Understanding Risk Models and Consumer Risk Profiles in Gambling Using Knowledge Extraction

Chris Percy – Lead Researcher, BetBuddy

BetBuddy

What motivates explainability?

New Approach

Create a suite of machine learning techniques that produce more explainable models, while maintaining a high level of learning performance

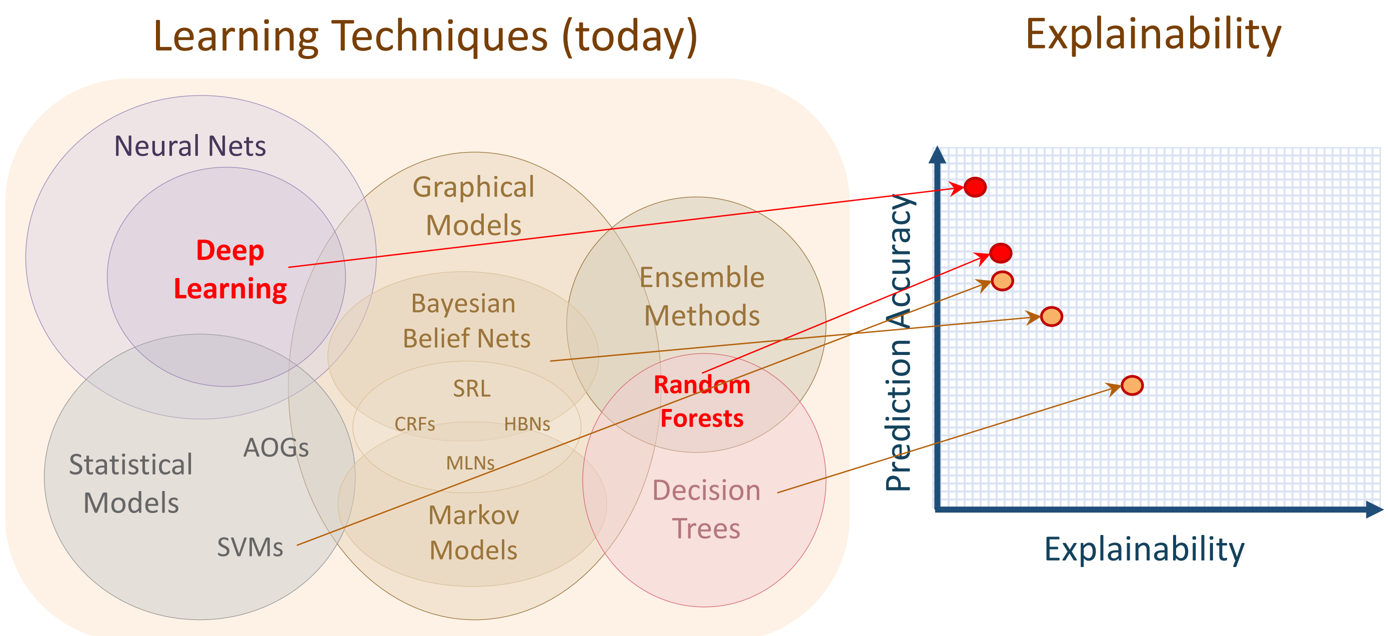

Source: Defense Advanced Research Projects Agency, 2016

“The current generation of AI systems offer tremendous benefits, but their effectiveness will be limited by the machine’s inability to explain its decisions and actions to users.

Explainable AI will be essential if users are to understand, appropriately trust, and effectively manage this incoming generation of artificially intelligent partners.”

What is explainable or human-like AI?

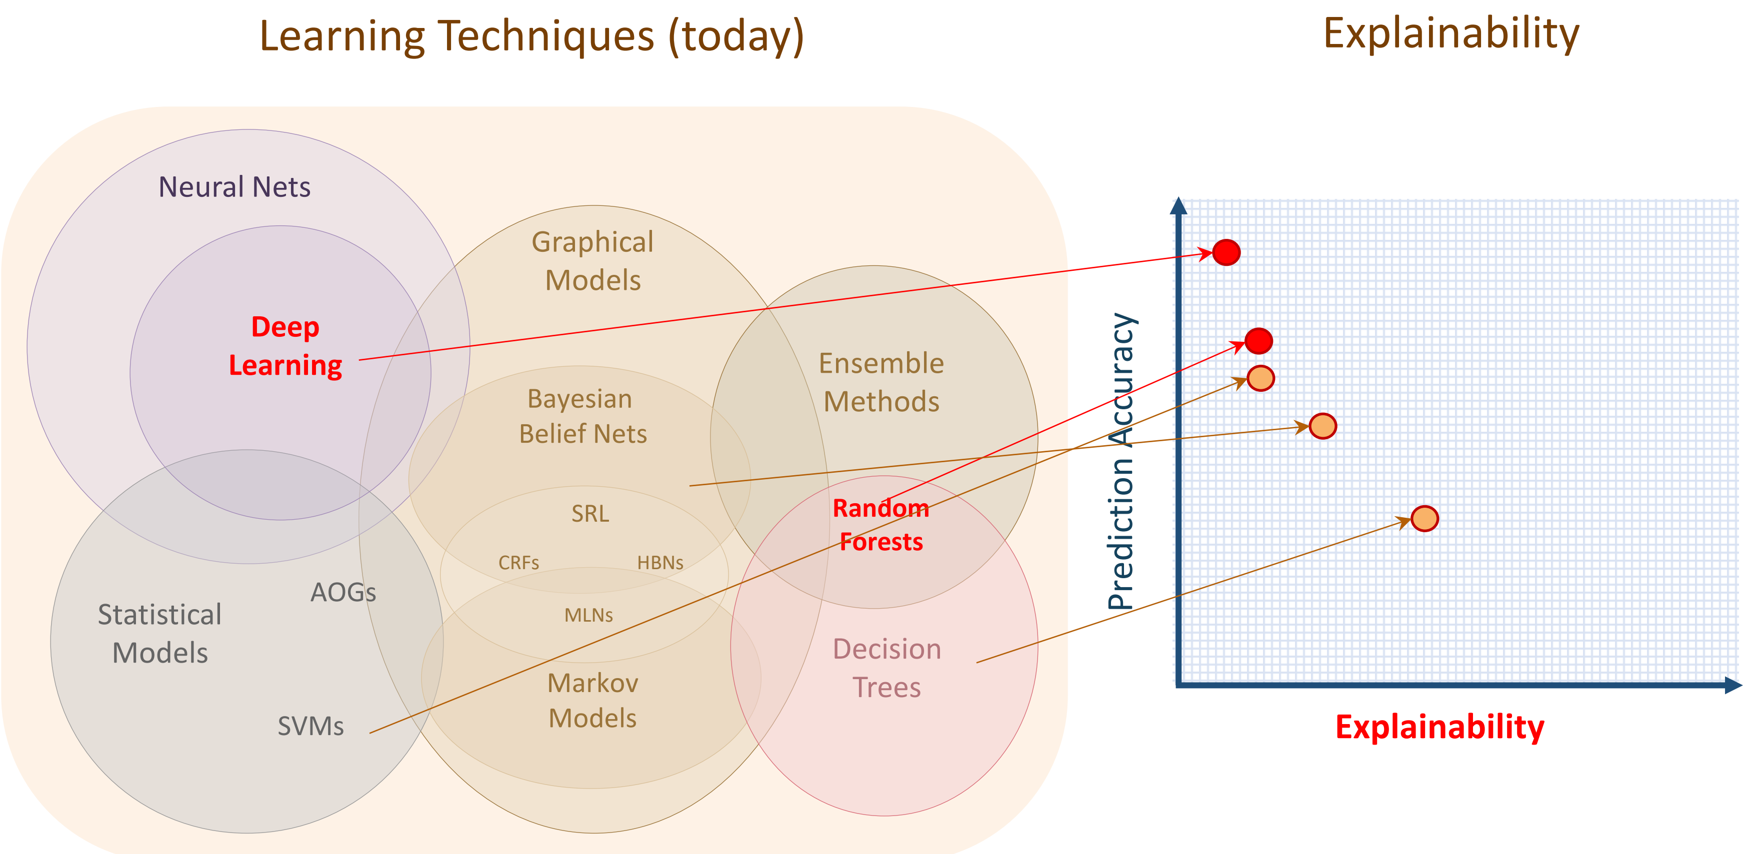

BetBuddy

What is explainable or human-like AI?

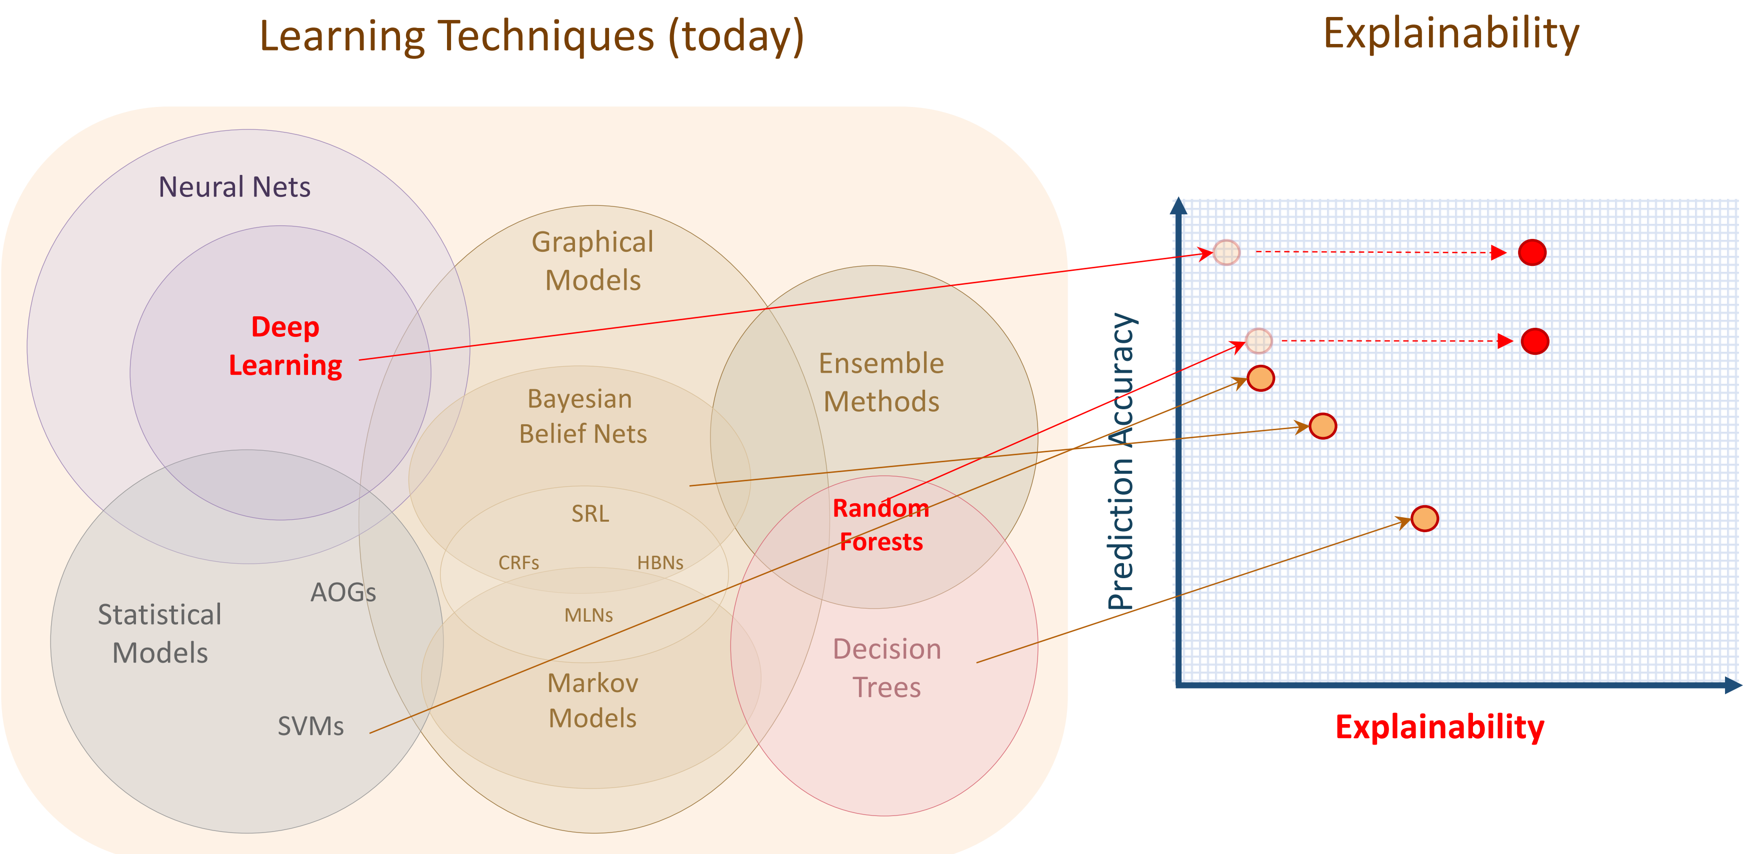

Where might explainable AI be useful?

User	Use Case
Physician	Validate 3rd party health diagnostics e.g., why is a patient included in a high-risk cluster?
Government & Enterprise	Understand strengths and weaknesses in critical AI applications e.g., security, terrorism, crime
Regulator	Scrutinise high-risk decision making models in industry e.g., gambling, mortgages, pay day loans, credit cards
AI Provider	Understand, test, and debug AI applications prior to deployment e.g., autonomous vehicles, image recognition
eCommerce	Explain an algorithm-driven profiling decision to a customer e.g., EU GDPR
Consultant	Improve AI-based applications for their enterprise customers

Where might explainable AI be useful?

User	Use Case
Physician	Validate 3rd party health diagnostics e.g., why is a patient included in a high-risk cluster?
Government & Enterprise	Understand strengths and weaknesses in critical AI applications e.g., security, terrorism, crime
Regulator	Scrutinise high-risk decision making models in industry e.g., **gambling**, mortgages, pay day loans, credit cards
AI Provider	Understand, test, and debug AI applications prior to deployment e.g., autonomous vehicles, image recognition
eCommerce	Explain an algorithm-driven profiling decision to a customer e.g., EU GDPR
Consultant	Improve AI-based applications for their enterprise customers

Regulators want industry to use AI for consumer protection

The Commission's position

We are somewhat concerned by the points raised in paragraph 21.1 – which appear to run counter to our expectation that licensees move towards a culture of evaluation or one in which data is used for non-commercial reasons as effectively as it is used for commercial reasons. One respondent noted that the effectiveness of existing requirements has yet to be measured. However, we expect licensees to provide ongoing feedback and insight as to how existing measures could be further enhanced and, where appropriate, developed.

There is, to some extent, a greater expectation on the remote gambling sector with its account based play and unlimited stake and prizes, to narrow the gap between customer retention and consumer protection measures. The expectation is that licensees will use datasets and player analytics (tools often used for commercial reasons) to monitor the effectiveness of gambling management controls and target those consumers most at risk of problem gambling. It is no longer sufficient to use player take up as the sole means through which to measure the effectiveness of social responsibility controls.

In relation to our statement of principles, it is not correct to suggest that our aim to deliver evidence based regulation restricts our duty to regulate in the public interest with regard to, and in pursuit of, the licensing objectives. The Commission's statement of principles sets out our approach as follows:

> *In interpreting the available evidence, the Commission will take a precautionary approach. For example, caution may be justified where evidence is mixed or inconclusive, and the Commission would not want to restrict its discretion by requiring conclusive evidence that something was unsafe before taking measures to restrict it.*

"The expectation is that licensees will use data sets and player analytics (tools often used for commercial reasons) to monitor the effectiveness of gambling management controls and target those customers most at risk of problem gambling."

However, there's a reluctance to rely on black box decision making

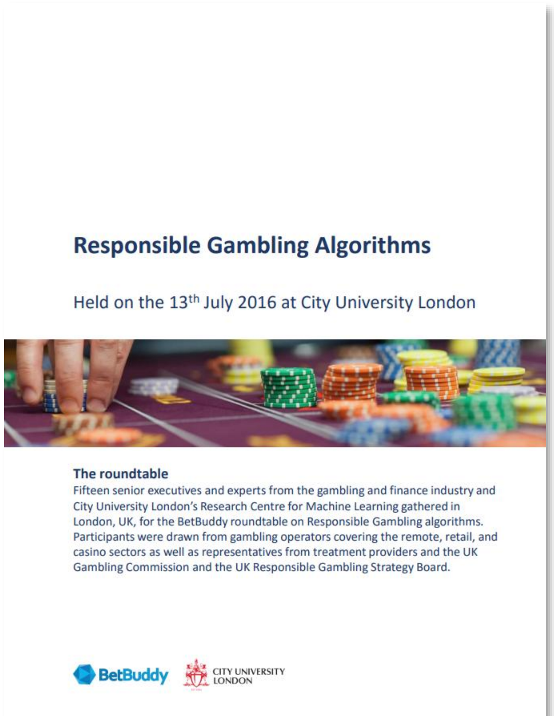

Responsible Gambling Algorithms

Held on the 13th July 2016 at City University London

The roundtable

Fifteen senior executives and experts from the gambling and finance industry and City University London's Research Centre for Machine Learning gathered in London, UK, for the BetBuddy roundtable on Responsible Gambling algorithms. Participants were drawn from gambling operators covering the remote, retail, and casino sectors as well as representatives from treatment providers and the UK Gambling Commission and the UK Responsible Gambling Strategy Board.

BetBuddy CITY UNIVERSITY LONDON

Transparency is important to enable an operator to manage any regulatory challenge with regards to a model that had been deployed

Roger Parkes

Global Head of Compliance, Betway

Instinctively, greater model understanding would be a higher priority compared with greater accuracy

Paul Hope

Programme Director, UK Gambling Commission

Results from a poll of experts from the New Horizons in Responsible Gambling Conference (2016) showed that industry, regulators, and treatment providers would prefer an algorithm that was 75% accurate (in predicting gambling related harm) and that was fully understandable, compared with an algorithm that was 90% accurate but was also a blackbox

BetBuddy

Two strands of research

TREPAN

Algorithms which seek to extract a compact decision tree from a given complex model. The original motivation was to represent a neural network model in a tree structure which could be more interpretable than a neural network classification model.

TREPAN generates decision rules of type M of N , N of N, 1 of N, or 1 of 1. In an M of N configuration, a tree node contains N distinct tests.

RANDOM FOREST TREE EXPLAINER

When considering a decision tree, it is intuitively clear that for each decision that a tree (or a forest) makes there is a path (or paths) from the root of the tree to the leaf, consisting of a series of decisions, guarded by a particular feature, each of which contribute to the final predictions.

How TREPAN works

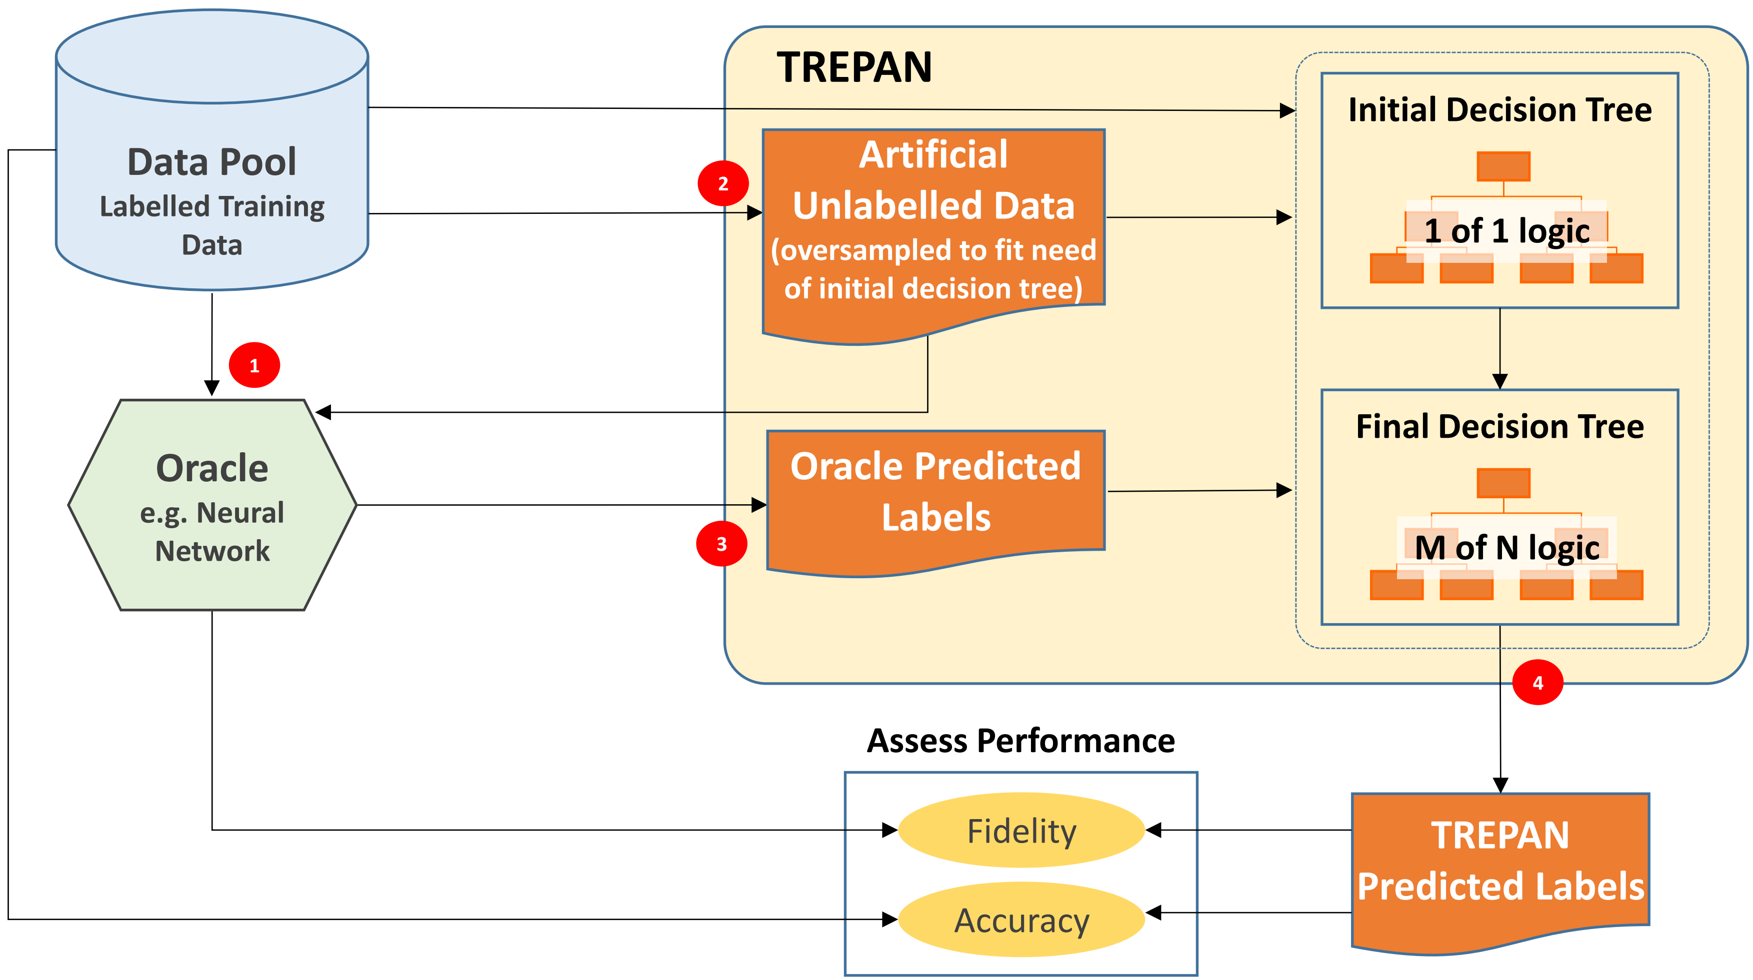

Key parameter choices in (1) Creation of and usage of oversampled data, (2) Growing 1-of-1 tree and (3) Pruning to M-of-N tree → Many different possible output trees – Domain expert needed to review

Early promise using TREPAN

Neural Net: 33 Features, 500+ Weights

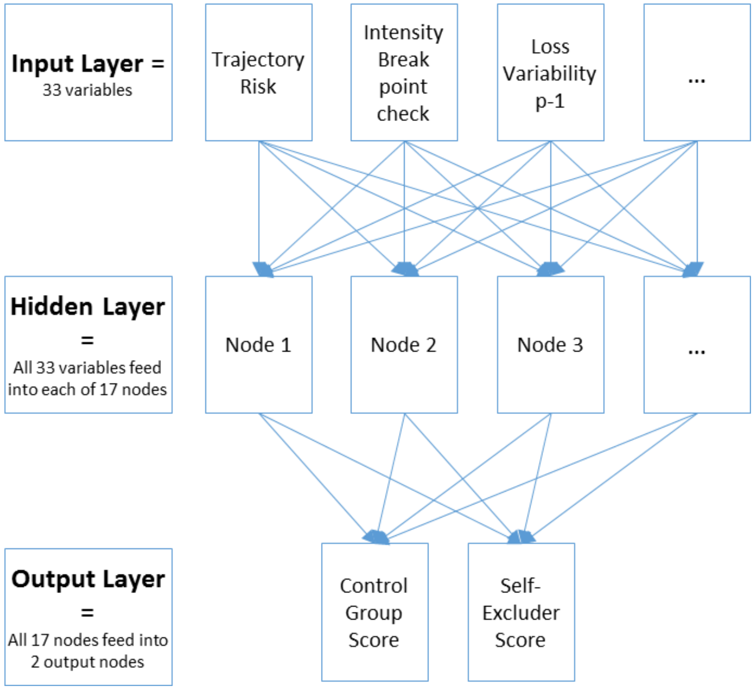

Compact Decision Tree: 8 Nodes

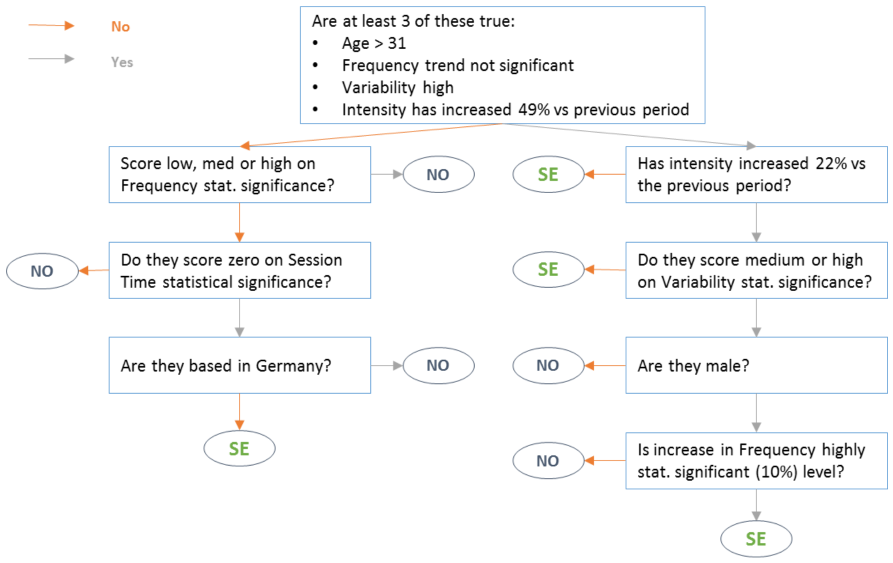

with competitive accuracy and fidelity

Enables extraction of rules and narratives for human validation

Compact Decision Tree

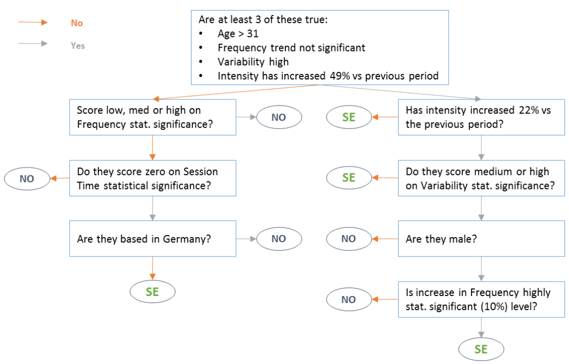

Rules and Narratives

Model Rule	Model Narrative	Viable Narrative?	Why
IF Age >31, Male, Variability High, Intensity Increase >49%, Frequency Increase Sig. >90%, THEN AT RISK	*Men aged over 31 who bet frequently, are increasing their number of bets, and who have a very volatile betting style, are at higher risk of problem gambling*	**OK**	Narrative supports addictions research
IF Age >31, Variability High, Intensity Increase >49%, Nationality ≠ German, THEN AT RISK	*People aged over 31 who have a very volatile betting pattern, are increasing their number of bets, and are not German are at higher risk of problem gambling*	**NO**	Narrative weak as nationality is not a valid indicator of addiction
Etc…			

Random forest tree explainer

Conceptual Approach

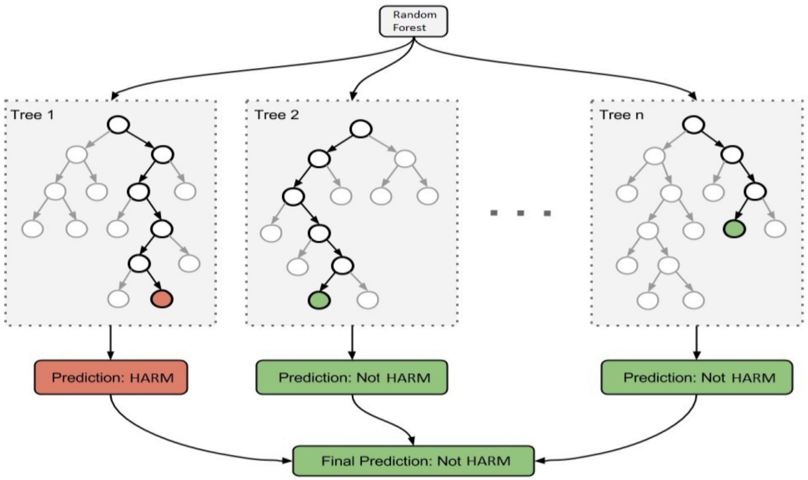

- A grid search optimized Random Forest classification model was built with a forest size of 200 binary trees having unlimited depth
- The primary focus was to extract sample-level information from the random forest model by traversing each node of each tree, which contains a either data feature and a threshold value based on which the tree is split into a further two branches; or a probability for the class at leaf nodes
- Each data sample consists of risk profile of a player based on his/her gambling pattern. Data samples are passed through each tree with in the forest model to arrive to the prediction
- Each unseen data sample is passed through the random forest model, and the number of occurrences of each feature (frequency value) in the decision path of each tree was recorded.

Random forest tree explainer

Weight Factor Approaches

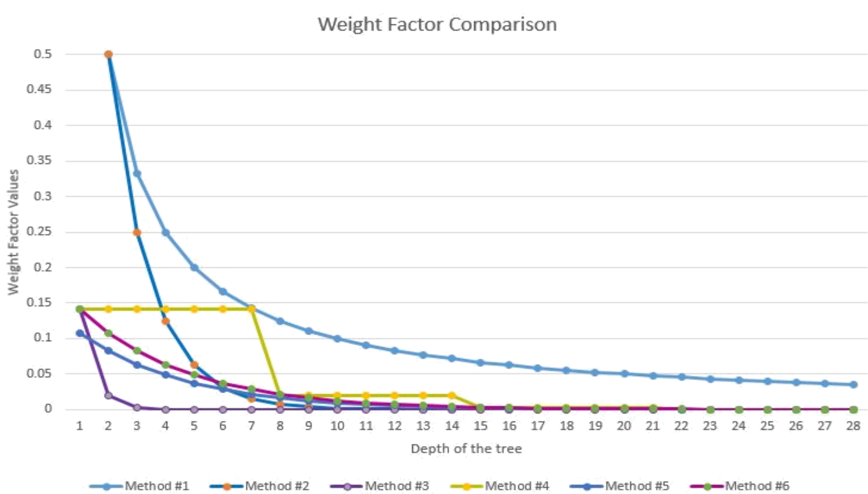

- Weight factors were calculated using several constraints, on which one approach was used to build the model. A few options were formulated for weight factor calculations and compared, as shown in the line graph

- Weight factors were applied to the recorded frequencies to obtain weighted feature frequencies for each unseen data sample.

Random forest tree explainer

Model Feature Prioritisation

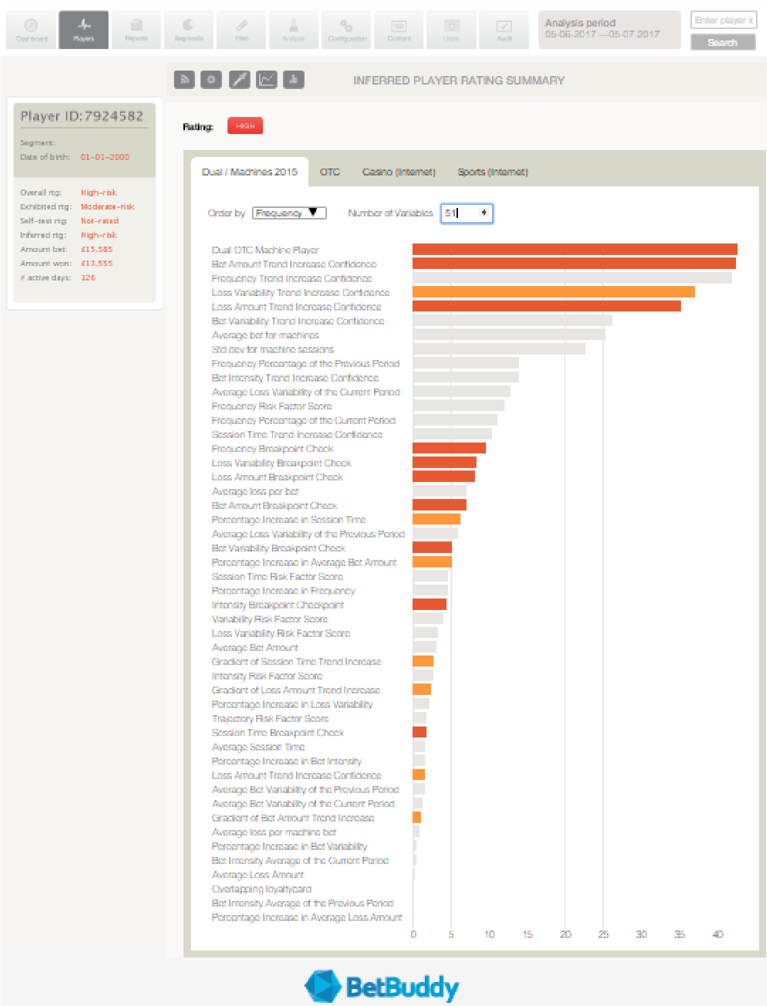

- Features are weighted at the sample level to build an individual knowledge graph for each player

Random forest tree explainer

Model Feature Prioritisation

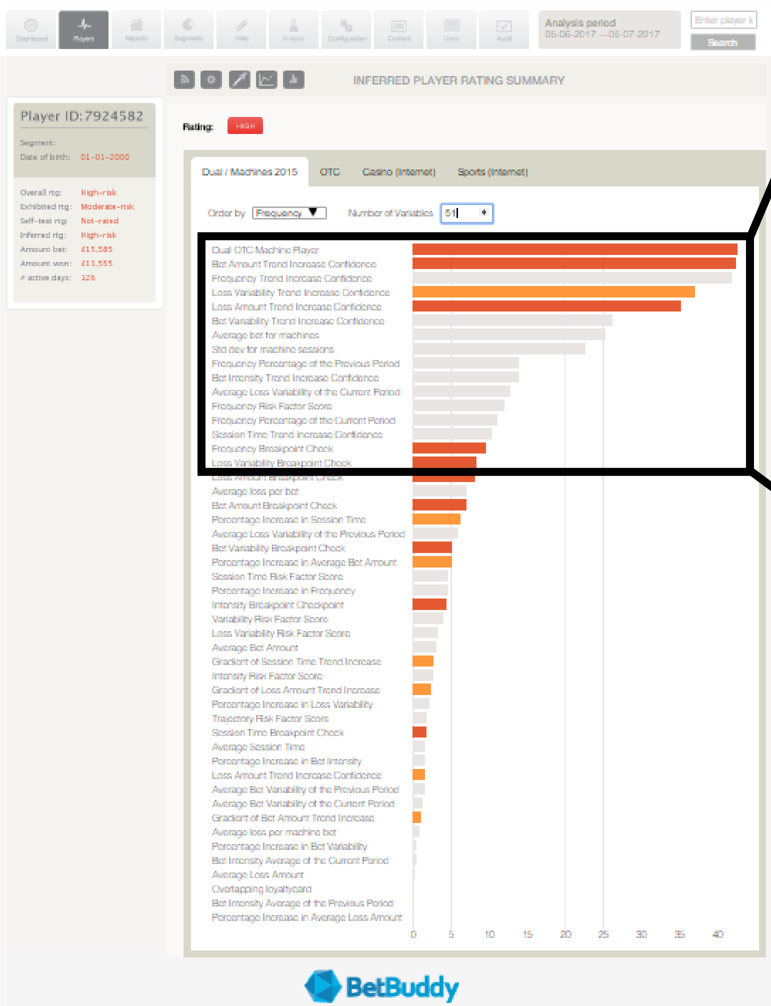

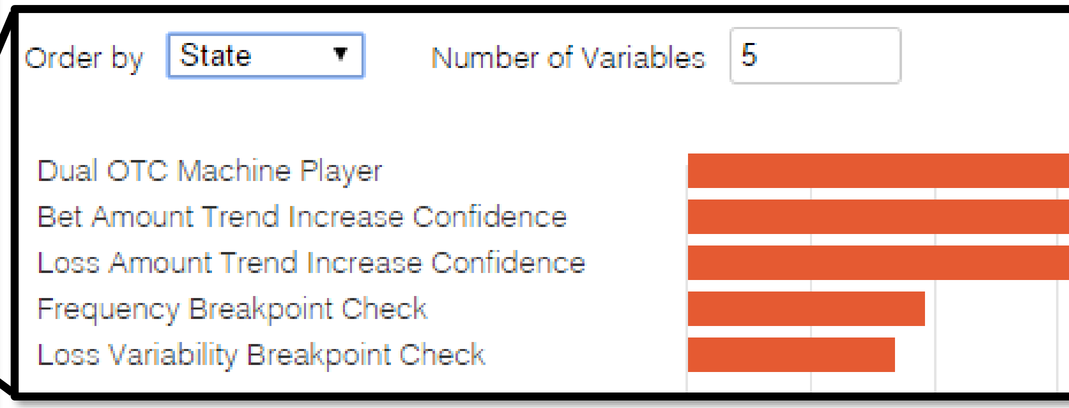

- Enables not only weighted feature frequency analysis at sample level but also categorisation of strength of feature value following distribution analysis in the training data sets
- Red / Amber classifications for targetted communications

Are all high risk players the same?

Player 1

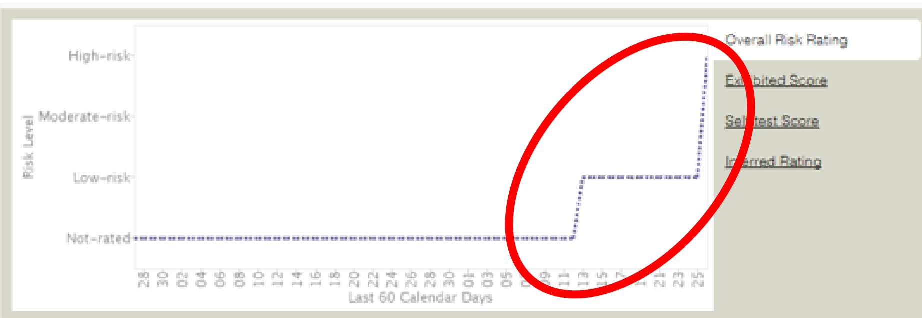

- New player (22 active gambling days)
- Risk profile increasing rapidly since opening account i.e., average daily bet amount of £743

Player 2

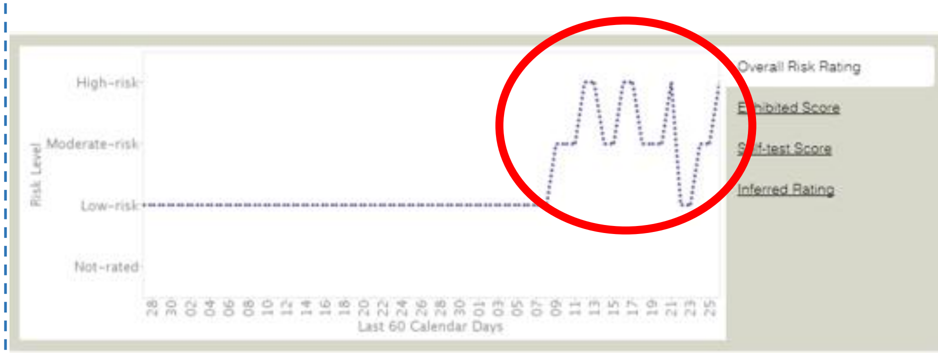

- Established player with 421 active gambling days
- Whilst risk profile has increased in last 2-3 weeks, appears to try to self-moderate
- Average daily bet amount of £249

Whilst both high on model prediction, they have different behavioural graphs

Risk Classification, Feature Weighting and Prioritisation

Player 1

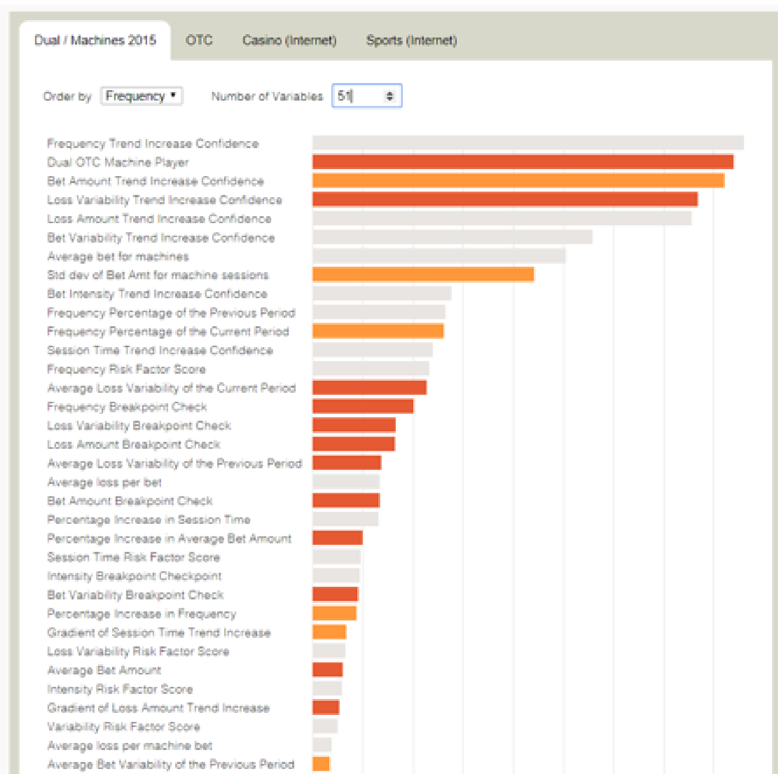

- Pattern-matching accuracy to other long-term self-excluders was 70%
- A large proportion of features ranked both Red/Amber and high on the tree interpreter

Player 2

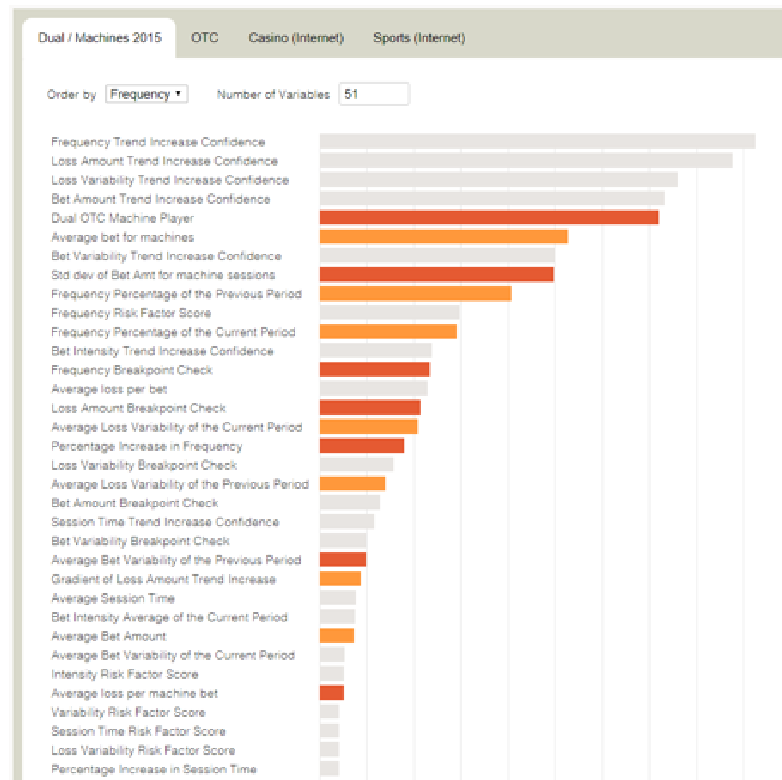

- Pattern-matching accuracy to other long-term self-excluders was 62%
- Reasonable proportion of features flagged Red/Amber, but only 1 in the top 5

Closer examination reveals distinct patterns driving risk

Player 1

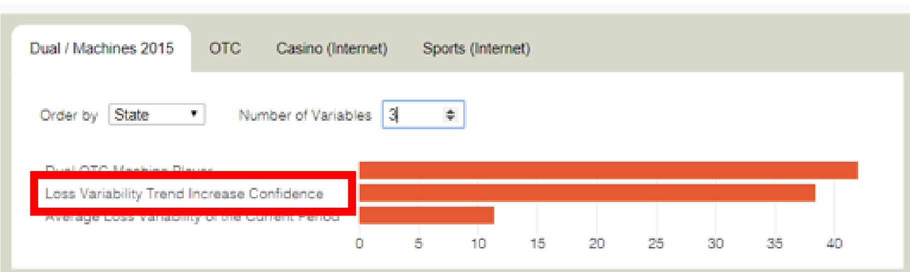

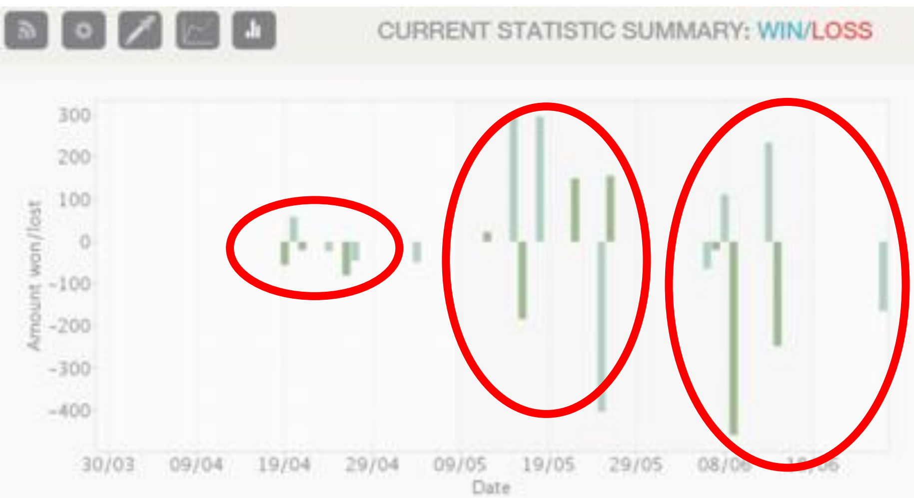

Highly variable loss patterns that are increasing

Closer examination reveals distinct patterns driving risk

Player 1

Player 2

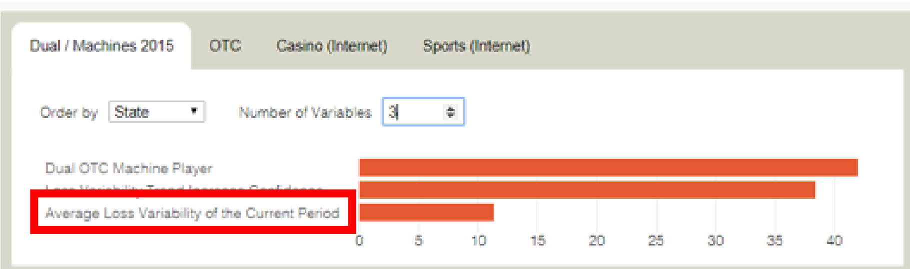

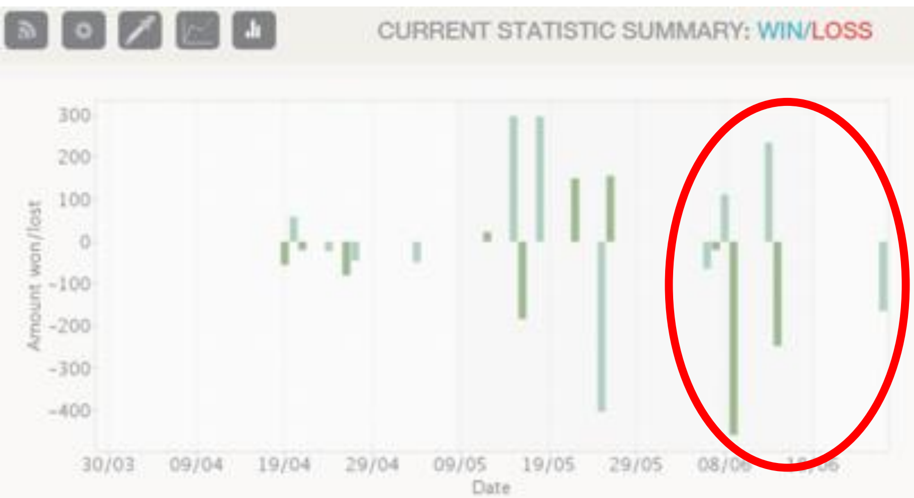

Recency in variable losses

Closer examination reveals distinct patterns driving risk

Player 1

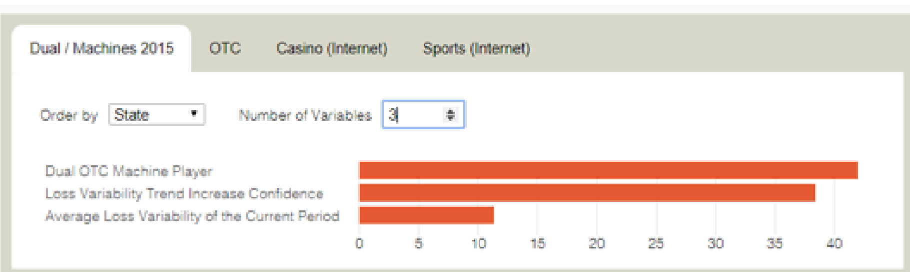

Player 2

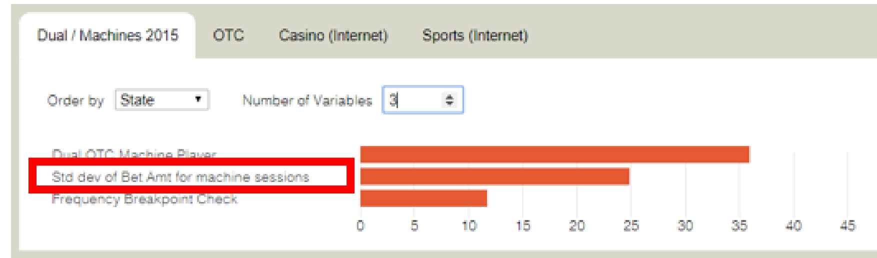

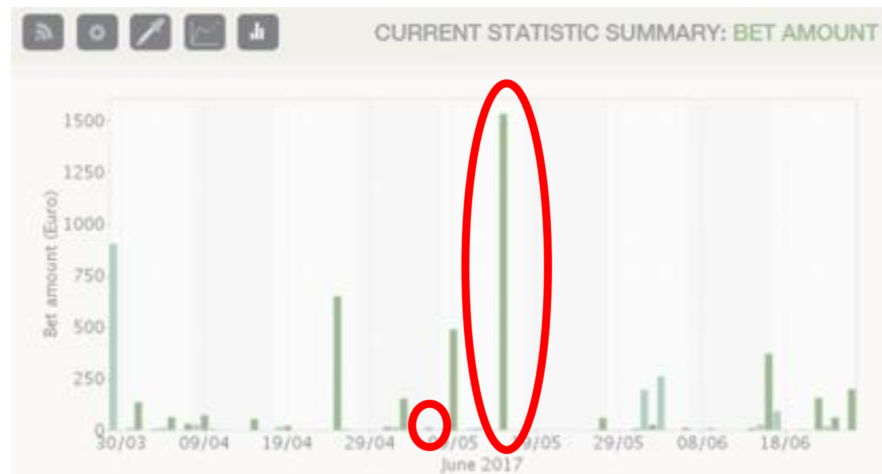

Highly variable staking patterns (not losses)

Closer examination reveals distinct patterns driving risk

Player 1

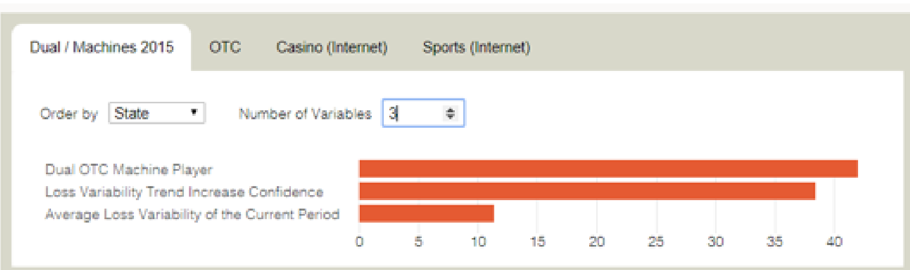

Player 2

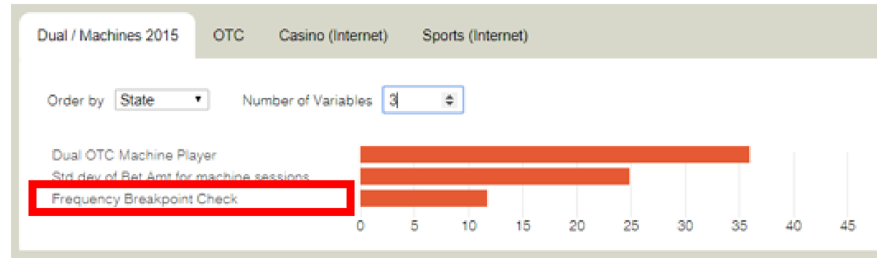

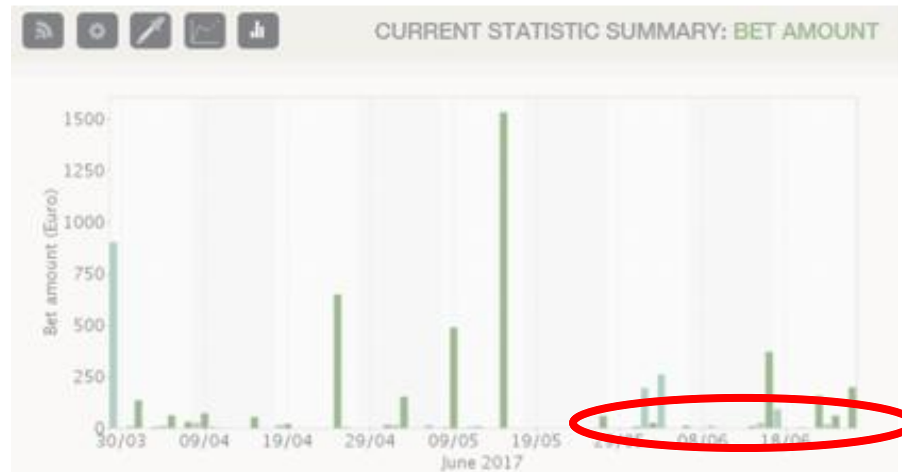

Frequency of play (every other day)

Should player 1 and 2 receive the same responsible gambling intervention?

Research and gambling data suggests personalisation may be beneficial

- Practical experience shows that non- problem gamblers (PGs) do not perceive RG tools as applicable for them, whilst PGs do not perceive themselves as PGs

- Problem gamblers (PGs) are also not a homogenous group therefore the types of communications sent may be perceived very differently by individuals

- Research shows personalized messages have shown to change behaviour in several areas such as smoking cessation, physical fitness, reduction in alcohol consumption, and diabetes management

- The term responsible gambling is often synonymous with problem gambling, which can be construed negatively by the customer

- Research into positive play behaviours suggests that players who play in a ‘positive manner’ exhibit different beliefs than those with a gambling problem. This points towards the potential for encouraging positive behaviours as well as discouraging high risk player behaviours.

Player Messaging Best Practice: UK Gov’t Behavioural Insights Team

Easy e.g.,

- Making messages clear often results in a significant increase in response rates – particularly breaking down complex goals

Attractive e.g.,

- People are more likely to do something there attention is drawn to – consider colour and images in messages
- People are drawn to incentives – are there incentives that can be designed into interventions?

Timely e.g.,

- The timing of interventions can have a significant impact – making them at a point close to when a behaviour is flagged could make them more relevant
- There is often a gap between intentions and actual behaviour – it may be useful to use further prompts if initial interventions do not result in a change

Social e.g.,

- We should be wary of reinforcing problematic behaviours by emphasising their high prevalence – specifically for normative comparisons.

Do Responsible Gambling Interactions have an effect on behaviour?

Do Responsible Gambling interactions (RGIs) achieve desired changes in behaviour? Early evaluation testing results from PlayOLG

- ~97,000 RGIs on PlayOLG during 2015/16 (viewed risk profile or self-test result)
- We wanted to know if RGIs have a positive impact on player behaviour
- Initially focused on a small subset of RGIs that met certain criteria e.g.,
    - active players
    - who used RGIs regularly (e.g., at least quarterly), and
    - who were rated either moderate or high risk at the time of the RGI
- Tested underlying behaviour (bet amount) in period leading up to the RGI and period after it.

Impact of RGIs on Behaviour (n = 1,455)

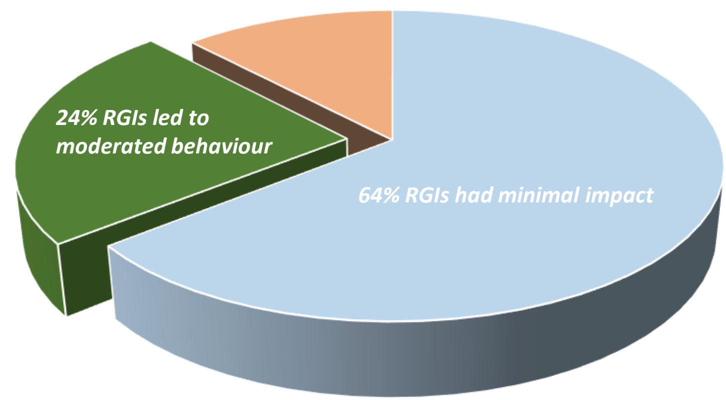

- 24% led to a moderation in underlying behaviour (i.e., improvement)
- 64% had minimal impact
- 12% led to an increase in underlying behaviour
- Whilst we cannot generalise, encouraging early indications RGIs can potentially help players.