposits[iii].

This diversity of categories suggests that optimised models do not rely on single aspects of player behaviour. Nonetheless, features connected with depositing activity (shaded dark blue) and play frequency (green) are particularly common key predictors across the models.

Fig. 2. Profiles of top 16 features

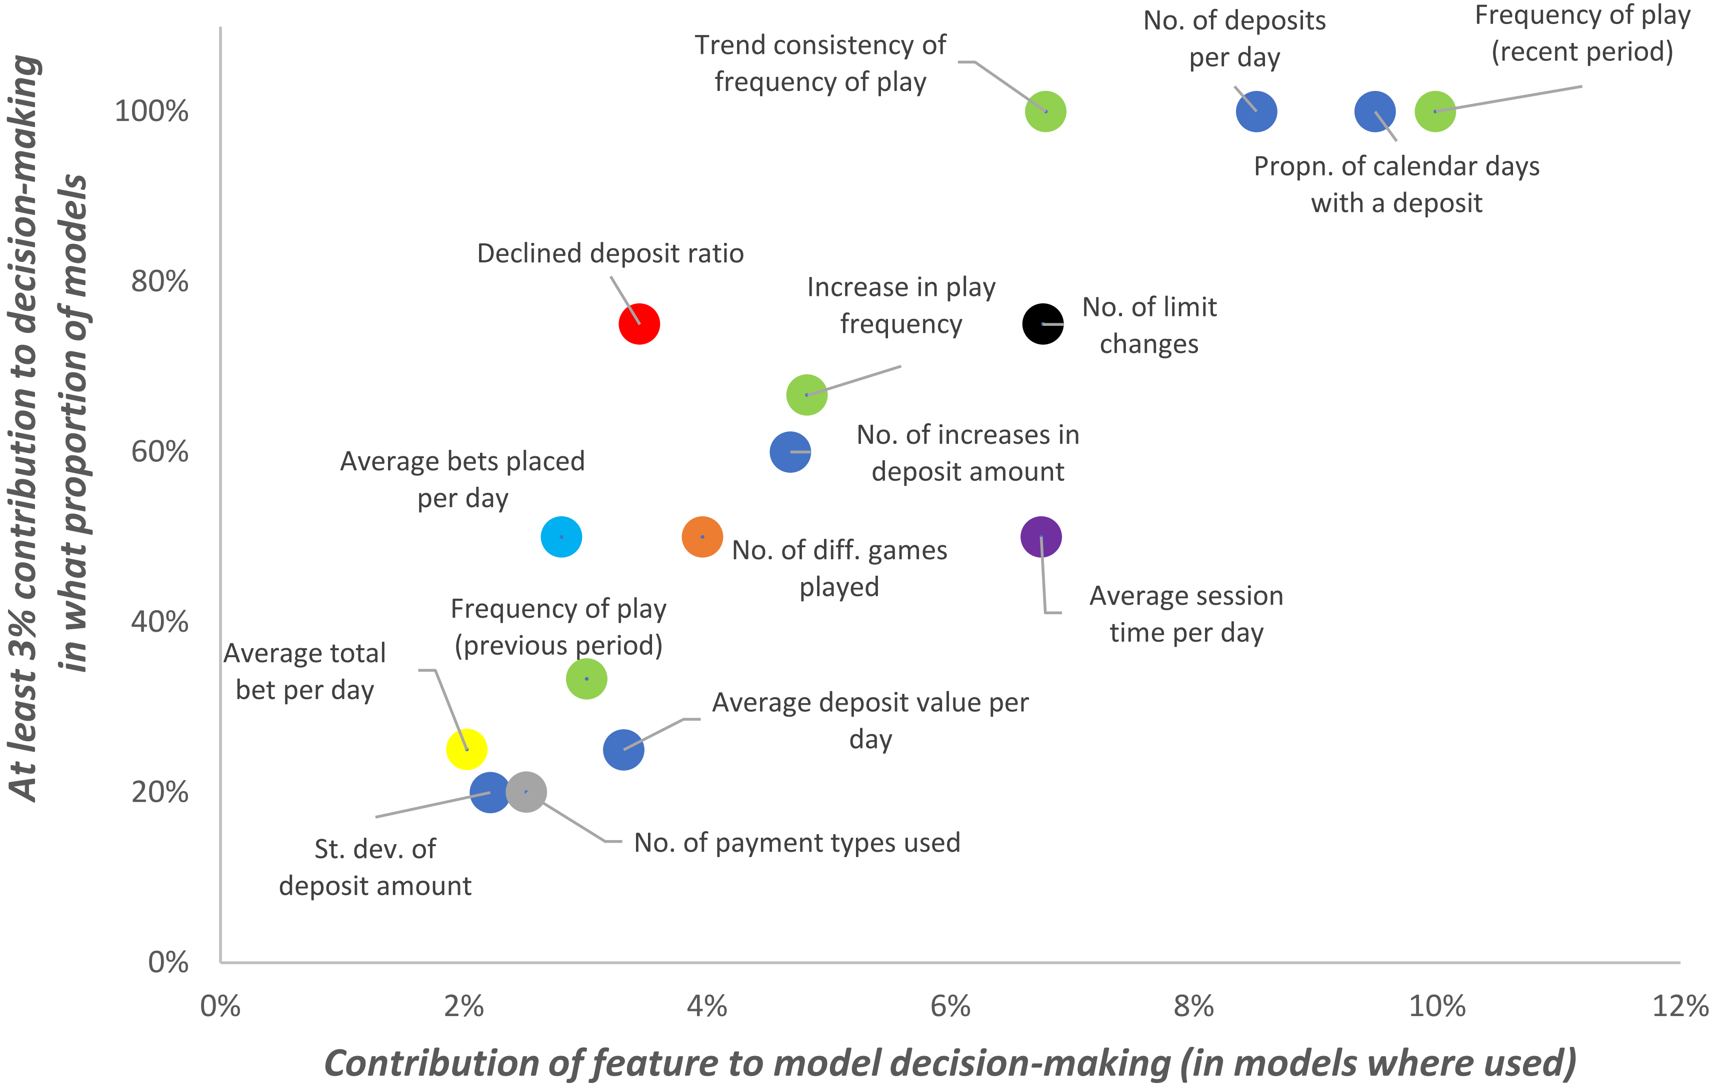

Colour coding groups features by type, e.g. those connected with deposits are shaded dark blue, those connected to play frequency are shaded green.

Another way of analysing features confirms the importance of deposits and play frequency seen in figure 2. In figure 3 below, we examine 60 main model features that align to particular categories.[iv] The advantage of this type of analysis is drawing on insights across many features, not just the top-performing ones, recognising that several small but separate contributions might add up to a large contribution overall. However, it is important to note the number of distinct features used in each category, since categories with more features will generally have a better chance of contributing to model decision-making (details are provided in the columns on the right hand side of figure 2).

Fig. 3. Analysis of aggregate importance by category of player behaviour

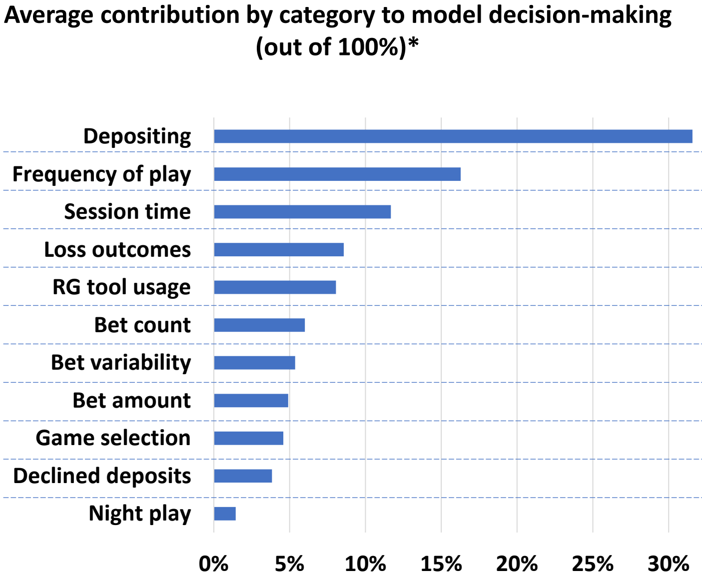

Category	*Number of distinct features available*	*Avg. # of features per model***	***Avg. model contribution per feature in category***
Depositing	*9*	*6.2*	***5.1%***
Frequency of play	*5*	*3.0*	***5.4%***
Session time	*4*	*4.0*	***2.9%***
Loss outcomes	*13*	*8.2*	***1.0%***
RG tool usage	*6*	*2.4*	***3.4%***
Bet count	*5*	*3.4*	***1.8%***
Bet variability	*4*	*4.0*	***1.3%***
Bet amount	*6*	*3.6*	***1.4%***
Game selection	*4*	*2.2*	***2.1%***
Declined deposits	*3*	*1.2*	***3.2%***
Night play	*1*	*1.0*	***1.4%***

* Averaged across models where at least one constituent feature is available; only categories with presence across 4+ models included
** Just those models contributing to the displayed average included

As well as understanding which features are typically most influential on model assessments of player risk, we would like to understand the consistency of top-performing features across the five models.

The top five features across each of the five models account for 36%-58% of model decision-making. Each model included at least one top five feature connected to recent depositing, such as the number of deposits, the proportion of days on which a deposit took place, the average deposit size, or measures of the increase in deposits. Recent frequency of play (proportion of calendar days on which they gamble with real money) is also included in three of the five models. However, beyond this there are variations in the top features.

Changes in play limit setting was available in four of the five models. It is the most important predictor in one model, influential but not in the top five of another two, and unimportant in the fourth. Other features related to responsible gambling tools in the models are not influential, so the difference in limit setting interpretation is unlikely to reflect a marginal preference for an alternative feature that captures similar underlying behaviours/risks. The same is true of number of different games played and different payment types used, which are key variables in one model only, despite being available in at least four.

Features can also be grouped by whether they look at the level of a particular activity (e.g. minutes played per day) or examine changes/trends over time (e.g. increase in average minutes played per day over the last month) and volatility over time (e.g. variance of daily play time).[v] When viewed through this lens, a diverse approach to analysing a particular behaviour appears to be important. Levels of activity are the most important type of feature, but measures of trend and variance account for over 40% of model decision-making in total. This is in line with BetBuddy's experience that including a diverse range of features and types of measurement typically results in better performing models.

Limitations to the analysis

Similar to all analyses of a particular set of features, insights on importance are drawn only relative to the other features chosen.[vi] It is always possible that other features not yet designed or analysed might be more effective. Additionally, the research focuses on one type of proxy for harm (serious self-exclusion[vii]), being the only proxy available at scale and in a standardised form across operators. However, certain players may be at risk but would not consider self-exclusion, or might not even have patterns of behaviour similar to others who do self-exclude. Finally, we emphasise that the analysis is not necessarily causal and that features can often be important in non-linear ways with counterintuitive interpretations once the values of other features are controlled for; see Percy et al (2019) for an example using night play.

What might the gambling industry do in response to this research?

This analysis provides an indication of which features are most commonly important when seeking to identify at risk players: deposit frequency and amount; and play frequency and increases in frequency.

Such key features might be priorities for the development of new models or business rules seeking to identify players at risk, e.g. as specific triggers for player interventions where a composite risk measure is not available. They might also be priority data points to use in personalising feedback to players to understand which aspects of their behaviour might be placing them at risk. However, with many other features important and varying across models, we recommend operators consider a broad range of features and aim to tailor models to their specific context wherever possible. For instance, general thresholds based on depositing might be straightforward to develop, but would misclassify some at-risk players compared to considering a full range of variables and tailoring models to product types and players.

Other modelling recommendations include incorporating financial transaction data (not just play data like wagering and outcomes)[viii] and examining diverse ways of summarising patterns of behaviour. For instance, recent levels of activity, previous levels of activity, the scale and consistency of behaviour trends over time, and variability in behaviour all prove useful to model outcomes. This strongly suggests that focusing just on current levels of activity (e.g. average bet amount) is insufficient to best detect risky play and that more diverse statistics should be considered (e.g. increase in bet amount, the volatility of bet amount etc.).

The industry and its stakeholders are advised to consider these results in the context of other research that has examined behavioural markers of harm. Early statistical work on this topic took place with Xuan & Shaffer (2009), Braverman & Shaffer (2010), and Dragicevic et al (2011, 2015). This work typically identified priority features in domains like play frequency, intensive betting, wager amount variability, the adoption of riskier gambling positions, increased losses, and increased wager size in the first month. Recent work, such as studies by PWC (2017) and the Behavioural Insights Team (2021a, 2021b), additionally emphasises behaviours like night play, spend relative to disposable income, declined deposits, and midweek betting, as well as daily triggers, demographic markers, and the value of personalising analysis by player segment.

The analysis presented here complements previous studies by taking a different approach, both in examining how highly complex, non-linear models interpret behavioural features and in contrasting the findings across diverse player bases from online operators.

How can I find out more?

To find out more or if you have any suggestions for future topics to be addressed via the Industry Research Brief, please contact the research team via protect@playtech.com.

References

Behavioural Insights Team. (2021a). *Patterns of Play - Gambling behaviour: What can bank transaction data tell us? A feasibility study (Part 1: Analysis of Monzo customer data).* London: GambleAware

Behavioural Insights Team. (2021b). *Patterns of Play - Gambling behaviour: What can bank transaction data tell us? A feasibility study (Part 2: Analysis of HSBC UK customer data).* London: GambleAware

Braverman, J. & Shaffer, H. (2010). How do gamblers start gambling: Identifying behavioral markers for high-risk internet gambling. *European journal of public health (22; 273-8).* 10.1093/eurpub/ckp232

Dragicevic, S., Tsogas G., & Kudic, A. (2011). Analysis of casino online gambling data in relation to behavioural risk markers for high-risk gambling and player protection. *International Gambling Studies, 11:3, 377-391*

Dragicevic, S., Percy, C., Kudic, A., & Parke, J. (2015). A descriptive analysis of demographic and behavioral data from internet gamblers and those who self-exclude from online gambling platforms. *Journal of Gambling Studies, 31(1)*

Future Anthem. (2021). *Research into the impact of casino games on player markers of harm: Summary Findings.* London: Future Anthem / Gamesys

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*

Percy, C., Garcez, A., Dragicevic, S., & Sarkar, S. (2019). Understanding the Risk Profile of Gambling Behaviour through Machine Learning Predictive Modelling and Explanation. *33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada*. Available via https://kr2ml.github.io/2019/papers/KR2ML_2019_paper_33.pdf

Percy, C., Garcez, A., Dragicevic, S., & Sarkar, S. (2020). Lessons Learned from Problem Gambling Classification: Indirect Discrimination and Algorithmic Fairness. Paper presented at AI for Social Good - AAAI Fall Symposium 2020.

PWC. (2017). *Remote Gambling Research: Interim report on Phase II*. London: GambleAware

Xuan Z., & Shaffer H. (2009). How do gamblers end gambling: Longitudinal analysis of Internet gambling behaviors prior to account closure due to gambling related problems. *J Gambling Stud 2009;25:239–52*

[i] Specifically, we look at the feature importance of the random forest algorithm, scaled such that total importance across all features in the model sums to 100%. A random forest is made up of many decision trees constructed with a random subsample of features and players, where the optimal decision rule on one feature is chosen to split the sample at each decision node in each tree. The importance of a feature is based on each time that feature was selected to split on and how much the squared error (over all trees) decreased as a result. For more details, please see https://docs.h2o.ai/h2o/latest-stable/h2o-docs/variable-importance.html

[ii] Test set performance refers to the performance of the model on data not used in the training of the model. This is important because "training set accuracy" can arbitrarily approach 100% with large models and such "over-trained" models typically perform less well in deployment. The standard metric we use is cross-validation AUROC, which captures the pay-off between false positives and missed positives across a full range of possible threshold values for specifying a particular player's score as at risk.

[iii] i.e. attempted deposits declined by the player's payment provider, often because the account does not have sufficient funds, card details are entered incorrectly (less common for regular deposits using saved card details), or because of a rule set up by the bank or player to prevent certain transactions.

[iv] Figure 3 excludes features under experimental review that are currently being tested in only one model, specifically those features connected to withdrawal activity, channel usage, weekday activity, and length of time as customer.

[v] Other distributional features of behavioural activity, e.g. higher moments like skewness, have also been examined but not proved major contributors to model decision-making to date.

[vi] This has several implications. For instance, removing a high impact feature might not necessarily cause significant declines in model performance, if there are other features available that capture similar underlying behavioural dynamics. If many similar features are included, they might be individually low impact for the model but collectively important (which is why figure 3 presents aggregate analysis by category), noting that the inclusion of many features will tend to add some value to decision-making as a result of chance variation if nothing else, especially for random forest models that repeatedly and randomly use only a sample of the available features in model construction. Other approaches, such as testing the impact of dropping features or creating orthogonal combinations of features via principal component analysis, can also be used to interpret feature importance. The impact of dropping features can also be tested in combination. For instance, with 40 features, there are 780 unique pairs of variables or 9880 unique triplets which could be dropped one by one to see the impact on the model. Combinations of four or more variables could also be dropped on a hypothesis-led basis. Note that linear covariance between features is unhelpful for exploring this phenomenon, limiting the value of PCA as traditionally applied, since the random forest algorithm can identify complex, non-linear relationships (see Percy et al, 2020 for a broader discussion).

[vii] Defined as previously regular players who request gambling operators to prevent their access to the platform for at least six months.

[viii] Depositing activity features are consistently more useful to model predictions than features connected to losses, despite the two being closely related and despite most models having access to more features that relate to losses than deposits.